

Simultaneous System Identification and Model Predictive Control with No Dynamic Regret

Hongyu Zhou Vasileios Tzoumas



Fig. 1: Simultaneous System Identification and Model Predictive Control: Tested Scenarios of Reference Trajectory Tracking subject to Various Unmodeled Aerodynamic Disturbances. In this paper, we focus on simultaneous system identification and model predictive control where the robots' capacity to achieve accurate and efficient tracking control is challenged by unknown both dynamics and (possibly state-dependent) disturbances. For example, the ability of a quadrotor to track a reference trajectory can be challenged by a variety of unknown aerodynamic disturbances: *ground effects* (left & right), *wind disturbances* (center & right), and *drag* (all). We provide a control algorithm that demonstrates an ability to simultaneously learn such unknown dynamics/disturbances in a self-supervised manner, based on the data collected on-the-go, and uses the learned models for predictive control. We prove that the algorithm guarantees bounded suboptimality against the optimal controller in hindsight, under technical assumptions that we describe later in the paper.

Abstract—We provide an algorithm for the simultaneous system identification and model predictive control of nonlinear systems. The algorithm has finite-time near-optimality guarantees and asymptotically converges to the optimal (non-causal) controller. Particularly, the algorithm enjoys sublinear *dynamic regret*, defined herein as the suboptimality against an optimal clairvoyant controller that knows how the unknown disturbances and system dynamics will adapt to its actions. The algorithm is self-supervised and applies to control-affine systems with unknown dynamics and disturbances that can be expressed in reproducing kernel Hilbert spaces. Such spaces can model external disturbances and modeling errors that can even be adaptive to the system's state and control input. For example, they can model wind and wave disturbances to aerial and marine vehicles, or inaccurate model parameters such as inertia of mechanical systems. We are motivated by the future of autonomy where robots will autonomously perform complex tasks despite real-world unknown disturbances such as wind gusts. The algorithm first generates random Fourier features that are used to approximate the unknown dynamics or disturbances. Then, it employs model predictive control based on the current learned model of the unknown dynamics (or disturbances). The model of the unknown dynamics is updated online using least squares based on the data collected while controlling the system. We validate our algorithm in both hardware experiments and physics-based simulations. The simulations include (i) a cart-pole aiming to maintain the pole upright despite inaccurate model parameters, and (ii) a quadrotor aiming to track reference trajectories despite unmodeled aerodynamic drag effects. The hardware experiments include a quadrotor aiming to track a circular trajectory despite unmodeled aerodynamic drag effects, ground effects, and wind disturbances. The code is open-sourced at <https://github.com/UM-iRaL/SSI-MPC>.

Manuscript received November 25, 2024; Revised March 13, 2025; Accepted April 21, 2025. This work was partially supported by NSF CAREER No. 2337412. This paper was recommended for publication by Editor-in-Chief Wolfram Burgard and Editor Jeannette Bohg upon evaluation of reviewers' comments. (Corresponding author: Hongyu Zhou.)

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109 USA; {zhouhy, vtzoumas}@umich.edu

Digital Object Identifier (DOI): see top of this page.

Index Terms—Online learning, adaptive model predictive control, regret optimization, random feature approximation.

I. INTRODUCTION

IN the future, mobile robots will automate fundamental tasks such as package delivery [1], target tracking [2], and inspection and maintenance [3]. Such tasks require accurate and efficient tracking control. But achieving accuracy and efficiency is challenging since such tasks often require the robots to operate under highly uncertain conditions, particularly, under unknown dynamics and (possibly state-dependent) disturbances. For example, they require the quadrotors to (i) pick up and carry packages of unknown weight, (ii) chase mobile targets at high speeds where the induced aerodynamic drag is hard to model, and (iii) inspect and maintain outdoor infrastructure exposed to turbulence and wind gusts.

State-of-the-art methods for control under unknown dynamics and disturbances typically rely on offline or online methods, or a mix of the two, including: robust control [4]–[10]; adaptive control and disturbance compensation [11]–[20]; offline learning for control [21]–[31]; offline learning for control with online adaption [32]–[38]; and online learning [39]–[51]. Among these, the offline learning methods require data collection for offline training, a process that can be expensive/time-consuming [52], instead of collecting data online and learning on the spot. For this reason, these methods may not also generalize to different scenarios than those used for offline learning [53], [54]. The robust control methods, given a known upper bound on the magnitude of the noise, can be conservative due to assuming worse-case disturbance realization [55], [56], instead of planning based on an accurate predictive model of the disturbance. Similarly, the adaptive control and the online learning control methods may exhibit sub-optimal performance due to only reacting to the history of observed disturbances, instead of planning based on an

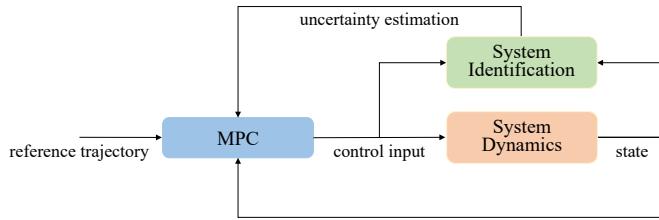


Fig. 2: Overview of Simultaneous System Identification and Model Predictive Control Pipeline. The pipeline is composed of two interacting modules: (i) a model predictive control (MPC) module, and (ii) an online system identification module. The MPC module uses the estimated unknown disturbances/dynamics from the system identification module to calculate the next control input. Given the control input and the observed new state, the online system identification module then updates the estimate of the unknown disturbances/dynamics.

accurate predictive model of the disturbance [11]–[13], [40], [49]. We discuss the related work in extension in Section VII.

In this paper, instead, we leverage the success of model predictive control methods for accurate tracking control [57]. To this end, we propose a self-supervised method to learn online a predictive model of the unknown uncertainties (Fig. 2). Therefore, the proposed method promises to enable: one-shot online learning, instead of offline (or episodic learning); online control that adapts to the actual disturbance realization, instead of the worst-case; and control planned over a look-ahead horizon of predicted system dynamics and disturbances, instead of their past. We elaborate on our contributions next.

Contributions. We provide a real-time and asymptotically-optimal algorithm for the simultaneous system identification and control of nonlinear systems. Specifically, the algorithm enjoys sublinear *dynamic regret*, defined herein as the suboptimality against an optimal clairvoyant controller that knows how the unknown disturbances and system dynamics will adapt to its actions. The said definition of dynamic regret differs from the standard definition of dynamic regret in the literature, e.g., [40]–[45] where the suboptimality of an Algorithm is measured against the controller that is optimal for the same realization of disturbances as the Algorithm experienced, instead of the realization that the optimal controller would experience given that the disturbances can be adaptive (Remark 3).

The algorithm is self-supervised and applies to control-affine systems with unknown dynamics and disturbances that can be expressed in reproducing kernel Hilbert spaces [58]. Such spaces can model external disturbances and modeling errors that are even adaptive to the system's state and control input. For example, they can model wind and wave disturbances to aerial and marine vehicles, or inaccurate model parameters such as inertia of mechanical systems.

Description of the algorithm: The algorithm is composed of two interacting modules (Fig. 2): (i) a model predictive control (MPC) module, and (ii) an online system identification module. At each time step, the MPC module uses the estimated unknown disturbances and system dynamics from the system identification module to calculate the next control input. Given the control input and the observed new state, the online system identification module then updates the estimate of the

unknown disturbances and system dynamics.

Approach of predictive online system identification: We enable the predictive online system identification module by employing online least-squares estimation via online gradient descent (OGD) [59], where the unknown disturbances/dynamics are parameterized as a linear combination of random Fourier features [15], [60], [61]. This allows us to maintain the computational efficiency of classical finite-dimensional parametric approximations that can be used in model predictive control while retaining the expressiveness of the RKHS with high probability. Using the finite-dimensional parametric approximation, the parameters are updated in a self-supervised manner using OGD at each time step. Such approximation then is used in MPC to select control input.

Performance guarantee: The algorithm asymptotically has *no dynamic regret*, that is, it asymptotically matches the performance of the optimal controller that knows a priori how the unknown disturbances and system dynamics will adapt to its actions. Particularly, we provide the following finite-time performance guarantee for Algorithm 1 (Theorem 1):

$$\text{dynamic regret} \leq \mathcal{O}\left(T^{\frac{3}{4}}\right).$$

Technical assumptions: The no-dynamic-regret guarantee holds true under the assumptions of bounds on MPC's cost function and value function, and Lipschitzness of MPC's cost function and unknown dynamics/disturbances, along with the previously stated assumption that the unknown disturbances and dynamics can be expressed in reproducing kernel Hilbert spaces. These assumptions are milder than the exponential stability or exponentially incremental stability assumption required in [14], [15], [62]–[64] and persistent excitation required in [62], [65].

Numerical Evaluations. We validate the algorithm in physics-based simulations and conduct an ablation study (Section V). Specifically, we validate our algorithm in simulated scenarios of (i) a cart-pole system that aims to stabilize at a setpoint despite inaccurate model parameters (Section V-B and Section V-C), and (ii) a quadrotor that aims to track reference trajectories subject to unmodeled aerodynamic forces (Section V-D). In the cart-pole experiments (Section V-B), We compare our algorithm with a nominal MPC (Nominal MPC) that ignores the unknown dynamics or disturbances, the non-stochastic MPC (NS-MPC) [45], and the Gaussian process MPC (GP-MPC) [24]. In the quadrotor experiment (Section V-D), we compare our algorithm with a Nominal MPC controller, and the Gaussian process MPC (GP-MPC) [23]. In the simulations, the algorithm achieves real-time control and superior tracking performance over the compared methods, demonstrating its ability to perform simultaneous online model learning and control. We also show that the performance of Algorithm 1 can be further improved by combining it with an incremental nonlinear dynamic inversion (INDI) inner-loop controller [16] in Section V-D. In addition, we conduct sensitivity analysis in Section V-C on the tuning parameters of Algorithm 1: how the number of random Fourier features and learning rate of OGD affect the performance of Algorithm 1.

Hardware Experiments. We implement the provided algorithm in a quadrotor (Section VI and Appendix I). In Section VI, we consider that the quadrotor is tasked to track a circular trajectory with diameter 1m at maximal speed

$v_m = 0.8 \text{ m/s}$ under (i) ground effects (Figure 1, left), (ii) wind disturbances (Figure 1, center), (iii) both ground effects and wind disturbances (Figure 1, right). In Appendix I, we consider the quadrotor is tasked to track the same circular trajectory with maximal speed either 0.8 m/s or 1.3 m/s , without ground effects and wind disturbances. In such cases, the quadrotor experiences unknown aerodynamic effects including body drag, rotor drag, and turbulent effects caused by the propellers. Across all hardware experiments, we compare Algorithm 1 with a nominal MPC (Nominal MPC), and the L1-adaptive MPC (L1-MPC) [17]. The algorithm achieves superior tracking performance over the compared methods under all tested scenarios, demonstrating its ability to perform simultaneous online system identification and predictive control.

II. SIMULTANEOUS SYSTEM IDENTIFICATION AND MODEL PREDICTIVE CONTROL

We formulate the problem of *Simultaneous System Identification and Model Predictive Control* (Problem 1). To this end, we use the following framework and assumptions.

Control-Affine Dynamics. We consider control-affine system dynamics of the form

$$x_{t+1} = f(x_t) + g(x_t)u_t + h(z_t), \quad t \geq 1, \quad (1)$$

where $x_t \in \mathbb{R}^{d_x}$ is the state, $u_t \in \mathbb{R}^{d_u}$ is the control input, $f : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ and $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x} \times \mathbb{R}^{d_u}$ are known locally Lipschitz functions, $h : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ is an unknown locally Lipschitz function, and $z_t \in \mathbb{R}^{d_z}$ is a vector of features chosen as a subset of $[x_t^\top u_t^\top]^\top$.¹ Particularly, $h(\cdot)$ represents unknown disturbances or system dynamics.

We refer to the undisturbed system dynamics $x_{t+1} = f(x_t) + g(x_t)u_t$ as the *nominal dynamics*.

Remark 1 (Extension to Systems with Unmodeled Process Noise). *In the Appendix, we extend the results in this paper to systems corrupted with unknown (potentially non-stochastic) noise that is additive to the system dynamics, i.e.,*

$$x_{t+1} = f(x_t) + g(x_t)u_t + h(z_t) + w_t, \quad (2)$$

where w_t represents, e.g., process noise.

The unknown disturbance (or system dynamics) $h(\cdot)$ may be adapting to the state and the control input. Assumptions on $h(\cdot)$ will be introduced in Section III-A.

Model Predictive Control (MPC). MPC selects a control input u_t by simulating the nominal system dynamics over a look-ahead horizon N , i.e., MPC selects u_t by solving the optimization problem [66], [67]:

$$\min_{u_t, \dots, u_{t+N-1}} \sum_{k=t}^{t+N-1} c_k(x_k, u_k) \quad (3a)$$

$$\text{subject to} \quad x_{k+1} = f(x_k) + g(x_k)u_k, \quad (3b)$$

$$u_k \in \mathcal{U}, \quad k \in \{t, \dots, t+N-1\}, \quad (3c)$$

¹The feature vector z_t can also include, for example, t [15, Remark 3.1] and $h(z_{t-1})$. We can also model time delay by using $x_{t+1} = f(x_t) + g(x_t)u_{t-\tau} + \tilde{h}(\tilde{z}_t) = f(x_t) + g(x_t)u_t + (g(x_t)u_{t-\tau} - g(x_t)u_t + \tilde{h}(\tilde{z}_t))$, where τ represents the latency. We can define $h(z_t) \triangleq g(x_t)u_{t-\tau} - g(x_t)u_t + \tilde{h}(\tilde{z}_t)$ with z_t including $x_t, u_t, x_{t-\tau}, u_{t-\tau}, \tilde{z}_t$.

where $c_t(\cdot, \cdot) : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}$ is the cost function, and \mathcal{U} is a compact set that represents constraints on the control input due to, e.g., controller saturation. We discuss how to incorporate state constraints in Appendix H.

The optimization problem in eq. (3) ignores the unknown disturbances $h(\cdot)$. To improve performance in the presence of $h(\cdot)$, in this paper we propose a method to estimate $h(\cdot)$ online so eq. (3) can be adapted to the optimization problem:

$$\min_{u_t, \dots, u_{t+N-1}} \sum_{k=t}^{t+N-1} c_k(x_k, u_k) \quad (4a)$$

$$\text{subject to} \quad x_{k+1} = f(x_k) + g(x_k)u_k + \hat{h}(z_k), \quad (4b)$$

$$u_k \in \mathcal{U}, \quad k \in \{t, \dots, t+N-1\}, \quad (4c)$$

where $\hat{h}(\cdot)$ is the estimate of $h(\cdot)$. Specifically, $\hat{h}(\cdot) \triangleq \hat{h}(\cdot; \hat{\alpha})$ where $\hat{\alpha}$ is a parameter that is updated online by our proposed method to improve the control performance.

We define the notion of MPC's *value function* and state the assumption on the cost function and value function.

Definition 1 (Value Function [68]). *Given a state x and parameter $\hat{\alpha}$, the value function $V_t(x; \hat{\alpha})$ is defined as the optimal value of eq. (5):*

$$\min_{u_t, \dots, u_{t+N-1}} \sum_{k=t}^{t+N-1} c_k(x_k, u_k) \quad (5a)$$

$$\text{subject to} \quad x_{k+1} = f(x_k) + g(x_k)u_k + \hat{h}(z_k), \quad (5b)$$

$$x_t = x, \quad u_k \in \mathcal{U}, \quad k \in \{t, \dots, t+N-1\}. \quad (5c)$$

Assumption 1 (Bounds on Cost Function and Value Function [68]). *For all $\hat{\alpha}$ in a compact set \mathcal{D} , there exist positive scalars $\underline{\lambda}$, $\bar{\lambda}$, and a continuous function $\sigma : \mathbb{R}^{d_x} \rightarrow \mathbb{R}_+$, such that (i) $c_t(x, u) \geq \underline{\lambda}\sigma(x), \forall x, u, t$; (ii) $V_t(x; \hat{\alpha}) \leq \bar{\lambda}\sigma(x), \forall x, t$, and (iii) $\lim_{\|x\| \rightarrow \infty} \sigma(x) \rightarrow \infty$.*

Under Assumption 1, the MPC policy in eq. (4) can be proved to ensure that the system in eq. (4b) is globally asymptotically stable [68].

Assumption 1 is milder than the exponential stability or exponentially incrementally stability assumption assumed, e.g., in [14], [15], [62]–[64]. It is common to assume bound conditions or stability for the nominal/estimated systems to facilitate the regret analysis [14], [40]–[45], [62]–[65], [69] or MPC stability analysis [15], [70]–[72].

Remark 2 (Discussion on Assumption 1). *Consider the quadratic cost $c_t(x_t, u_t) = x_t^\top Q x_t + u_t^\top R u_t$. For linear systems, Assumption 1 is satisfied [68, Lemma 1]. For nonlinear systems, Assumption 1 is satisfied when the quadratic cost is (exponentially/asymptotically) controllable to zero with respect to $\sigma : \mathbb{R}^{d_x} \rightarrow \mathbb{R}_+$ [68, Section. III]. Particularly, for the nonlinear case: a) when $R = 0$, the (exponentially/asymptotically) controllability holds if MPC can drive $x \rightarrow 0$; b) for $R \neq 0$, it is non-trivial to verify Assumption 1 and a necessary condition is $(x, u) = (0, 0)$ being an equilibrium point or $R \triangleq R(x) = 0$ when $x = 0$ [68].*

Assumption 2 (Lipschitzness). *We Assume that $c_t(x, u)$ is locally Lipschitz in x and u , $\hat{h}(\cdot)$ is locally Lipschitz in $\hat{\alpha}$.*

Assumption 2 will be used to establish the Lipschitzness of the value function $V_t(x; \hat{\alpha})$ with respect to the initial state x and parameter $\hat{\alpha}$.

Control Performance Metric. We design u_t to ensure a control performance that is comparable to an optimal clairvoyant (non-causal) policy that knows the disturbance function h a priori. Particularly, we consider the metric below.

Definition 2 (Dynamic Regret). *Assume a total time horizon of operation T , and loss functions c_t , $t = 1, \dots, T$. Then, dynamic regret is defined as*

$$\text{Regret}_T^D = \sum_{t=1}^T c_t(x_t, u_t, h(z_t)) - \sum_{t=1}^T c_t(x_t^*, u_t^*, h(z_t^*)), \quad (6)$$

where we made the dependence of the cost c_t to the unknown disturbance h explicit, u_t^* is the optimal control input in hindsight, i.e., the optimal (non-causal) input given a priori knowledge of the unknown function h and x_{t+1}^* is the state reached by applying the optimal control inputs (u_1^*, \dots, u_T^*) .

Remark 3 (Adaptivity of h). *In the definition of regret in eq. (6), h adapts (possibly differently) to the state and control sequences $(x_1, u_1), \dots, (x_T, u_T)$ and $(x_1^*, u_1^*), \dots, (x_T^*, u_T^*)$ since h is a function of the state and the control input. In particular, the proposed control algorithm selects control input using eq. (4) with estimated \hat{h} , and the optimal controller with the knowledge of h and (potentially) with a look-ahead horizon different than N . Both controllers then evolve to the next state experiencing the real disturbances, defined by $h(\cdot)$. Since $h(\cdot)$ is a function of the state and control input, the proposed controller witnesses $h(z_t)$, while the optimal controller witnesses $h(z_t^*)$. This is in contrast to previous definitions of dynamic regret, e.g., [40]–[45] and references therein, where the optimal state x_{t+1}^* is reached given the same realization of h as of x_{t+1} , i.e., $x_{t+1}^* = f(x_t^*) + g(x_t^*) u_t^* + h(z_t)$.*

Problem 1 (Simultaneous System Identification and Model Predictive Control (SSI-MPC)). *Suppose x_1 is known and x_t can be measured. At each $t = 1, \dots, T$, estimate the unknown disturbance $\hat{h}(\cdot)$, and identify a control input u_t by solving eq. (4), such that Regret_T^D is sublinear.*

A sublinear dynamics regret means $\lim_{T \rightarrow \infty} \text{Regret}_T^D / T \rightarrow 0$, which implies the algorithm asymptotically converges to the optimal (non-causal) controller.

III. ALGORITHM

We present the algorithm for Problem 1 (Algorithm 1). The algorithm is sketched in Figure 2. The algorithm is composed of two interacting modules: (i) an MPC module, and (ii) an online system identification module. At each $t = 1, 2, \dots$, the MPC module uses the estimated $\hat{h}(\cdot)$ from the system identification module to calculate the control input u_t . Given the current control input u_t and the observed new state x_{t+1} , the online system identification module updates the estimate $\hat{h}(\cdot)$. To this end, it employs online least-squares estimation via online gradient descent, where $h(\cdot)$ is parameterized as a linear combination of random Fourier features.

To rigorously present the algorithm, we thus first provide background information on random Fourier features for approximating an $h(\cdot)$ (Section III-A), and on online gradient descent for estimation (Section III-B).

A. Function Approximation via Random Fourier Features

We overview the randomized approximation algorithm introduced in [15] for approximating an $h(\cdot)$ under Assumption 3 and Assumption 4. The algorithm is based on random Fourier features [60], [61] and their extension to vector-valued functions [73], [74]. By being randomized, the algorithm is computationally efficient while retaining the expressiveness of the RKHS with high probability.

We assume the following for the unknown disturbances or system dynamics $h(\cdot)$.

Assumption 3 (Function Space of h [75]). *$h : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x}$ lies in a subspace of a Reproducing Kernel Hilbert Space (RKHS) \mathcal{H} where the kernel K is considered known [76] and can be written via a feature map $\Phi : \mathbb{R}^{d_z} \times \Theta \rightarrow \mathbb{R}^{d_x \times d_1}$ as*

$$K(z_1, z_2) = \int_{\Theta} \Phi(z_1, \theta) \Phi(z_2, \theta)^T d\nu(\theta), \quad (7)$$

where $d_1 \leq d_x$, ν is a known probability measure on a measurable space Θ .

The subspace defined in Assumption 3 will be rigorously defined later in eq. (12).

Assumption 4 (Operator-Valued Bochner's Theorem [73]). *The measurable space Θ is a subset of \mathbb{R}^{d_z+1} such that $\theta \in \Theta$ can be written as $\theta = (w, b)$, where $w \in \mathbb{R}^{d_z}$ and $b \in \mathbb{R}$. Also, the feature map can be factorized as $\Phi(z, \theta) = B(w)\phi(w^T z + b)$, where $B : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_x \times d_1}$ and $\phi : \mathbb{R} \rightarrow [-1, 1]$ is a 1-Lipschitz function.*

Examples of kernels that satisfy Assumption 3 and Assumption 4 include the Gaussian and Laplace kernels. Additional examples of kernels can be found in [77, Section 5.2].

Under Assumption 3, a function h can be written as [75]

$$h(\cdot) = \int_{\Theta} \Phi(\cdot, \theta) \alpha(\theta) d\nu(\theta), \quad (8)$$

with $\|h\|_{\mathcal{H}}^2 = \|\alpha\|_{L_2(\Theta, \nu)}^2 \triangleq \int_{\Theta} \|\alpha(\theta)\|^2 d\nu(\theta)$, where $\alpha : \Theta \rightarrow \mathbb{R}^{d_1}$ is a square-integrable signed density. The corresponding Hilbert space where h satisfies eq. (8) is referred to as \mathcal{F}_2 [75], [78].

Equation (8) implies that $h(\cdot)$ is an integral of $\Phi(\cdot, \theta) \alpha(\theta)$ over the base measure ν , thus, we can obtain a finite-dimensional approximation of $h(\cdot)$ by

$$h(\cdot) \approx \hat{h}(\cdot; \alpha) \triangleq \frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i, \quad (9)$$

where $\theta_i \sim \nu$ are drawn i.i.d. from the base measure ν , $\alpha_i \triangleq \alpha(\theta_i)$ are parameters to be learned, and M is the number of sampling points that decides the approximation accuracy.

To rigorously establish the relationship between the number of random features M and the accuracy of uniformly approximation of $h(\cdot)$, we first define $B_{\Phi}(\delta)$ as any function on a fixed compact set $\mathcal{Z} \subset \mathbb{R}^{d_z}$ such that, for any $\delta \in (0, 1)$,

$$\mathbb{P}_{\theta \sim \nu} \left(\sup_{z \in \mathcal{Z}} \|\Phi(z, \theta)\|_{\text{op}} > B_{\Phi}(\delta) \right) \leq \delta. \quad (10)$$

Then, we define a truncated Φ for any $\mu \in (0, 1)$ as

$$\Phi_\mu(z, \theta) \triangleq \Phi(z, \theta) \mathbf{1}\{\|\Phi(z, \theta)\|_{\text{op}} \leq B_\Phi(\mu)\}. \quad (11)$$

We aim to approximate $h(\cdot)$ over a subset of \mathcal{F}_2 , particularly,

$$\mathcal{F}_2(B_h) \triangleq \left\{ h(\cdot) = \int_{\Theta} \Phi(\cdot, \theta) \alpha(\theta) d\nu(\theta) \middle| \alpha \in \mathcal{D}(B_h) \right\}, \quad (12)$$

where $\mathcal{D}(B_h) \triangleq \{\alpha(\theta) \mid \text{ess sup}_{\theta \in \Theta} \|\alpha(\theta)\| \leq B_h\}$. For simplicity, we will drop the dependence of \mathcal{D} on B_h .

Under Assumption 3 and Assumption 4, [15] extends the approximation theory of [61] to vector-valued functions:

Proposition 1 (Uniformly Approximation Error [15]). *Assume $h \in \mathcal{F}_2(B_h)$. Let $\delta \in (0, 1)$ and $\mu = \frac{\delta}{2M}$. With probability at least $1 - \delta$, there exist $\{\alpha_i\}_{i=1}^M \in \mathcal{D}$ such that*

$$\begin{aligned} & \left\| h(\cdot) - \frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i \right\|_\infty \\ & \leq \frac{B_h}{\sqrt{M}} \left[2B_\Phi \left(\frac{\delta}{2M} \right) \left(2B_Z \sqrt{\mathbb{E} \|w_i\|^2} \right. \right. \\ & \quad \left. \left. + 2\sqrt{d_1} + \sqrt{\log \frac{2}{\delta}} + \sqrt{\frac{\delta}{2} \mathbb{E} \|B(w)\|_{\text{op}}^2} \right) \right] \end{aligned} \quad (13)$$

where $B_Z \triangleq \sup_{z \in \mathcal{Z}} \|z\|$.

Proposition 1, therefore, indicates that the uniformly approximation error scales $\mathcal{O}\left(\frac{1}{\sqrt{M}}\right)$.

Assumption 5 (Small Uniformly Approximation Error). *For a chosen M value, we assume that the uniformly approximation error in Proposition 1 is negligible, i.e., $h(\cdot)$ can be expressed as $\frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i$, where $\alpha_i \in \mathcal{D}$.*

Remark 4 (Beyond Random Fourier Features). *Beyond random Fourier features, h can be learned by any models given that (i) $\hat{h}(\cdot; \alpha)$ is convex in α , (ii) Assumption 2 and Assumption 5 hold, and (iii) $\hat{h}(\cdot; \alpha)$ and $\nabla_\alpha \hat{h}$ are uniformly bounded for all $\alpha \in \mathcal{D}$.² This includes, but not limited to:*

- trained neural network model as feature map [36], [37], where $h(z_t) \triangleq \Phi(z_t, \theta) \alpha$, with $\Phi(\cdot, \theta)$ the uniformly bounded neural network model (e.g., using tanh function or sigmoid function as the last layer activation function), θ the parameter of neural network learned offline, and α the parameter to be updated online.
- Koopman model [19], where $\Phi(h(z_t)) \triangleq F\Phi(h(z_{t-1})) + H\Psi(h(z_{t-1}), z_t)$ and $h(z_t) \triangleq C\Phi(h(z_t))$ with $\Phi(\cdot)$ locally Lipschitz nonlinear function, $\Psi(\cdot, \cdot)$ uniformly bounded nonlinear function, C a constant matrix that maps $\Phi(h(\cdot))$ back to $h(\cdot)$, $h(\cdot)$ uniformly bounded unknown disturbances, and F and H the parameters to be learned.

The resulting performance of MPC depends on the specific models, e.g., their expressiveness and generalization to different disturbances, and the computational complexity.

²The finite-dimensional approximation of $h(\cdot)$ using random Fourier features satisfies the third condition per Assumption 4.

B. Online Least-Squares Estimation

Given a data point $(z_t, h(z_t))$ observed at time t , we employ an online least-squares algorithm that updates the parameters $\hat{\alpha}_t \triangleq [\alpha_{i,t}^\top, \dots, \alpha_{M,t}^\top]^\top$ to minimize the approximation error $l_t = \|h(z_t) - \hat{h}(z_t)\|^2$, where $\hat{h}(\cdot) \triangleq \frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \hat{\alpha}_{i,t}$ and $\Phi(\cdot, \theta_i)$ is the random Fourier feature as in Section III-A. Specifically, the algorithm used the online gradient descent algorithm (OGD) [59]. At each $t = 1, \dots, T$, it makes the steps:

- Given $(z_t, h(z_t))$, formulate the estimation loss function (approximation error):

$$l_t(\hat{\alpha}_t) \triangleq \left\| h(z_t) - \frac{1}{M} \sum_{i=1}^M \Phi(z_t, \theta_i) \hat{\alpha}_{i,t} \right\|^2.$$

- Calculate the gradient of $l_t(\hat{\alpha}_t)$ with respect to $\hat{\alpha}_t$:

$$\nabla_t \triangleq \nabla_{\hat{\alpha}_t} l_t(\hat{\alpha}_t).$$

- Update using gradient descent with learning rate η :

$$\hat{\alpha}'_{t+1} = \hat{\alpha}_t - \eta \nabla_t.$$

- Project each $\hat{\alpha}'_{i,t+1}$ onto \mathcal{D} :

$$\hat{\alpha}_{i,t+1} = \Pi_{\mathcal{D}}(\hat{\alpha}'_{i,t+1}) \triangleq \underset{\alpha \in \mathcal{D}}{\text{argmin}} \|\alpha - \hat{\alpha}'_{i,t+1}\|^2.$$

The above online least-squares estimation algorithm enjoys an $\mathcal{O}(\sqrt{T})$ regret bound, per the regret bound of OGD [59].

Proposition 2 (Regret Bound of Online Least-Squares Estimation [59]). *Assume $\eta = \mathcal{O}(1/\sqrt{T})$. Then,*

$$\text{Regret}_T^S \triangleq \sum_{t=1}^T l_t(\alpha_t) - \sum_{t=1}^T l_t(\alpha^*) \leq \mathcal{O}(\sqrt{T}), \quad (14)$$

where $\alpha^* \triangleq \underset{\alpha \in \mathcal{D}}{\text{argmin}} \sum_{t=1}^T l_t(\alpha)$ is the optimal parameter that achieves lowest cumulative loss in hindsight.

The online least-squares estimation algorithm thus asymptotically achieves the same estimation error as the optimal parameter α^* since $\lim_{T \rightarrow \infty} \text{Regret}_T^S/T = 0$.

Remark 5 (Time-Varying Optimal Parameters). *While we focus on the case of time-invariant optimal parameter α^* , our analysis naturally generalizes to the case of time-variant optimal parameter $(\alpha_1^*, \dots, \alpha_T^*)$, by using dynamic regret analysis of online least-square estimation [79].*

C. Algorithm for Problem 1

We describe the algorithm for SSI-MPC. The pseudo-code is in Algorithm 1. The algorithm is composed of three steps, initialization, control, and estimation, where the control and estimation steps influence each other at each time steps (Fig. 2):

- *Initialization steps:* Algorithm 1 first initializes the system state x_1 and parameter $\hat{\alpha}_1 \in \mathcal{D}$ (line 1). Then given the number of random Fourier features, Algorithm 1 randomly samples θ_i and formulates $\Phi(\cdot, \theta_i)$, where $i \in \{1, \dots, M\}$ (line 2).

Algorithm 1: Simultaneous System Identification and Model Predictive Control.

Input: Number of random Fourier features M ; base measure ν ; domain set \mathcal{D} ; gradient descent learning rate η .
Output: At each time step $t = 1, \dots, T$, control input u_t .

- 1: Initialize $x_1, \hat{\alpha}_{i,1} \in \mathcal{D}$;
- 2: Randomly sample $\theta_i \sim \nu$ and formulate $\Phi(\cdot, \theta_i)$, where $i \in \{1, \dots, M\}$;
- 3: **for** each time step $t = 1, \dots, T$ **do**
- 4: Apply control input u_t by solving eq. (4) with $\hat{h}(\cdot) \triangleq \frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \hat{\alpha}_{i,t}$;
- 5: Observe state x_{t+1} , and calculate disturbance via $h(z_t) = x_{t+1} - f(x_t) - g(x_t)u_t$;
- 6: Formulate estimation loss $l_t(\hat{\alpha}_t) \triangleq \|h(z_t) - \frac{1}{M} \sum_{i=1}^M \Phi(z_t, \theta_i) \hat{\alpha}_{i,t}\|^2$;
- 7: Calculate gradient $\nabla_t \triangleq \nabla_{\hat{\alpha}_t} l_t(\hat{\alpha}_t)$;
- 8: Update $\hat{\alpha}'_{t+1} = \hat{\alpha}_t - \eta \nabla_t$;
- 9: Project $\hat{\alpha}'_{i,t+1}$ onto \mathcal{D} , i.e., $\hat{\alpha}_{i,t+1} = \Pi_{\mathcal{D}}(\hat{\alpha}'_{i,t+1})$, for $i \in \{1, \dots, M\}$;
- 10: **end for**

- **Control steps:** Then, at each t , given the current estimate $\hat{h}(\cdot) \triangleq \frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \hat{\alpha}_{i,t}$, Algorithm 1 applies the control inputs u_t obtained by solving eq. (4) (line 4).
- **Estimation steps:** The system then evolves to state x_{t+1} , and, $h(z_t)$ is calculated upon observing x_{t+1} (line 5). Afterwards, the algorithm formulates the loss $l_t(\hat{\alpha}_t) \triangleq \|h(z_t) - \sum_{i=1}^M \Phi(z_t, \theta_i) \hat{\alpha}_{i,t}\|^2$ (line 6), and calculates the gradient $\nabla_t \triangleq \nabla_{\hat{\alpha}_t} l_t(\hat{\alpha}_t)$ (line 7). Algorithm 1 then updates the parameter $\hat{\alpha}_t$ to $\hat{\alpha}'_{t+1}$ (line 8) and, finally, projects each $\hat{\alpha}'_{i,t+1}$ back to the domain set \mathcal{D} (line 9).

IV. NO-REGRET GUARANTEE

We present the sublinear regret bound of Algorithm 1.

Theorem 1 (No-Regret). *Assume Algorithm 1's learning rate is $\eta = \mathcal{O}(1/\sqrt{T})$. Then, Algorithm 1 achieves*

$$\text{Regret}_T^D \leq \mathcal{O}\left(T^{\frac{3}{4}}\right). \quad (15)$$

Theorem 1 serves as a finite-time performance guarantee as well as implies that Algorithm 1 converges to the optimal (non-causal) control policy since $\lim_{T \rightarrow \infty} \text{Regret}_T^D/T \rightarrow 0$.

We have no proof that the bound in eq. (15) is tight. We will consider this analysis in our future work. Notably, the bound matches that of the adaptive control method in [14, Corollary 5.7]. However, in contrast to the results herein:

- The control performance metric in [14] and, thus, the definition of dynamic regret, is only a function of the state. In contrast, in this paper, the performance control metric and dynamic regret depend on both the state and the control input. Therefore, Theorem 1 provides a performance guarantee over a broader state and control optimization setting.
- The method in [14] is based on adaptive control instead of model predictive control.
- The method in [14] assumes exponentially incrementally stability. In contrast, we only require Assumption 1.

V. NUMERICAL EVALUATIONS

We evaluate Algorithm 1 in extensive simulated scenarios of control under uncertainty, where the controller aims to track

a reference setpoint/trajectory despite unknown disturbance. We first detail how Algorithm 1 is implemented empirically in Section V-A. For the experiments, we first consider a cart-pole aiming to stabilize around a setpoint despite inaccurate model parameters, i.e., inaccurate cart mass, pole mass, and pole length (Section V-B). Then, we consider the same setup of cart-pole and conduct the parameter sensitivity analysis, i.e., how the number of random Fourier features and the learning rate affect the performance of Algorithm 1 (Section V-C). We consider a quadrotor aiming to track given reference trajectories subject to unknown aerodynamic forces (Section V-D).

A. Empirical Implementation of Algorithm 1

We employ Algorithm 1 as follows:

- We use the quadratic cost function in MPC since it is widely applied and intuitive to tune, despite the difficulty of verifying if Assumption 1 holds.
- We assume $h(\cdot)$ can be fitted via $\frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i$, per Assumption 5.
- We choose the kernel K to be the Gaussian kernel. For its randomized approximation, we obtain $\{\theta_i\}_{i=1}^M$ i.i.d. by sampling w_i from a Gaussian distribution, and b_i from an uniform distribution from $[0, 2\pi]$ [60]. Unless specified, we sample w_i from a standard Gaussian distribution.
- We use different sets of parameters for each entry of $h(\cdot)$: Suppose $h(\cdot) = [h_1(\cdot), h_2(\cdot)]^\top \in \mathbb{R}^2$, then we approximate $h_1(\cdot)$ and $h_2(\cdot)$ by $\frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i^{(1)}$ and $\frac{1}{M} \sum_{i=1}^M \Phi(\cdot, \theta_i) \alpha_i^{(2)}$, respectively; i.e., $h_1(\cdot)$ and $h_2(\cdot)$ share the same set of random Fourier features but have different parameters α . We make this design choice to reduce the feature map to $\Phi(z, \theta) = \phi(w^\top z + b)$ and avoid the tuning of $B(w)$.

B. Cart-Pole Scenario

Simulation Setup. We consider a cart-pole system, where a cart of mass m_c connects via a prismatic joint to a 1D track, while a pole of mass m_p and length $2l$ is hinged to the cart. The state vector \mathbf{x} includes the horizontal position of the cart x , the velocity of the cart \dot{x} , the angle of the pole with respect to vertical θ , and the angular velocity of the pole $\dot{\theta}$. The control input is the force F applied to the center of mass of the cart. The goal of the cart-pole is to stabilize at $(x, \dot{x}, \theta, \dot{\theta}) = (0, 0, 0, 0)$. The dynamics of the cart-pole are [80]:

$$\begin{aligned} \ddot{x} &= \frac{m_p l (\dot{\theta}^2 \sin \theta - \ddot{\theta} \cos \theta) + F}{m_c + m_p}, \\ \ddot{\theta} &= \frac{g \sin \theta + \cos \theta \left(\frac{-m_p l \dot{\theta}^2 \sin \theta - F}{m_c + m_p} \right)}{l \left(\frac{4}{3} - \frac{m_p \cos^2 \theta}{m_c + m_p} \right)}, \end{aligned} \quad (16)$$

where g is the acceleration of gravity.

To control the system, we will employ MPC at 15Hz with a look-ahead horizon $N = 20$. We use quadratic cost functions with $Q = \text{diag}([5.0, 0.1, 5.0, 0.1])$ and $R = 0.1$. We use the fourth-order Runge-Kutta (RK4) method [81] for discretizing the above continuous-time dynamics. The true system parameters are $m_c = 1.0$, $m_p = 0.1$, and $l = 0.5$ but the parameters for the nominal system dynamics are scaled

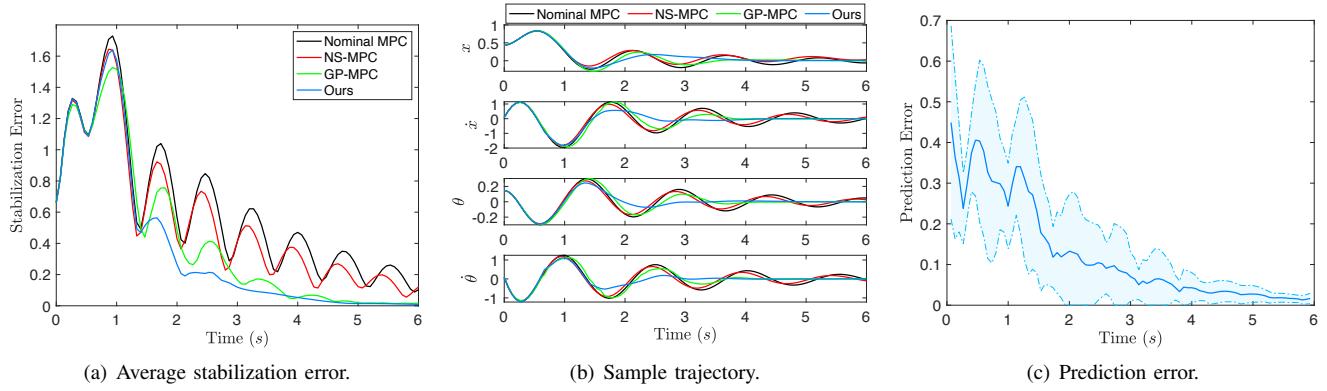


Fig. 3: **Simulation Results of the Cart-Pole Stabilization Experiment in Section V-B.** (a) and (b) demonstrate that Algorithm 1 achieves stabilization in the least time among all tested algorithms. GP-MPC comes second but it incurs a larger deviation from the stabilization goal $(0, 0, 0, 0)$ than Algorithm 1. NS-MPC and Nominal MPC have similar performance, showing that the state-of-the-art non-stochastic control methods are insufficient when the unknown disturbance is adaptive. (c) shows that as Algorithm 1 collects more data, the prediction error decreases.

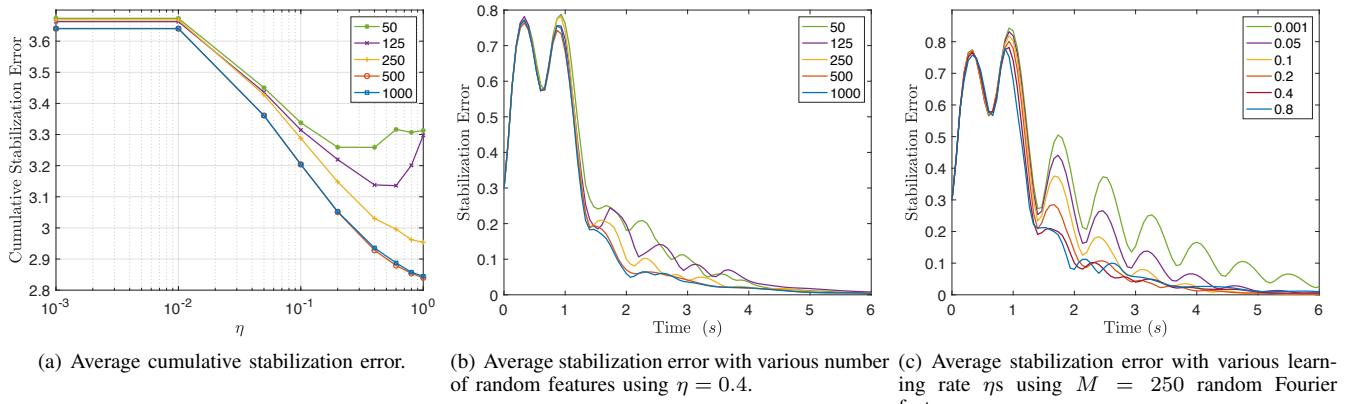


Fig. 4: **Simulation Results for the Sensitivity Analysis in Section V-C over the Cart-Pole System.** The results suggest that large M and η achieves better performance. However, M cannot be arbitrarily large as it increases the computational complexity of solving eq. (4), shown in Table II.

TABLE I: **Computational Performance across Different Methods for the Cart-Pole System in Section V-B.** The table reports the average and standard deviation of computational time in millisecond.

| | MPC | NS-MPC | GP-MPC | Ours |
|-----------|-----------------|-----------------|-------------------|-------------------|
| Time (ms) | 9.63 ± 3.34 | 9.94 ± 3.39 | 24.28 ± 38.24 | 16.28 ± 21.77 |

to 75% of the said true values. We use $z_t = [x_t^\top \ u_t^\top]^\top$. We use $M = 75$ random Fourier features and $\eta = 0.25$, and initialize $\hat{\alpha}$ as a zero vector —an ablation study follows in the next section. We simulate the setting for 6s and perform the simulation 10 times with random initialization sampled uniformly from $x \in [-1, 1]$, $\dot{x} \in [-0.1, 0.1]$, $\theta \in [-0.2, 0.2]$, $\dot{\theta} \in [-0.1, 0.1]$. The simulation environment is based on [82] in Pybullet [83]. We use CasADi [84] to solve eq. (4).

Compared Algorithms. We compare Algorithm 1 with an MPC that uses the (wrong) nominal system parameters (Nominal MPC), the non-stochastic MPC (NS-MPC) [45], and the Gaussian process MPC (GP-MPC) [24]. In more detail, the Nominal MPC uses the nominal dynamics to select control input by solving eq. (3). The NS-MPC augments the Nominal

TABLE II: **Computational Performance across Different M for our Method over the Cart-Pole System using $\eta = 0.4$ in Section V-C.** The table reports the average value and standard deviation of computational time in milliseconds.

| M | 50 | 125 | 250 | 500 | 1000 |
|-----------|-------------------|-------------------|-------------------|--------------------|--------------------|
| Time (ms) | 14.32 ± 18.80 | 23.84 ± 40.36 | 31.32 ± 71.47 | 53.16 ± 148.94 | 96.39 ± 299.99 |

MPC with an additional control v_t input that is updated by running the online gradient descent algorithm to minimize the state cost $\mathbf{x}_{t+1} Q \mathbf{x}_{t+1}^\top$. The GP-MPC learns $\hat{h}(\cdot)$ with a sparse Gaussian process (GP) [85], [86] whose data points are collected online, *i.e.*, GP fixes its hyperparameters and collects data points $(z_t, h(z_t))$ online.

Performance Metric. We evaluate the performance of Nominal MPC, NS-MPC, GP-MPC, and Algorithm 1 in terms of their stabilization error $\|\mathbf{x}_t\|^2$ and computational time. Also, we evaluate the prediction accuracy of Algorithm 1 as we collect more data online via estimation error e_t .

Results. The results are given in Figure 3 and Table I. Figure 3(a) and Figure 3(b) demonstrate that Algorithm 1

achieves stabilization the fastest. GP-MPC comes second but it incurs a larger deviation from the stabilization goal $(0, 0, 0, 0)$ than Algorithm 1. NS-MPC and Nominal MPC have similar performance, showing that the state-of-the-art non-stochastic control methods are insufficient when the unknown disturbance is adaptive. Figure 3(c) shows a fast convergence of the estimation error: the error decreases to around 0.1 within 2s (*i.e.*, less than 30 iterations), which benefits our stabilization goal as shown in Figure 3(b). In Table I, Nominal MPC is the most computationally efficient, followed by NS-MPC since the online gradient descent update of v_t only needs one projection step. Our method maintains the expressiveness of random Fourier features such that it learns the unknown dynamics due to incorrect model parameters and is computationally more efficient than GP-MPC.

C. Sensitivity Analysis over Cart-Pole Scenario

Simulation Setup. We consider the same simulation setup as in section V-B. In addition, we run the simulation with various number of random Fourier features and learning rate: we use $M \in \{50, 125, 250, 500, 1000\}$, and $\eta \in \{0.001, 0.01, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1\}$.

Performance Metric. We evaluate the performance of Algorithm 1 in terms of the stabilization error $\|\mathbf{x}_t\|^2$ and cumulative stabilization error $\sum_{t=1}^T \|\mathbf{x}_t\|^2$. Also, we consider the computational time for calculating control inputs with different M while fixing η .

Results. The results are given in Figure 4 and Table II. First, we notice that Algorithm 1 has similar performance as Nominal MPC for small η , *i.e.*, $\eta \leq 0.01$, regardless of M . The reason is that very small η will keep $\hat{\alpha}$ to stay around zero. Figure 4 suggests that large M and η achieve better performance. For $M = 50$ and $M = 125$, increasing η to 1 causes worse performance, due to overshoot behaviors in estimating $\hat{\alpha}$ and, therefore, steady-state error. This can be avoided by increasing M . Figure 4(a) shows that the cumulative stabilization error cannot be improved after $M \geq 500$. Further, Table II shows that increasing M makes the optimization problem in eq. (4) computationally expensive and prevents the method from real-time implementations, as the frequency of MPC is less than 3 Hz with $M \geq 250$ in Python.

Another tuning parameter is the standard deviation of the Gaussian distribution from which w_i are sampled. This depends on the magnitude of feature vector z_t , since they multiply together per the definition of feature map $\Phi(z, \theta) = \phi(w^\top z + b)$. Hence, if z_t has a large magnitude, the standard deviation might need to be small. Alternatively, we can normalize z_t based on its magnitude and start tuning with a standard Gaussian distribution.

D. Quadrotor Scenario

Simulation Setup. The quadrotor dynamics, the simulation environment, and the controller setup are as follows:

1) Quadrotor Dynamics:

$$\dot{\mathbf{p}} = \mathbf{v}, \quad m\dot{\mathbf{v}} = mg + \mathbf{f} + \mathbf{f}_a, \quad (17)$$

$$\dot{\mathbf{q}} = \frac{1}{2}\mathbf{q} \otimes \begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix}, \quad \mathcal{J}\dot{\boldsymbol{\omega}} = -\boldsymbol{\omega} \times \mathcal{J}\boldsymbol{\omega} + \boldsymbol{\tau}, \quad (18)$$

where $\mathbf{p} \in \mathbb{R}^3$ and $\mathbf{v} \in \mathbb{R}^3$ are position and velocity in the inertial frame, \mathbf{q} is the quaternion, \otimes is the quaternion multiplication operator, $\boldsymbol{\omega} \in \mathbb{R}^3$ is the body angular velocity, m is the quadrotor mass, \mathcal{J} is the inertia matrix of the quadrotor, \mathbf{g} is the gravity vector, $\mathbf{f} = \mathbf{R}[0 \ 0 \ T]^\top \in \mathbb{R}^3$ and $\boldsymbol{\tau} \in \mathbb{R}^3$ are the total thrust and body torques from the four rotors, T is the thrust from the four rotors along the z -axis of the body frame, and $\mathbf{f}_a \in \mathbb{R}^3$ is the aerodynamic force.

2) Gazebo Environment and Control Architecture: We employ the AscTec Hummingbird quadrotor model using the RotorS simulator in Gazebo [87]. The RotorS simulator implements the rotor drag as aerodynamic effects, which is a linear mapping with respect to the body frame velocity [88]. We use the following control architecture: the MPC is running at 50Hz, takes as input the reference trajectory, and outputs the desired total thrust and desired body angular velocity to the low-level body rate controller [89], [90]; the body rate controller is running at 200Hz, then converts them into motor speed commands to control the quadrotor in the simulator; the state of the quadrotor is available at 100Hz.

3) Control Design: The MPC uses look-ahead horizon $N = 10$ simulating the quadrotor dynamics for 1s. We use quadratic cost functions with $Q = \text{diag}([0.5\mathbf{I}_3, 0.1\mathbf{I}_4, 0.05\mathbf{I}_3, 0.01\mathbf{I}_3])$ and $R = \mathbf{I}_4$. We use the RK4 method [81] for discretization. We use as the feature z_t the velocity, quaternion, body angular velocity, and normalized thrust from four rotors. We sample w_i from a Gaussian distribution with standard deviation 0.01. We use $M = 50$ random Fourier features and $\eta = 0.25$, and initialize $\hat{\alpha}$ as a zero vector. We use CasADi [84] and acados [91] to solve eq. (4).

4) Benchmark Experiment Setup: We consider that the quadrotor is tasked to track a prescribed trajectory at different maximal speeds v_m , affecting the aerodynamic forces; we consider four types of reference trajectories: Circle, Wrapped Circle, Lemniscate, and Wrapped Lemniscate, showed in Figure 5. We simulate each reference trajectory and each maximal speed 5 times. We use as the performance metric the root mean squared error (RMSE) in position.

Compared Algorithms. We compare Algorithm 1 with: a nominal MPC that assumes no aerodynamic forces (Nominal MPC), and the Gaussian process MPC (GP-MPC) [23]. The GP model in [23] is pre-trained. We adopt the default training procedure per [23]'s open-sourced code: a Nominal MPC is given 10 random trajectories with maximal speeds $v_m \in [6.95, 15.62]$; then the Nominal MPC commands the quadrotor to track these trajectories and collects the training dataset with 3556 data points; finally, the GP model is trained such that \mathbf{f}_a is predicted based on body velocities. The prediction of the GP model is used only for the first step over the look-ahead horizon. We also combine Algorithm 1 with incremental non-linear dynamic inversion (INDI) [16] to account for unknown disturbance in the attitude dynamics in eq. (18).

Results. The results are given in Figure 6 and Figure 7. In Figure 6, Algorithm 1 demonstrates improved performance over Nominal MPC and GP-MPC in terms of the tracking error over all tested reference trajectories and maximal speeds. The limitation of GP-MPC appears to be: (i) due to computational complexity, the prediction of GP is only used for one step in MPC; by contrast, Algorithm 1 incorporates $\hat{h}(\cdot)$ over the entire look-ahead horizon N ; (ii) GP-MPC may not

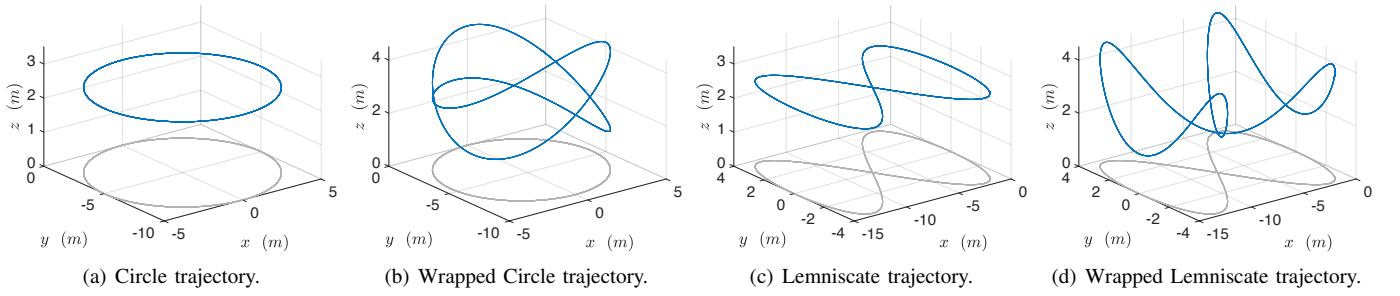


Fig. 5: **Reference Trajectory for the Quadrotor Experiments in Section V-D.** The blue lines represent the reference trajectories in 3D. The gray lines are the projection of reference trajectories onto the ground.

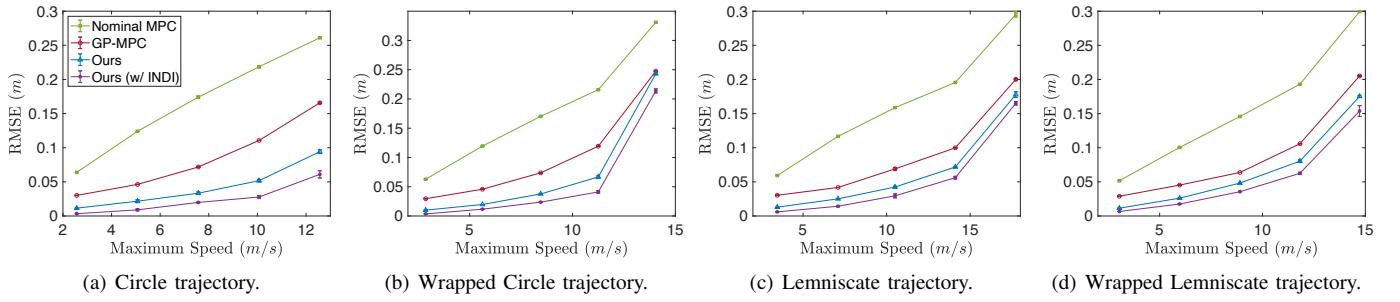


Fig. 6: **Tracking Performance Comparison for the Quadrotor Experiments in Section V-D.** Algorithm 1 demonstrates improved performance compared to Nominal MPC and GP-MPC in terms of tracking error over all tested reference trajectories and maximal speeds. Algorithm 1 with INDI achieves the best performance as INDI provides better tracking in attitude dynamics.

perform well if the trajectories in the training dataset are different from the executed ones; in contrast, Algorithm 1 aims to collect data and learn $\hat{h}(\cdot)$ online, thus alleviates such generalization errors. Additionally, the performance of Algorithm 1 is further improved with INDI. However, due to the low frequency in INDI control loop ($200Hz$) and state estimate ($100Hz$), the improvement is marginal especially in high-speed scenarios.³ Figure 7 presents the trajectories of Nominal MPC, GP-MPC, and Algorithm 1 tracking a Circle trajectory with $v_m = 10.07 m/s$, and a Lemniscate trajectory with $v_m = 14.11 m/s$, from $t = 40s$ to $t = 50s$ when the quadrotor reaches the maximal speeds. In Circle trajectory, Algorithm 1 clearly achieves the best tracking performance. In Lemniscate trajectory, Algorithm 1 has good tracking performance except in the corners where the quadrotor needs to turn; while Nominal MPC and GP-MPC have tracking errors along the whole trajectory.

VI. HARDWARE EXPERIMENTS

We evaluate Algorithm 1 in extensive real-world scenarios of control under uncertainty. For the experiments, we use a quadrotor as shown in Figure 8 to track a circular trajectory despite ground effects and wind disturbances (Figure 1). Additional results under aerodynamic drag effects and turbulent effects are given in Appendix I.

³To achieve accurate tracking performance, INDI requires high-frequency control update (*e.g.*, $500Hz$ in [16], $300Hz$ in [57]) and measurements (*e.g.*, $500Hz$ motor speed and IMU measurements in [57]) to accurately estimate the external disturbances.

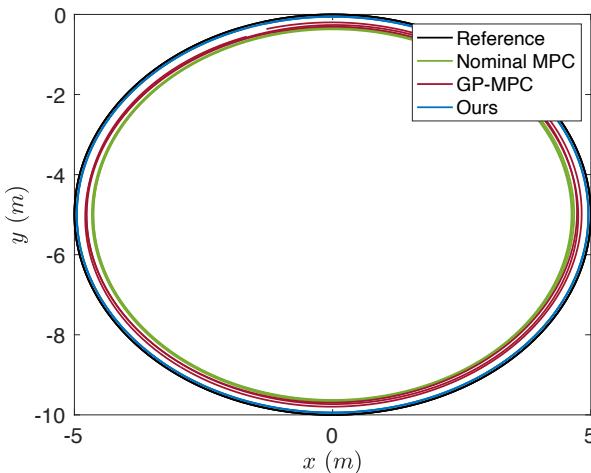
Quadrotor. The quadrotor dynamics are modeled as in eqs. (17) and (18) with weight $0.68kg$ and $0.22m$ diagonal motor-to-motor distance. The quadrotor is equipped with an Nvidia Jetson Orin NX 16GB [92] to handle on-board computation. We use the low-level controller in Holybro Pix32 v6 running PX4-Autopilot firmware [93]. The communication between Jetson Orin NX and PX4 is through MAVROS [94]. A Vicon motion capture system provides the pose of the quadrotor at $30Hz$, and we use PX4's Extended Kalman filter (EKF) to fuse the Vicon measurements with IMU readings onboard to estimate the quadrotor's odometry. The odometry is sent to the controller at $100Hz$ and MPC runs at $50Hz$.

Control Design. We use the same MPC parameters and feature z_t as in Section V-D. We use $M = 25$ random Fourier features and $\eta = 0.05$, and initialize $\hat{\alpha}$ as a zero vector. Those parameters are fixed for all hardware experiments. We use CasADi [84] and acados [91] to solve eq. (4).

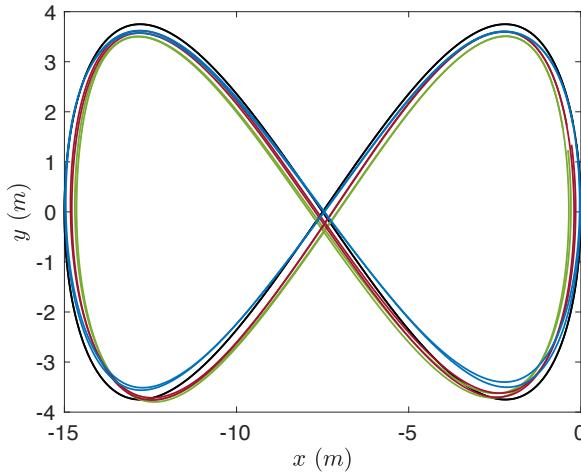
Compared Algorithms. We compare Algorithm 1 with a nominal MPC that assumes no aerodynamic forces (Nominal MPC), and the L1-adaptive MPC (L1-MPC) [17]. The L1-MPC uses the L1 adaptive control to estimate the aerodynamic force and moment in eqs. (17) and (18). Then, the MPC calculates the control input based on the estimated aerodynamic force and moment. We adopt the default parameters of estimation using L1 adaptive control per [17]'s open-sourced code.

Benchmark Experiment Setup. We consider that the quadrotor is tasked to track a circular trajectory with diameter $1m$ at maximal speed $v_m = 0.8 m/s$. We consider the following disturbance scenarios:

- *Ground Effects:* We use an elevated foam mat to create



(a) Sample trajectories in tracking a Circle trajectory with $v_m = 10.07 \text{ m/s}$.



(b) Sample trajectories in tracking a Lemniscate trajectory with $v_m = 14.11 \text{ m/s}$.

Fig. 7: Tracking Performance Comparison for the Quadrotor Experiments of Circle and Lemniscate Trajectories in Section V-D. The plots present the trajectories of Nominal MPC, GP-MPC, and Algorithm 1 tracking a Circle trajectory with $v_m = 10.07 \text{ m/s}$, and a Lemniscate trajectory with $v_m = 14.11 \text{ m/s}$, from $t = 40\text{s}$ to $t = 50\text{s}$ when the quadrotor reaches the maximal speeds. The trajectory of Algorithm 1 with INDI is omitted as it overlaps with Algorithm 1 w/o INDI.

the ground effects to the quadrotor (Figure 1, left). The surface of the foam mat is 0.78m high and the circular trajectory is set to be 0.85m in height. Half of the circular trajectory is above the foam mat and the other half is outside. This setting creates significant ground effects when the quadrotor is above the foam mat and a sharp transition between w/ and w/o the ground effects.

- **Wind Disturbances:** We use a fan to create wind disturbances (Figure 1, center). The circular trajectory is at 1m height and its center is 3m away from the fan. The quadrotor experiences wind speed from 2.0 m/s to 5.4 m/s along the circular trajectory, depending on its distance to the fan. The trajectory ends at a position with wind speed from 3.5 m/s to 5.0 m/s .



Fig. 8: The quadrotor for the hardware experiments.

- **Ground Effects + Wind Disturbances:** We combine the previous two scenarios and now the quadrotor suffers from both ground effects and wind disturbances (Figure 1, right). The circular trajectory is at 0.85m and the wind speed varies from 1.5 m/s to 5.0 m/s along the circular trajectory. The trajectory ends at a position with wind speed from 3.3 m/s to 4.7 m/s .

Additional unknown disturbances come from the (i) EKF estimation errors, (iii) battery's voltage drop during the flight, which causes a decrease in motors' thrust, (iii) communication latency between Jetson Orin NX and PX4, and (iv) modeling error of quadrotor dynamics. The resulting unknown h is a combination of multi-source heterogeneous disturbances. We conduct each scenario 5 times. We use as the performance metric the RMSE in position.

Results. The results are given in Figure 9, Figure 10, Figure 11, and Figure 12. Figure 9 shows that our algorithm achieves the best performance in average RMSE over all tested scenarios. Nominal MPC is the worst as it does not account for unknown disturbances, and it has increased z -direction tracking error during the flight due to insufficient thrust caused by voltage drop (Figure 10 and Figure 11, top right). Nominal MPC fails under both ground effects and wind disturbances (Figure 12) due to strong wind and voltage drop. L1-MPC has improved performance over Nominal MPC by compensating for the estimated disturbances. However, it does not account for the future evolution of the unknown disturbances and has large tracking error when the unknown disturbances change abruptly, *e.g.*, a sharp transition between the ground and free space (Figure 10, top right). In contrast, Algorithm 1 learns the model of unknown disturbances for predictive control, and therefore, achieves the best tracking performance even when the unknown disturbances change abruptly. Algorithm 1 achieves up to 75% and 60% improvement in average RMSE over Nominal MPC and L1-MPC, respectively. In addition, Algorithm 1 using the same set of parameters can learn different unknown disturbances, *i.e.*, ground effects (Figure 10), wind disturbances (Figure 11), a combination of ground effects and wind disturbances (Figure 12), and voltage drop. This is enabled by the self-supervised learning framework using the data collected online.

VII. RELATED WORK

We discuss work on system identification and control across the research topics of adaptive control; robust control; control

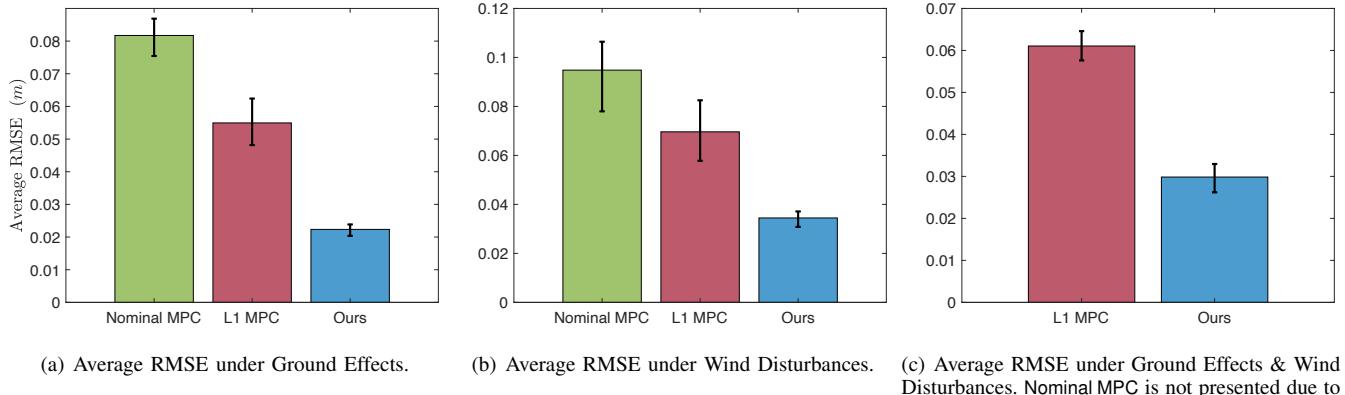


Fig. 9: **Tracking Performance Comparison for the Hardware Quadrotor Experiments in Section VI.** The error bar represents the minimum and maximum RMSE. Algorithm 1 (ours) demonstrates improved performance compared to Nominal MPC and L1-MPC in terms of tracking error over all tested scenarios.

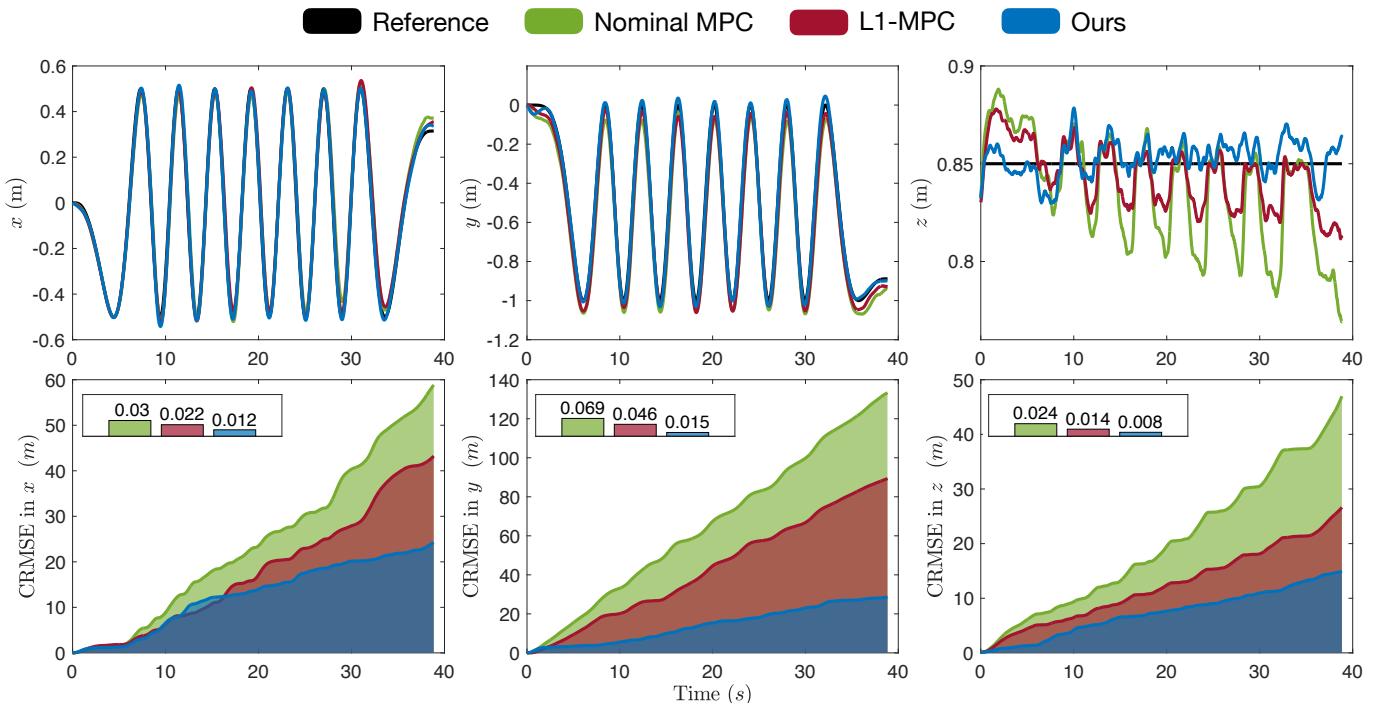


Fig. 10: **Tracking Performance Comparison for the Hardware Quadrotor Experiments Under Ground Effects in Section VI.** The plots present sample trajectories of Nominal MPC, L1-MPC, and Algorithm 1 in x -, y -, and z - position (top) and the corresponding the cumulative RMSE (CRMSE) (bottom). The inserted plots present the average RMSE (m).

based on offline learning; control based on offline learning with online adaptation; and control based on online learning.⁴

Adaptive control: Adaptive control methods often assume parametric uncertainty additive to the known system dynamics,

⁴The related work also includes papers on *sequential online model learning and predictive control* [62], [65]. These works assume no knowledge of system dynamics and instead learn the whole system dynamics online. Therefore, their methods first need to use random control inputs to excite the systems and estimate the unknown system dynamics; then they design control policy using the estimated system dynamics. In contrast, in this paper, we consider only partially unknown dynamics, and we simultaneously estimate unknown dynamics/disturbances and control the system. Finally, [62], [65] require persistent excitation to obtain regret guarantees, while we do not.

e.g., parametric uncertainty in the form of unknown coefficients multiplying known basis functions [11]–[13]. These coefficients are updated online and generate an adaptive control input to compensate for the estimated disturbances. These methods often require persistent excitation to guarantee the exponential stability of the closed-loop system [11]–[13]. In contrast, inspired by [15], our method handles unknown dynamics and disturbances that can be expressed in reproducing kernel Hilbert space (RKHS), and does not require persistent excitation to enjoy near-optimality guarantees. [16]–[18] bypass the assumption of parametric uncertainty and persistent excitation, and directly estimate the value of unknown

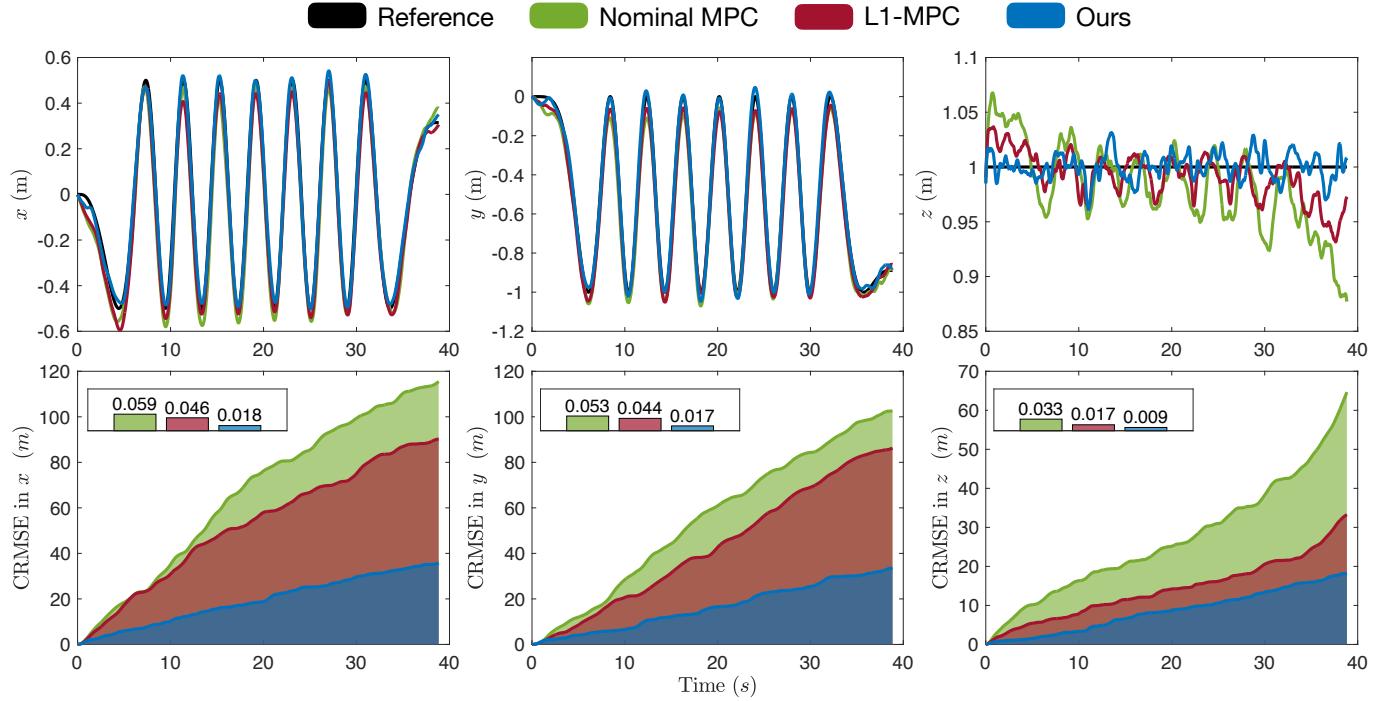


Fig. 11: **Tracking Performance Comparison for the Hardware Quadrotor Experiments Under Wind Disturbances in Section VI.** The plots present the sample trajectories of Nominal MPC, L1-MPC, and Algorithm 1 in x -, y -, and z - position (top) and the corresponding CRMSE (bottom). The inserted plots present the average RMSE (m).

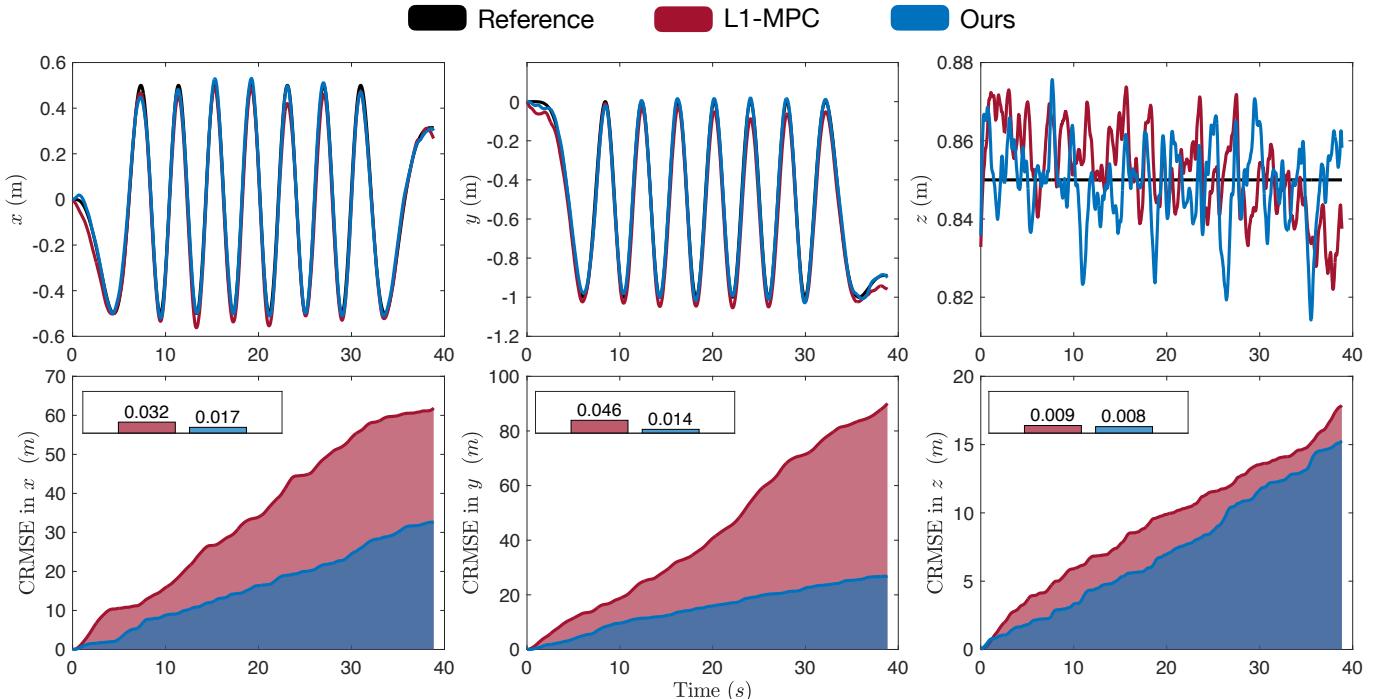


Fig. 12: **Tracking Performance Comparison for the Hardware Quadrotor Experiments under both Ground Effects and Wind Disturbances in Section VI.** The plots present sample trajectories of L1-MPC and Algorithm 1 in x -, y -, and z - position (top) and the corresponding CRMSE (bottom). The inserted plots present the average RMSE (m). Nominal MPC is not presented since it causes crashes under both ground effects and wind disturbances.

disturbances. However, the methods therein, as well as, the relevant method in [14], [15], [19], focus on adaptive control to compensate for the estimated disturbances, instead of learning a model of the disturbances to enable model predictive control.

Robust control: Robust control algorithms select control inputs upon simulating the future system dynamics across a lookahead horizon [4]–[10]. To this end, they assume a worst-case realization of disturbances, given an upper bound to the magnitude of the disturbances [4]–[10]. However, assuming the worst-case disturbances can be conservative. In addition, [5]–[10] focus on linear dynamical systems, and [7]–[10] may be computationally expensive. Hence, the application of these approaches is often limited to low-dimensional linear systems.

Offline learning for control: This line of work trains neural-network, Gaussian-process, or Koopman models using training data collected offline [21]–[31]. [25]–[27] train the models with data collected from random environments such that the controller can exhibit robustness to unseen environments [53], [54]. In this paper, instead, we require no offline data collection and training, employ one-shot online learning in a self-supervised manner based on data collected online.

Offline learning and online adaption for control: Such control methods have two components: offline learning and online adaptation [32]–[38]. The offline learning aims to train a neural network model using data collected from different environments, *e.g.*, the disturbances of a quadrotor under different wind speeds, so that the learned neural network model captures the underlying features across different environments. The online adaptation aims to update online either all the parameters of the neural network or only the parameters of the last layer to better predict unknown disturbances when encountering new environments. Similar to data-driven control with only offline learning, the methods need big data for effective training. Also, updating all parameters online can be computationally heavy for real-time control [32]–[35]. [36]–[38] only update the last layer parameters, but either directly compensate the estimated disturbances without predicting their future evolution to optimize the controller [36], [37] or it is computationally expensive to use a neural network in on-board model predictive control [38]. In contrast, we propose to use random Fourier features to approximate online a model the unknown dynamics and disturbances. Our approach allows us to maintain the computational efficiency of classical finite-dimensional parametric approximations that can be used in model predictive control while retaining the expressiveness of the RKHS with high probability.

Online learning for control: We discuss two lines of online learning algorithms: non-stochastic control algorithms [39]–[45] and contextual optimization-based control algorithms [46]–[51]. These methods often quantify the performance of the algorithms through *regret*, *i.e.*, the suboptimality against an optimal clairvoyant controller that knows the unknown disturbances and dynamics. Non-stochastic control algorithms select control inputs based on past information only since they assume no model that can be used to simulate the future evolution of the disturbances [39]–[45]. Instead, they provide controllers with bounded regret guarantees, upon employing the OCO framework to capture the non-stochastic control problem as a sequential game between a controller and an adversary [59]. They rely on the knowledge of a known

upper bound to the magnitude of the disturbances, and a pre-stabilizing controller. The proposed approaches have been observed to be sensitive to the choice of the pre-stabilizing controller and the tuning parameters of OGD, *e.g.*, in [40], [42], [44]. Contextual optimization-based control algorithms ignore the system identification step, and optimize control by minimizing a one-step surrogate cost function, also providing regret guarantees [46]–[51]. Specifically, such surrogate function is model based on past information, and depends on the current control action to be optimized and the known current state. Then the current control is selected by minimizing this one-step surrogate cost function due to the lack of a predictive model. In contrast, our method uses OGD to learn the model of the disturbances instead of optimizing the online controller, and uses model predictive control to optimize online the control input based on the learned online disturbance model.⁵ In addition, we consider a stronger definition of regret than [40]–[45], [49], [50] by (i) allowing the unknown dynamics and disturbances adaptive to the state and control, instead of the same disturbances for the designed controller and the optimal controller [40]–[45], (ii) allowing the optimal controller to select control input with multi-step prediction into the future, instead of only one-step prediction [49], [50].

VIII. CONCLUSION

We provided Algorithm 1 for the problem of *Simultaneous System Identification and Model Predictive Control* (Problem 1). Algorithm 1 guarantees no dynamic regret against an optimal clairvoyant (non-causal) policy that knows the disturbance function h *a priori* (Theorem 1). The algorithm uses random Fourier features to approximate the unknown dynamics or disturbances in reproducing kernel Hilbert spaces. Then, it employs model predictive control based on the current learned model of the unknown dynamics (or disturbances). The model of the unknown dynamics is updated online in a self-supervised manner using least squares based on the data collected while controlling the system.

We validate Algorithm 1 in simulated and hardware experiments. Our physics-based simulations include (i) a cart-pole aiming to maintain the pole upright despite inaccurate model parameters (Sections V-B and V-C), and (ii) a quadrotor aiming to track reference trajectories despite unmodeled aerodynamic drag effects (Section V-D). Our hardware experiments involve a quadrotor aiming to track a circular trajectory despite unmodeled aerodynamic drag effects, ground effects, and wind disturbances (Section VI and Appendix I). We observed that our method demonstrated better tracking performance and computational efficiency than the state-of-the-art methods GP-MPC [23], [24], NS-MPC [45], and L1-MPC [17].

ACKNOWLEDGEMENTS

We thank Torbjørn Cunis, Ilya Kolmanovsky, and Miguel Castroviejo-Fernandez from the University of Michigan, for their invaluable discussion on the Lipschitzness of value

⁵In [95], a predictive control method with online learning is provided. It focuses on learning unknown trajectories of a target, where the target trajectory does not adapt to the pursuer's motion. In contrast, we focus on learning unknown dynamics/disturbances that are adaptive to the controller's state and control input. In addition, [95] focuses on linear systems and linear MPC, while we utilize nonlinear MPC for nonlinear control-affine systems.

function in nonlinear MPC. We thank Devansh Agrawal, Hardik Parwana, and Dimitra Panagou from the University of Michigan, for their invaluable help in hardware experiments.

APPENDIX

A. Extension to Systems with Unmodeled Process Noise

In this section, we provide the extension of Algorithm 1 to systems corrupted with also unmodeled (potentially non-stochastic) noise that is additive to the system dynamics, *i.e.*, *i.e.*

$$x_{t+1} = f(x_t) + g(x_t)u_t + h(z_t) + w_t, \quad (19)$$

where w_t represents the unmodeled process noise.

Then, Algorithm 1 guarantees the following.

Corollary 1 (Dynamic Regret in Systems with Process Noise). *Assume Algorithm 1's learning rate is $\eta = \mathcal{O}(1/\sqrt{T})$. Then, for the system in eq. (19), Algorithm 1 achieves*

$$\text{Regret}_T^D \leq \mathcal{O}\left(T^{\frac{3}{4}}\right) + L\sqrt{T \sum_{t=1}^T \|w_t\|^2}. \quad (20)$$

Compared to Theorem 1, the regret bound in Corollary 1 has an additional term that depends on the energy of the noise sequence (w_1, \dots, w_T) , *i.e.*, $\sum_{t=1}^T \|w_t\|^2$. Specifically, when the energy is less than $\mathcal{O}(T)$, we achieve sublinear regret.

B. Proof of Theorem 1

We use $x_{t_1:t_2}$ to denote the sequence $(x_{t_1}, \dots, x_{t_2})$, and $\Phi(\cdot) \triangleq \frac{1}{M}[\Phi(\cdot, \theta_1), \dots, \Phi(\cdot, \theta_M)]$. We use \hat{x} to denote the one-step-ahead prediction given an estimate $\hat{\alpha}$, *i.e.*, $\hat{x}_{t+1} = f(x_k) + g(x_k)u_k + \Phi(z_k)\hat{\alpha}_t = x_{t+1} + \Phi(z_k)\hat{\alpha} - h(z_k)$.

Then, the dynamic regret in Definition 2 can be upper bounded as follows,

$$\begin{aligned} \text{Regret}_T^D &= \sum_{t=1}^T c_t(x_t, u_t) - \sum_{t=1}^T c_t(x_t^*, u_t^*) \\ &\leq \sum_{t=1}^T V_t(x_t; \hat{\alpha}_t) - \sum_{t=1}^T c_t(x_t^*, u_t^*) \\ &\leq \sum_{t=1}^T V_t(x_t; \hat{\alpha}_t), \end{aligned} \quad (21)$$

where the first inequality holds by definition of $V_t(x_t; \hat{\alpha}_t)$, and the second inequality holds by $\sum_{t=1}^T c_t(x_t^*, u_t^*) \geq 0$.

Hence, we aim to bound $\sum_{t=1}^T V_t(x_t; \hat{\alpha}_t)$. To this end, we first establish the relationship between $V_t(x_t; \hat{\alpha}_t)$ and $V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1})$ using the following lemmas (the proofs of them are given in Appendix D and Appendix E).

Lemma 1. *Suppose Assumption 1, Assumption 3, and Assumption 4, hold, then for all $N > (\bar{\lambda}/\underline{\lambda})^2 + 1$, $V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t)$ and $V_t(x_t; \hat{\alpha}_t)$ satisfy*

$$V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t) \leq \epsilon V_t(x_t; \hat{\alpha}_t), \quad (22)$$

where $\epsilon = 1 - \left(1 - \frac{(\bar{\lambda}/\underline{\lambda})^2}{N-1}\right)$ ($\underline{\lambda}/\bar{\lambda} \in (0, 1)$).

Lemma 1 indicates that the system in eq. (4b) is globally asymptotic stable [68] when the MPC policy is applied: the

value function keeps decreasing and converges to zero, thus, the state must converge to zero.

Lemma 2. *Suppose Assumption 2, Assumption 3, Assumption 4, and Assumption 5, hold, then $V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1})$ and $V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t)$ for eq. (1) satisfy*

$$V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1}) - V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t) \leq L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right). \quad (23)$$

Lemma 2 is achieved upon using Assumption 2 to upper bound $|V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1}) - V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t)|$ with $\|x_{t+1} - \hat{x}_{t+1}\|$ and $\|\hat{\alpha}_t - \hat{\alpha}_{t+1}\|$.

Using Lemma 1 and Lemma 2, we have

$$V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1}) \leq \epsilon V_t(x_t; \hat{\alpha}_t) + L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right). \quad (24)$$

Then, we can bound $V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1})$ by

$$\begin{aligned} &V_{t+1}(x_{t+1}; \hat{\alpha}_{t+1}) \\ &\leq \epsilon V_t(x_t; \hat{\alpha}_t) + L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right) \\ &\leq \epsilon^2 V_{t-1}(x_{t-1}; \hat{\alpha}_{t-1}) + \epsilon L \left(\sqrt{l_{t-1}(\hat{\alpha}_{t-1})} + \|\eta \nabla_{t-1}\| \right) \\ &\quad + L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right) \\ &\leq \epsilon^t V_1(x_1; \hat{\alpha}_1) + \sum_{k=1}^t \epsilon^{t-k} L \left(\sqrt{l_k(\hat{\alpha}_k)} + \|\eta \nabla_k\| \right) \\ &\leq \epsilon^t \bar{\lambda} \sigma(x_1) + \sum_{k=1}^t \epsilon^{t-k} L \left(\sqrt{l_k(\hat{\alpha}_k)} + \|\eta \nabla_k\| \right), \end{aligned} \quad (25)$$

where the last inequality uses Assumption 1 (ii).

Therefore, we have

$$\begin{aligned} &\sum_{t=1}^T V_t(x_t; \hat{\alpha}_t) \\ &\leq \sum_{t=1}^T \epsilon^{t-1} \bar{\lambda} \sigma(x_1) + \sum_{t=1}^T \sum_{k=1}^{t-1} \epsilon^{t-1-k} L \left(\sqrt{l_k(\hat{\alpha}_k)} + \|\eta \nabla_k\| \right) \\ &\leq \sum_{t=1}^T \epsilon^{t-1} \bar{\lambda} \sigma(x_1) + \left(\sum_{t=0}^T \epsilon^t \right) \left(\sum_{t=1}^T L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right) \right) \\ &\leq \frac{\bar{\lambda}}{1-\epsilon} \sigma(x_1) + \frac{1}{1-\epsilon} \left(\sum_{t=1}^T L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right) \right), \end{aligned} \quad (26)$$

where the second inequality holds by adding positive terms to $\sum_{t=1}^T \sum_{k=1}^{t-1} \epsilon^{t-1-k} L \left(\sqrt{l_k(\hat{\alpha}_k)} + \|\eta \nabla_k\| \right)$ to complete $\left(\sum_{t=0}^T \epsilon^t \right) \left(\sum_{t=1}^T L \left(\sqrt{l_t(\hat{\alpha}_t)} + \|\eta \nabla_t\| \right) \right)$.

Since x_1 is bounded, we have

$$\frac{\bar{\lambda}}{1-\epsilon} \sigma(x_1) = \mathcal{O}(1). \quad (27)$$

We now upper bound the term $\sum_{t=1}^T L \sqrt{l_t(\hat{\alpha}_t)}$. By the Cauchy-Schwarz inequality, we have

$$\sum_{t=1}^T L \sqrt{l_t(\hat{\alpha}_t)} \leq L \sqrt{T} \sqrt{\sum_{t=1}^T l_t(\hat{\alpha}_t)}. \quad (28)$$

From Proposition 2, we have $\sum_{t=1}^T l_t(\hat{\alpha}_t) \leq \mathcal{O}(\sqrt{T})$. Thereby, we obtain

$$\sum_{t=1}^T L\sqrt{l_t(\hat{\alpha}_t)} \leq \mathcal{O}\left(T^{\frac{3}{4}}\right). \quad (29)$$

Next, we consider the term $\sum_{t=1}^T L\|\eta\nabla_t\|$. Since $\Phi(\cdot)$, α , $\hat{\alpha}_t$ are uniformly bounded for all t , then ∇_t is bounded by a constant G . Then

$$\sum_{t=1}^T L\|\eta\nabla_t\| \leq \sum_{t=1}^T LG\|\eta\| \leq \mathcal{O}\left(\sqrt{T}\right), \quad (30)$$

where the last inequality holds by $\eta = \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$.

Combining eqs. (27), (29) and (30) gives the result in Theorem 1. \square

C. Proof of Corollary 1

For eq. (19) where w_t is the unmodeled process noise, we follow similar steps of the proof of Theorem 1, with the difference of eq. (29):

$$\sum_{t=1}^T L\sqrt{l_t(\hat{\alpha}_t)} \leq \mathcal{O}\left(T^{\frac{3}{4}}\right) + L\sqrt{T\sum_{t=1}^T \|w_t\|^2}, \quad (31)$$

due to $\sum_{t=1}^T l_t(\hat{\alpha}_t) - \sum_{t=1}^T l_t(\alpha) = \sum_{t=1}^T l_t(\hat{\alpha}_t) - \sum_{t=1}^T \|w_t\|^2 \leq \mathcal{O}(\sqrt{T})$ per Proposition 2.

Hence, we obtain

$$\text{Regret}_T^D \leq \mathcal{O}\left(T^{\frac{3}{4}}\right) + L\sqrt{T\sum_{t=1}^T \|w_t\|^2}. \quad (32)$$

\square

D. Proof of Lemma 1

Consider that at time step t , the system in eq. (1) is at state x_t and the MPC problem becomes

$$\begin{aligned} & \underset{u_t, \dots, u_{t+N-1}}{\text{argmin}} \sum_{k=t}^{t+N-1} c_k(x_k, u_k) \\ & \text{subject to} \quad x_{k+1} = f(x_k) + g(x_k)u_k + \Phi(z_k)\hat{\alpha}_t, \\ & \quad u_k \in \mathcal{U}, k \in \{t, \dots, t+N-1\}. \end{aligned} \quad (33)$$

Let $\hat{u}_{0:N-1}^t$ and $\hat{x}_{t:t+N}^t$ be the optimal control input and state sequences to the above problem, where the superscript denotes the MPC problem is solved at time step t . We use $\hat{x}_{t+k}^t \triangleq \psi_t(k, x_t, \hat{u}_{0:N-1}^t; \hat{\alpha}_t)$ as the state reached from x_t by applying $\hat{u}_{0:k-1}^t$ to the system dynamics with parameter $\hat{\alpha}_t$. By definition, $\psi_t(k, x_t, \hat{u}_{0:N-1}^t; \hat{\alpha}_t)$ satisfies

$$\psi_t(k, x_t, \hat{u}_{0:N-1}^t; \hat{\alpha}_t) = \psi_t(k, x_t, \hat{u}_{0:k-1}^t; \hat{\alpha}_t), \quad (34)$$

$$\psi_t(k, x_t, \hat{u}_{0:k-1}^t; \hat{\alpha}_t) = \psi_{t+1}(k-1, \hat{x}_{t+1}^t, \hat{u}_{1:k-1}^t; \hat{\alpha}_t), \quad (35)$$

where the first equality holds since \hat{x}_{t+k}^t does not depend on $\hat{u}_{k:N-1}^t$, and the second equality holds by the evolution of system dynamics $x_{k+1} = f(x_k) + g(x_k)u_k + \Phi(z_k)\hat{\alpha}_t$.

Then, we have

$$\begin{aligned} & V_t(x_t; \hat{\alpha}_t) \\ &= \sum_{k=0}^{N-1} c_{t+k}(\psi_t(k, x_t, \hat{u}_{0:k-1}^t; \hat{\alpha}_t), \hat{u}_k^t) \\ &= c_t(x_k, \hat{u}_0^t) + \sum_{k=0}^{N-2} c_{t+k+1}(\psi_t(k+1, x_t, \hat{u}_{0:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ &= c_t(x_k, \hat{u}_0^t) + \sum_{k=0}^{N-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ &= c_t(x_k, \hat{u}_0^t) + \sum_{k=0}^{j-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ & \quad + \sum_{k=j-1}^{N-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ &= c_t(x_k, \hat{u}_0^t) + \sum_{k=0}^{j-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ & \quad + \sum_{k=0}^{N-j-1} c_{t+k+j}(\psi_{t+1}(k+j-1, \hat{x}_{t+1}^t, \hat{u}_{1:k+j-1}^t; \hat{\alpha}_t), \hat{u}_{k+j}^t) \\ &= c_t(x_k, \hat{u}_0^t) + \sum_{k=0}^{j-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ & \quad + \sum_{k=0}^{N-j-1} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t). \end{aligned} \quad (36)$$

Similarly,

$$\begin{aligned} & V_{t+1}(\hat{x}_{t+1}; \hat{\alpha}_t) = V_{t+1}(\hat{x}_{t+1}^t; \hat{\alpha}_t) \\ &= \sum_{k=0}^{N-1} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{0:k-1}^{t+1}; \hat{\alpha}_t), \hat{u}_k^{t+1}) \\ &= \sum_{k=0}^{j-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{0:k-1}^{t+1}; \hat{\alpha}_t), \hat{u}_k^{t+1}) \\ & \quad + \sum_{k=j-1}^{N-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{0:k-1}^{t+1}; \hat{\alpha}_t), \hat{u}_k^{t+1}) \\ &= \sum_{k=0}^{N-j} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{0:k-1}^{t+1}; \hat{\alpha}_t), \hat{u}_k^{t+1}) \\ & \quad + \sum_{k=0}^{N-j} c_{t+k+j}(\psi_{t+1}(k+j-1, \hat{x}_{t+1}^t, \hat{u}_{0:k+j-2}^{t+1}; \hat{\alpha}_t), \hat{u}_{k+j-1}^{t+1}) \\ &\leq \sum_{k=0}^{N-j} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ & \quad + \min_{v_{0:N-j}} \sum_{k=0}^{N-j} c_{t+k+j}(\psi_{t+1}(k, \hat{x}_{t+1}^t, v_{0:k-1}; \hat{\alpha}_t), v_k) \\ &\leq \sum_{k=0}^{j-2} c_{t+k+1}(\psi_{t+1}(k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) \\ & \quad + V_{t+j}(\hat{x}_{t+j}; \hat{\alpha}_t) \end{aligned}$$

$$\leq \sum_{k=0}^{j-2} c_{t+k+1} (\psi_{t+1} (k, \hat{x}_{t+1}^t, \hat{u}_{1:k}^t; \hat{\alpha}_t), \hat{u}_{k+1}^t) + \bar{\lambda} \sigma (\hat{x}_{t+j}^t), \quad (37)$$

where the first inequality is due to sub-optimal control sequence $\hat{u}_{1:k}^t$ and the definition of $\hat{u}_{j-1:N-1}^t$, the second inequality holds by definition of $V_{t+j} (\hat{x}_{t+j}^t; \hat{\alpha}_t)$, and the last inequality is due to Assumption 1 (ii).

Combining eq. (36) and eq. (37) gives

$$\begin{aligned} & V_{t+1} (\hat{x}_{t+1}; \hat{\alpha}_t) - V_t (x_t; \hat{\alpha}_t) \\ & \leq \bar{\lambda} \sigma (\hat{x}_{t+j}^t) - c_t (x_t, \hat{u}_0^t) \\ & \leq \bar{\lambda} \sigma (\hat{x}_{t+j}^t) - \underline{\lambda} \sigma (x_t), \end{aligned} \quad (38)$$

where the last inequality uses Assumption 1 (i).

Using Assumption 1 and $\hat{x}_{t+j}^t = \psi_t (j, \hat{x}_t^t, \hat{u}_{0:j-1}^t; \hat{\alpha}_t)$, for any $j \in [1, N-1]$, we have

$$\begin{aligned} & \underline{\lambda} \sum_{k=0}^{N-1} \sigma (\hat{x}_{t+j}^t) \\ & \leq \sum_{k=0}^{N-1} c_{t+k} (\psi_t (j, \hat{x}_t^t, \hat{u}_{0:j-1}^t; \hat{\alpha}_t), \hat{u}_k^t) \\ & = V_t (x_t; \hat{\alpha}_t) \\ & \leq \bar{\lambda} \sigma (x_t). \end{aligned} \quad (39)$$

Hence, there exists $j \in [1, N-1]$, such that $\sigma (\hat{x}_{t+j}^t) \leq \frac{\bar{\lambda}/\underline{\lambda}}{N-1} \sigma (x_t)$. Then, eq. (38) becomes

$$\begin{aligned} & V_{t+1} (\hat{x}_{t+1}; \hat{\alpha}_t) - V_t (x_t; \hat{\alpha}_t) \\ & \leq \left(\frac{(\bar{\lambda}/\underline{\lambda})^2}{N-1} - 1 \right) \underline{\lambda} \sigma (x_t) \\ & \leq \left(\frac{(\bar{\lambda}/\underline{\lambda})^2}{N-1} - 1 \right) (\underline{\lambda}/\bar{\lambda}) V_t (x_t; \hat{\alpha}_t) \end{aligned} \quad (40)$$

Then, for all $N > (\bar{\lambda}/\underline{\lambda})^2 + 1$, we have

$$V_{t+1} (\hat{x}_{t+1}; \hat{\alpha}_t) \leq \epsilon V_t (x_t; \hat{\alpha}_t), \quad (41)$$

where $\epsilon \triangleq 1 - \left(1 - \frac{(\bar{\lambda}/\underline{\lambda})^2}{N-1} \right) (\underline{\lambda}/\bar{\lambda}) < 1$. \square

E. Proof of Lemma 2

We first establish the Lipschitzness of the value function and boundedness of state by the following lemmas.

Lemma 3 (Lipschitzness of Value Function). *Under Assumption 2, the value function is Lipschitz in the initial condition x and the parameter $\hat{\alpha}$ over bounded domain sets, i.e., there exists a constant L such that $|V_t (x; \hat{\alpha}) - V_t (x'; \hat{\alpha}')| \leq L (\|x - x'\| + \|\hat{\alpha} - \hat{\alpha}'\|)$, $\forall x, x', \hat{\alpha}, \hat{\alpha}'$ such that $\|x - x'\| \leq R_1$ and $\|\hat{\alpha} - \hat{\alpha}'\| \leq R_2$ for some constants $R_1 \geq 0$ and $R_2 \geq 0$.*

Lemma 4 (Boundedness of State). *Under Assumption 1 and Assumption 5, the state x_t is bounded for all t .*

Then we have

$$\begin{aligned} & V_{t+1} (x_{t+1}; \hat{\alpha}_{t+1}) - V_{t+1} (\hat{x}_{t+1}; \hat{\alpha}_t) \\ & \leq L (\|\hat{x}_{t+1} - x_{t+1}\| + \|\hat{\alpha}_{t+1} - \hat{\alpha}_t\|) \\ & = L (\|\Phi (z_t) \hat{\alpha}_t - \Phi (z_t) \alpha\| + \|\hat{\alpha}_{t+1} - \hat{\alpha}_t\|) \\ & = L \left(\sqrt{l_t (\hat{\alpha}_t)} + \|\hat{\alpha}_{t+1} - \hat{\alpha}_t\| \right) \\ & \leq L \left(\sqrt{l_t (\hat{\alpha}_t)} + \|\eta \nabla_t\| \right), \end{aligned} \quad (42)$$

where the first inequality holds per Assumption 2, the first equality holds per Assumption 5, the second inequality holds by definitions of $l_t (\hat{\alpha}_t)$, and the second inequality holds by the Pythagorean theorem [59, Theorem 2.1]. \square

F. Proof of Lemma 3

Given that $c_t (x_t, u_t)$ and $x_{k+1} = f (x_k) + g (x_k) u_k + h (z_k)$ are locally Lipschitz in x_t , u_t , and $\hat{\alpha}$, we have that the objective function in eq. (5) is locally Lipschitz in the initial condition x , u_k and $\hat{\alpha}$ by substituting eq. (5b) into eq. (5a). Also, since \mathcal{U} is compact, the objective function is Lipschitz in $\mathcal{U} \times \dots \times \mathcal{U}$ and the value function is Lipschitz in x and $\hat{\alpha}$. Then the proof follows the same steps in the proof of [66, Theorem C. 29] \square

G. Proof of Lemma 4

Under Assumption 5, since $\Phi (\cdot)$, α , $\hat{\alpha}_t$ are uniformly bounded for all t , then $l_t (\hat{\alpha}_t)$ and ∇_t are bounded for all t . From eq. (25), we have:

$$\begin{aligned} & V_{t+1} (x_{t+1}; \hat{\alpha}_{t+1}) \\ & \leq \epsilon^t \bar{\lambda} \sigma (x_1) + \sum_{k=1}^t \epsilon^{t-k} L \left(\sqrt{l_k (\hat{\alpha}_k)} + \|\eta \nabla_k\| \right) \\ & \leq \epsilon^t \bar{\lambda} \sigma (x_1) + \frac{\mu}{1 - \epsilon} \\ & \leq \bar{\lambda} \sigma (x_1) + \frac{\mu}{1 - \epsilon}, \end{aligned} \quad (43)$$

where the second inequality holds with μ as the uniform upper bound of $L \left(\sqrt{l_t (\hat{\alpha}_t)} + \|\eta \nabla_t\| \right)$ for all t , and the third inequality holds by $\epsilon \in (0, 1)$.

Under Assumption 1, we have

$$\begin{aligned} \sigma (x_t) & \leq \frac{1}{\underline{\lambda}} c_t (x_t, u_t) \leq \frac{1}{\underline{\lambda}} V_t (x_t; \hat{\alpha}_t) \\ & \leq \frac{\bar{\lambda}}{\underline{\lambda}} \sigma (x_1) + \frac{\mu}{\underline{\lambda} (1 - \epsilon)}, \end{aligned} \quad (44)$$

where we obtain that $\sigma (x_t)$ is bounded by constant $\frac{\bar{\lambda}}{\underline{\lambda}} \sigma (x_1) + \frac{\mu}{\underline{\lambda} (1 - \epsilon)}$. Hence, x_t is bounded for all t . \square

H. Discussion on State Constraints

In the MPC problem in eq. (4), we can combine several approaches with the current framework to handle state constraints. For example:

- Safety filter with robust control barrier function [96]. Since this can arbitrarily modify the control and the resulting state, the no-regret guarantee (Theorem 1) can be difficult to derive.

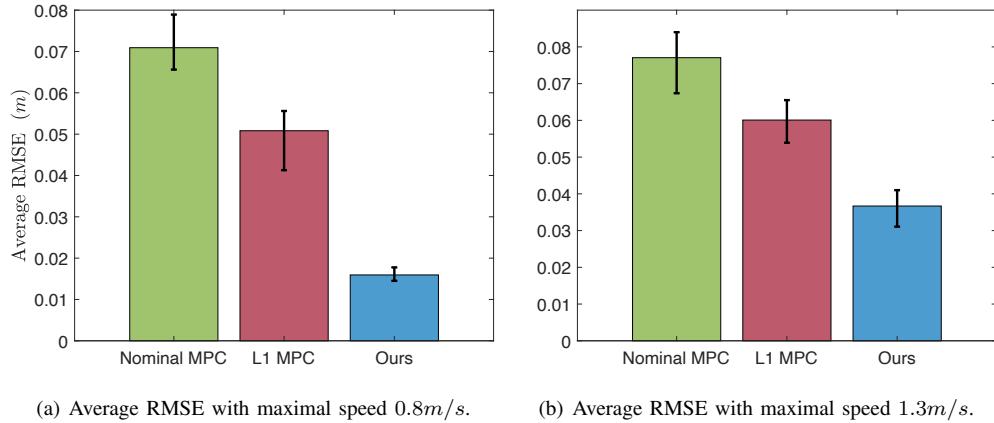


Fig. 13: **Tracking Performance Comparison for the Quadrotor Experiments in Section VI.** the quadrotor tracking the circular trajectory at maximal speeds $0.8m/s$ and $1.3m/s$ w/o adding ground effects and wind disturbances. The error bar represents the minimum and maximum RMSE. Algorithm 1 demonstrates improved performance compared to Nominal MPC and L1-MPC in terms of tracking error over all tested scenarios.

- Robust safety in robust MPC with an assumption of bounded $\|h(\cdot)\|$ [4], [97], [98]. If the Lipschitzness of the value function can be obtained [99, Theorem 4.14], we expect that Theorem 1 (no-regret guarantee) would still hold when the u_t in dynamic regret definition is obtained by solving MPC (eq. (4)) with tightened state constraints.
- Uncertainty quantification using the set-membership algorithm [100] for the case of time-invariant optimal parameter α^* . Under the assumption of bounded $\|h(\cdot)\|$, this method essentially prunes the uncertainty set $\mathcal{D}(B_h)$ at every time step using the observed $h(z_t)$, and the pruned uncertainty set \mathcal{D}_t is non-expanding. Then in line 9 of Algorithm 1, we need to project $\hat{\alpha}'_{i,t+1}$ onto \mathcal{D}_t , and the Proposition 2 remains the same. We can apply an approach similar to the robust safety method for deriving tightened state constraints and regret guarantees.

I. Additional Results of Hardware Experiments

In this section, we provide additional hardware experiments results to Section VI of the quadrotor tracking the circular trajectory at speeds $0.8m/s$ and $1.3m/s$ w/o ground effects and wind disturbances. In this case, the quadrotor suffers from unknown body drag, rotor drag, and turbulent effects caused by the propellers. The results are given in Figure 13. Algorithm 1 demonstrates improved performance compared to Nominal MPC and L1-MPC in terms of tracking error in both cases of maximal speeds.

REFERENCES

- [1] E. Ackerman, "Amazon promises package delivery by drone: Is it for real?" *IEEE Spectrum, Web*, 2013.
- [2] J. Chen, T. Liu, and S. Shen, "Tracking a moving target in cluttered environments using a quadrotor," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 446–453.
- [3] D. Seneviratne, L. Ciani, M. Catelani, D. Galar *et al.*, "Smart maintenance and inspection of linear assets: An industry 4.0 approach," *Acta Imeko*, vol. 7, pp. 50–56, 2018.
- [4] D. Q. Mayne, M. M. Seron, and S. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.
- [5] G. Goel and B. Hassibi, "Regret-optimal control in dynamic environments," *arXiv preprint:2010.10473*, 2020.
- [6] O. Sabag, G. Goel, S. Lale, and B. Hassibi, "Regret-optimal full-information control," *arXiv preprint:2105.01244*, 2021.
- [7] A. Martin, L. Furieri, F. Dörfler, J. Lygeros, and G. Ferrari-Trecate, "Safe control with minimal regret," in *Learning for Dynamics and Control Conference (L4DC)*, 2022, pp. 726–738.
- [8] A. Didier, J. Sieber, and M. N. Zeilinger, "A system level approach to regret optimal control," *IEEE Control Systems Letters (L-CSS)*, 2022.
- [9] H. Zhou and V. Tzoumas, "Safe control of partially-observed linear time-varying systems with minimal worst-case dynamic regret," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 8781–8787.
- [10] A. Martin, L. Furieri, F. Dörfler, J. Lygeros, and G. Ferrari-Trecate, "On the guarantees of minimizing regret in receding horizon," *IEEE Transactions on Automatic Control*, 2024.
- [11] J.-J. E. Slotine, "Applied nonlinear control," *PRENTICE-HALL google schola*, vol. 2, pp. 1123–1131, 1991.
- [12] M. Krstic, P. V. Kokotovic, and I. Kanellakopoulos, *Nonlinear and adaptive control design*. John Wiley & Sons, Inc., 1995.
- [13] P. A. Ioannou and J. Sun, *Robust adaptive control*. PTR Prentice-Hall Upper Saddle River, NJ, 1996, vol. 1.
- [14] N. M. Boffi, S. Tu, and J.-J. E. Slotine, "Regret bounds for adaptive nonlinear control," in *Learning for Dynamics and Control*. PMLR, 2021, pp. 471–483.
- [15] —, "Nonparametric adaptive control and prediction: Theory and randomized algorithms," *Journal of Machine Learning Research*, vol. 23, no. 281, pp. 1–46, 2022.
- [16] 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, 2020.
- [17] Z. Wu, S. Cheng, P. Zhao, A. Gahlawat, K. A. Ackerman, A. Lakshmanan, C. Yang, J. Yu, and N. Hovakimyan, "L1 quad: L1 adaptive augmentation of geometric control for agile quadrotors with performance guarantees," *arXiv preprint arXiv:2302.07208*, 2023.
- [18] E. Das and J. W. Burdick, "Robust control barrier functions using uncertainty estimation with application to mobile robots," *arXiv preprint arXiv:2401.01881*, 2024.
- [19] J. Jia, W. Zhang, K. Guo, J. Wang, X. Yu, Y. Shi, and L. Guo, "Evolver: Online learning and prediction of disturbances for robot control," *IEEE Transactions on Robotics*, 2023.
- [20] Y. Rahman, A. Xie, J. B. Hoagg, and D. S. Bernstein, "A tutorial and overview of retrospective cost adaptive control," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 3386–3409.
- [21] C. Sánchez-Sánchez and D. Izzo, "Real-time optimal control via deep neural networks: study on landing problems," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 5, pp. 1122–1135, 2018.
- [22] A. Carron, E. Arcari, M. Wermelinger, L. Hewing, M. Hutter, and M. N. Zeilinger, "Data-driven model predictive control for trajectory

tracking with a robotic arm,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3758–3765, 2019.

[23] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, “Data-driven MPC for quadrotors,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3769–3776, 2021.

[24] L. Hewing, J. Kabzan, and M. N. Zeilinger, “Cautious model predictive control using gaussian process regression,” *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2019.

[25] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.

[26] F. Ramos, R. C. Possas, and D. Fox, “Bayessim: adaptive domain randomization via probabilistic inference for robotics simulators,” *arXiv preprint arXiv:1906.01728*, 2019.

[27] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.

[28] B. O. Koopman, “Hamiltonian systems and transformation in hilbert space,” *Proceedings of the National Academy of Sciences*, vol. 17, no. 5, pp. 315–318, 1931.

[29] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Sparse identification of nonlinear dynamics with control (sindyc),” *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 710–715, 2016.

[30] I. Abraham, G. De La Torre, and T. D. Murphey, “Model-based control using koopman operators,” *arXiv preprint arXiv:1709.01568*, 2017.

[31] D. Bruder, X. Fu, R. B. Gillespie, C. D. Remy, and R. Vasudevan, “Data-driven control of soft robots using koopman operator theory,” *IEEE Transactions on Robotics*, vol. 37, no. 3, pp. 948–961, 2020.

[32] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.

[33] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, “Information theoretic MPC for model-based reinforcement learning,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 1714–1721.

[34] 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,” *arXiv preprint arXiv:1803.11347*, 2018.

[35] S. Belkhale, 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, no. 2, pp. 1471–1478, 2021.

[36] 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, 2019, pp. 9784–9790.

[37] 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.

[38] A. Saviolo, J. Frey, A. Rathod, M. Diehl, and G. Loianno, “Active learning of discrete-time dynamics for uncertainty-aware model predictive control,” *IEEE Transactions on Robotics*, 2023.

[39] E. Hazan and K. Singh, “Introduction to online nonstochastic control,” *arXiv preprint arXiv:2211.09619*, 2022.

[40] N. Agarwal, B. Bullins, E. Hazan, S. Kakade, and K. Singh, “Online control with adversarial disturbances,” in *International Conference on Machine Learning (ICML)*, 2019, pp. 111–119.

[41] P. Zhao, Y.-X. Wang, and Z.-H. Zhou, “Non-stationary online learning with memory and non-stochastic control,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2022, pp. 2101–2133.

[42] P. Gradu, E. Hazan, and E. Minasyan, “Adaptive regret for control of time-varying dynamics,” *arXiv preprint:2007.04393*, 2020.

[43] H. Zhou, Z. Xu, and V. Tzoumas, “Efficient online learning with memory via frank-wolfe optimization: Algorithms with bounded dynamic regret and applications to control,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 8266–8273.

[44] H. Zhou and V. Tzoumas, “Safe non-stochastic control of linear dynamical systems,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 5033–5038.

[45] H. Zhou, Y. Song, and V. Tzoumas, “Safe non-stochastic control of control-affine systems: An online convex optimization approach,” *IEEE Robotics and Automation Letters*, 2023.

[46] A. Krause and C. Ong, “Contextual gaussian process bandit optimization,” *Advances in neural information processing systems*, vol. 24, 2011.

[47] F. Berkenkamp, A. P. Schoellig, and A. Krause, “Safe controller optimization for quadrotors with gaussian processes,” in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 491–496.

[48] S. R. Chowdhury and A. Gopalan, “On kernelized multi-armed bandits,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 844–853.

[49] F. Djeumou, A. P. Vinod, E. Goubault, S. Putot, and U. Topcu, “On-the-fly control of unknown systems: From side information to performance guarantees through reachability,” *IEEE Transactions on Automatic Control*, vol. 68, no. 8, pp. 4857–4872, 2022.

[50] A. P. Vinod, A. Israel, and U. Topcu, “On-the-fly control of unknown nonlinear systems with sublinear regret,” *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2729–2742, 2022.

[51] M. Valko, N. Korda, R. Munos, I. Flaounas, and N. Cristianini, “Finite-time analysis of kernelised contextual bandits,” *arXiv preprint arXiv:1309.6869*, 2013.

[52] H. A. Pierson and M. S. Gashler, “Deep learning in robotics: a review of recent research,” *Advanced Robotics*, vol. 31, no. 16, pp. 821–835, 2017.

[53] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2018.

[54] Y. S. Abu-Mostafa, M. Magdon-Ismail, and H.-T. Lin, *Learning from data*. AMLBook New York, 2012, vol. 4.

[55] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice hall Upper Saddle River, NJ, 1998, vol. 104.

[56] E. Lavretsky and K. A. Wise, “Robust adaptive control,” in *Robust and adaptive control: With aerospace applications*. Springer, 2012, pp. 317–353.

[57] S. Sun, A. Romero, P. Foehn, E. Kaufmann, and D. Scaramuzza, “A comparative study of nonlinear MPC and differential-flatness-based control for quadrotor agile flight,” *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3357–3373, 2022.

[58] F. Cucker and S. Smale, “On the mathematical foundations of learning,” *Bulletin of the American mathematical society*, vol. 39, no. 1, pp. 1–49, 2002.

[59] E. Hazan *et al.*, “Introduction to online convex optimization,” *Foundations and Trends in Optimization*, vol. 2, no. 3–4, pp. 157–325, 2016.

[60] A. Rahimi and B. Recht, “Random features for large-scale kernel machines,” *Advances in neural information processing systems*, vol. 20, 2007.

[61] ———, “Uniform approximation of functions with random bases,” in *2008 46th annual allerton conference on communication, control, and computing*. IEEE, 2008, pp. 555–561.

[62] S. Lale, K. Azizzadenesheli, B. Hassibi, and A. Anandkumar, “Model learning predictive control in nonlinear dynamical systems,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 757–762.

[63] A. Karapetyan, E. C. Balta, A. Tsiamis, A. Iannelli, and J. Lygeros, “On the regret of recursive methods for discrete-time adaptive control with matched uncertainty,” *arXiv preprint arXiv:2404.02023*, 2024.

[64] M. Nonhoff, J. Köhler, and M. A. Müller, “Online convex optimization for constrained control of nonlinear systems,” *arXiv preprint arXiv:2412.00922*, 2024.

[65] D. Muthirayan, J. Yuan, D. Kalathil, and P. P. Khargonekar, “Online learning for predictive control with provable regret guarantees,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 6666–6671.

[66] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: Theory, computation, and design*. Nob Hill Publishing, 2017, vol. 2.

[67] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

[68] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, “Model predictive control: for want of a local control lyapunov function, all is not lost,” *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 546–558, 2005.

[69] Y. Li, S. Das, and N. Li, “Online optimal control with affine constraints,” in *AAAI Conference on Artificial Intelligence (AAAI)*, vol. 35, no. 10, 2021, pp. 8527–8537.

[70] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi, “Tube-based robust nonlinear model predictive control,” *International journal of robust and nonlinear control*, vol. 21, no. 11, pp. 1341–1353, 2011.

[71] J. Köhler, R. Soloperto, M. A. Müller, and F. Allgöwer, “A computationally efficient robust model predictive control framework for uncertain nonlinear systems,” *IEEE Transactions on Automatic Control*, vol. 66, no. 2, pp. 794–801, 2020.

[72] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, “Data-driven model predictive control with stability and robustness guarantees,” *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1702–1717, 2020.

[73] R. Brault, M. Heinonen, and F. Buc, "Random fourier features for operator-valued kernels," in *Asian Conference on Machine Learning*. PMLR, 2016, pp. 110–125.

[74] H. Q. Minh, "Operator-valued bochner theorem, fourier feature maps for operator-valued kernels, and vector-valued learning," *arXiv preprint arXiv:1608.05639*, 2016.

[75] F. Bach, "Breaking the curse of dimensionality with convex neural networks," *Journal of Machine Learning Research*, vol. 18, no. 19, pp. 1–53, 2017.

[76] C. Carmeli, E. De Vito, A. Toigo, and V. Umanitá, "Vector valued reproducing kernel hilbert spaces and universality," *Analysis and Applications*, vol. 8, no. 01, pp. 19–61, 2010.

[77] M. Nonhoff and M. A. Müller, "On the relation between dynamic regret and closed-loop stability," *arXiv preprint arXiv:2209.05964*, 2022.

[78] Y. Bengio, N. Roux, P. Vincent, O. Delalleau, and P. Marcotte, "Convex neural networks," *Advances in neural information processing systems*, vol. 18, 2005.

[79] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Interna. Conf. on Machine Learning (ICML)*, 2003, pp. 928–936.

[80] R. V. Florian, "Correct equations for the dynamics of the cart-pole system," *Center for Cognitive and Neural Studies (Coneural), Romania*, p. 63, 2007.

[81] P. E. Kloeden, E. Platen, P. E. Kloeden, and E. Platen, *Stochastic differential equations*. Springer, 1992.

[82] Z. Yuan, A. W. Hall, S. Zhou, L. Brunke, M. Greeff, J. Panerati, and A. P. Schoellig, "Safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11142–11149, 2022.

[83] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.

[84] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, pp. 1–36, 2019.

[85] J. Quinonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate gaussian process regression," *The Journal of Machine Learning Research*, vol. 6, pp. 1939–1959, 2005.

[86] E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," *Advances in neural information processing systems*, vol. 18, 2005.

[87] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, *Robot Operating System (ROS): The Complete Reference (Volume 1)*. Cham: Springer International Publishing, 2016, ch. RotorS—A Modular Gazebo MAV Simulator Framework, pp. 595–625. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-26054-9_23

[88] P. Martin and E. Salaün, "The true role of accelerometer feedback in quadrotor control," in *2010 IEEE international conference on robotics and automation*. IEEE, 2010, pp. 1623–1629.

[89] M. Faessler, D. Falanga, and D. Scaramuzza, "Thrust mixing, saturation, and body-rate control for accurate aggressive quadrotor flight," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 476–482, 2016.

[90] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 620–626, 2017.

[91] R. Verschueren, G. Frison, D. Kouzoupis, J. Frey, N. v. Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados—a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, no. 1, pp. 147–183, 2022.

[92] D. C. Youvan, "Developing and deploying ai applications on nvidia jetson orin nx: A comprehensive guide," 2024.

[93] "Px4 autopilot: Open source autopilot for drones." [Online]. Available: <https://docs.px4.io/main/en/index.html>

[94] "Mavros: Communication driver for autopilots with mavlink communication protocol." [Online]. Available: <https://wiki.ros.org/mavros#Overview>

[95] A. Tsiamis, A. Karapetyan, Y. Li, E. C. Balta, and J. Lygeros, "Predictive linear online tracking for unknown targets," *arXiv preprint arXiv:2402.10036*, 2024.

[96] M. Jankovic, "Robust control barrier functions for constrained stabilization of nonlinear systems," *Automatica*, vol. 96, pp. 359–367, 2018.

[97] P. Bouffard, A. Aswani, and C. Tomlin, "Learning-based model predictive control on a quadrotor: Onboard implementation and experimental results," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 279–284.

[98] A. Aswani, P. Bouffard, and C. Tomlin, "Extensions of learning-based model predictive control for real-time application to a quadrotor helicopter," in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 4661–4666.

[99] S. Dempe, *Foundations of bilevel programming*. Springer Science & Business Media, 2002.

[100] M. Milanese and A. Vicino, "Optimal estimation theory for dynamic systems with set membership uncertainty: An overview," *Automatica*, vol. 27, no. 6, pp. 997–1009, 1991.



Hongyu Zhou received the B.Eng. degree in Naval Architecture and Ocean Engineering from the Huazhong University of Science and Technology in 2018, and the M.Sc. degree in Marine Technology from the Norwegian University of Science and Technology in 2020. He received M.Sc. degree in Robotics from the University of Michigan in 2025. He is currently pursuing a Ph.D. degree in Aerospace Engineering at the University of Michigan. Hongyu works on algorithms to achieve agile, efficient, and safe robot motion control in uncertain and dynamic environments, with formal performance guarantees. His research interests include optimal control, adaptive control, model learning for control, and online learning.



Vasileios Tzoumas (IEEE Senior Member) received his Ph.D. in Electrical and Systems Engineering at the University of Pennsylvania (2018). He holds a Master of Arts in Statistics from the Wharton School of Business at the University of Pennsylvania (2016), a Master of Science in Electrical Engineering from the University of Pennsylvania (2016), and a diploma in Electrical and Computer Engineering from the National Technical University of Athens (2012). Vasileios is as an Assistant Professor in the Department of Aerospace Engineering, University of Michigan, Ann Arbor. Previously, he was at the Massachusetts Institute of Technology (MIT), in the Department of Aeronautics and Astronautics, and in the Laboratory for Information and Decision Systems (LIDS), where he was a research scientist (2019–2020) and a post-doctoral associate (2018–2019). Vasileios works on algorithms and innovative hardware for scalable and reliable cyber-physical systems in resource-constrained, uncertain, and contested environments via resource-aware decision-making, online learning, and resilient adaptation. Vasileios is a recipient of an NSF CAREER Award, the Best Paper Award in Robot Vision at the 2020 IEEE International Conference on Robotics and Automation (ICRA), an Honorable Mention from the 2020 IEEE Robotics and Automation Letters (RA-L), and was a Best Student Paper Award finalist at the 2017 IEEE Conference in Decision and Control (CDC).