

Adaptive Robust Model Predictive Control via Uncertainty Cancellation

Rohan Sinha, James Harrison, Spencer M. Richards, and Marco Pavone

Abstract—We propose a learning-based robust predictive control algorithm that compensates for significant uncertainty in the dynamics for a class of discrete-time systems that are nominally linear with an additive nonlinear component. Such systems are commonly used to model the nonlinear effects of an unknown environment on a nominal linear system. Inspired by certainty equivalent “estimate-and-cancel” control laws pioneered in classical adaptive control, we optimize over a class of nonlinear feedback policies to significantly improve performance in the presence of uncertainties of large magnitude, a setting in which existing learning-based predictive control algorithms often struggle to guarantee safety. In contrast to previous work in robust adaptive model predictive control, our approach allows us to take advantage of structure (i.e., the numerical predictions) in the a priori unknown dynamics learned online through function approximation. Our approach also extends typical nonlinear adaptive control methods to systems with state and input constraints even when we cannot directly cancel the additive uncertain function from the dynamics. Moreover, we apply contemporary statistical estimation techniques to certify the system’s safety in the form of persistent constraint satisfaction with high probability. Finally, we show in simulation that our method can accommodate more significant unknown dynamics terms than existing methods.

Index Terms—Adaptive Control, Machine Learning, Meta-Learning, Model Predictive Control, Robust Control

I. INTRODUCTION

Developing control systems capable of autonomous operation in diverse, unstructured environments requires control algorithms that learn from experience. Therefore, rapid advances in machine learning algorithms (e.g., see [2], [3]) have driven a concomitant explosion in research on the use of learning algorithms to control dynamical systems (e.g., see [4]–[22]). These *learning-based* control algorithms leverage operational data to improve closed-loop performance, typically by refining estimates of the environments’ nonlinear and a priori unknown effects on the dynamics online.

The authors are with the Autonomous Systems Lab at Stanford University, Stanford, CA, {rhnsinha, jharrison, spenrich, pavone}@stanford.edu. This research was supported in part by the National Science Foundation (NSF) via Cyber-Physical Systems (CPS) award #1931815, and the National Aeronautics and Space Administration (NASA) via University Leadership Initiative grant #80NSSC20M0163 and via an Early Stage Innovations grant. Spencer M. Richards and James Harrison were also supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). This article solely reflects our own opinions and conclusions, and not those of any NSF, NASA, or NSERC entity. This paper extends an earlier conference version of this work [1].




Fig. 1. High-level illustration of our approach. We learn the effect of environmental uncertainty online and curb its influence through certainty equivalent adaptive control. Then, we account for any remaining uncertainty and guarantee constraint satisfaction using tube MPC.

Upon deployment, these methods should provide rigorous safety guarantees while quickly adapting in the face of uncertainty. The problem of safe learning in control was initially considered in adaptive control theory, a mature discipline that has historically emphasized safety in the form of closed-loop stability guarantees [23]–[25]. These classical adaptive control methods have seen continued interest in recent work on learning to control for their simplicity and asymptotic convergence behavior for systems with matched uncertainty (e.g., see [4], [10], [18], [26]–[28]). However, these approaches generally cannot guarantee the satisfaction of constraints on states and inputs, even though this finer-grained notion of safety is essential in practice to ensure unsafe regions of the state space are avoided under limits on the control authority.

In contrast, model predictive control (MPC) algorithms were developed to provide such set-avoidance guarantees [29]. Therefore, many recently proposed learning-based control algorithms strive to integrate online learning algorithms with robust MPC strategies (e.g., see [5]–[8], [11], [13]–[15], [20], [30], [31]). However, guaranteeing constraint satisfaction then generally requires estimating uncertainty in learned quantities and propagating this uncertainty forward in time to characterize the set (or distribution) of possible trajectories when performing trajectory optimization. This is a challenging problem when we estimate the dynamics with expressive models like deep neural networks. Thus, we reach a central tension underlying modern learning-based control; we desire complex nonlinear models with the broad representational capacity necessary for autonomy in diverse and a priori unmodelled environments, but these models are not readily integrated into constrained control algorithms. This tension often results in

learning-based control algorithms for constrained systems that are either too conservative (e.g., yielding limited performance to remain safe) or too fragile (e.g., infeasible in the face of considerable uncertainty).

In this work, we leverage classical adaptive control techniques to reduce the over-conservatism and fragility of learning-based control algorithms for systems subject to state and input constraints. We combine a simple nonlinear control law inspired by “estimate-and-cancel” methods in nonlinear adaptive control with robust MPC techniques to control a system in an uncertain environment, represented as an unknown nonlinear term in the dynamics. This strategy results in a simple control algorithm that is recursively feasible, input-to-state stable, and can safely leverage expressive nonlinear models. We can view our approach through the lens of both adaptive control and robust adaptive MPC. On the one hand, we extend classical adaptive cancellation-based methods to a setting with uncertain, unmatched dynamics subject to state and input constraints. On the other hand, we introduce a simple nonlinear feedback law to construct an adaptive robust MPC scheme that can reduce the conservatism of existing approaches by taking advantage of the learned structure in a priori unknown dynamics. We demonstrate on various simulated systems that our method reduces the conservatism and increases the feasible domain of the resulting robust MPC problem compared to typical adaptive robust MPC methods.

A. Related Work

We briefly review two significant paradigms for the control of uncertain systems, namely *adaptive control* and *robust control*. We then discuss recent works that combine ideas from both paradigms, oftentimes leveraging modern methods in machine learning.

1) *Adaptive Control*: Adaptive control concerns the joint design of a parametric feedback controller and a parameter adaptation law to improve closed-loop performance over time when the dynamics are partially unknown [24], [25]. Design of these components for nonlinear systems commonly relies on expressing unknown dynamics terms as linear combinations of *known* basis functions, i.e., *features* [24]. The adaptation law updates the feature weights online, and the controller applies part of the control signal to cancel the estimated term from the dynamics [23], [24], [32], [33]. These simple methods can achieve tracking convergence up to an error threshold that depends on the representation capacity of the features relative to the true dynamics [32], [33]. Recent works propose combining high-capacity parametric and non-parametric models from machine learning with classical adaptive control designs. This includes deep neural networks via online back-propagation [10], Gaussian processes [26], and Bayesian neural networks [18], [34] via online Bayesian updates and meta-learned features [4], [19]. However, these approaches are fundamentally limited by common assumptions in classical adaptive control, namely that uncertain dynamics terms can be stably canceled by the control input in their entirety, i.e., that these terms are *matched uncertainties* [24], [25], [32], [33]. Moreover, most of these works do not consider state and input constraints, which are essential to safe control in practice. We generalize these

classical adaptive methods to incorporate safety constraints even if the uncertainty is not fully matched.

2) *Robust Control*: Robust control seeks consistent performance despite uncertainty in the dynamics. In this work, we consider the robust control of constrained discrete-time systems using tools from predictive control. In particular, robust MPC algorithms for linear systems consider the control of a system subject to bounded noise or uncertain dynamics terms, i.e., *disturbances*, as an optimization program with explicit state and input constraints. Some methods optimize the worst-case performance of the controller [35], while others tighten the constraints to accommodate the set of all possible trajectories induced by the disturbances and optimize the nominal predicted trajectory instead [36], [37]. To account for future information gain and reduce conservatism, these methods either fix a disturbance feedback policy [36] or optimize over state feedback policies [38].

3) *Adaptive Robust MPC (ARMPC)*: ARMPC, often referred to as learning-based MPC, incorporates the online estimation (i.e., learning) from adaptive control methods into robust MPC to satisfy constraints in the presence of process noise and model uncertainty during learning. Recent years have seen a flurry of work on nonlinear predictive control methods that apply contemporary machine learning techniques to learn uncertain dynamics online [11], [13], [20], [22]. These methods typically result in non-convex programs for trajectory optimization under the learned dynamics, while relying on conservative approximate methods for uncertainty propagation to guarantee constraint satisfaction. However, it is unclear how to construct the necessary components, i.e., the robust positive invariant and the terminal cost function, for predictive control to make claims of persistent constraint satisfaction (i.e., safety) or stability for arbitrary nonlinear systems. Some methods ignore these topics and do not make rigorous safety guarantees [11]. Other works, such as [20], [22], assume these ingredients already exist or only consider trajectory optimization tasks where a goal region needs to be reached in a finite number of time steps [13]. Moreover, iterative methods used to solve for local minima of non-convex programs can be computationally prohibitive and often have limited performance guarantees.

To make rigorous safety guarantees, we will focus on adaptive robust methods for systems that are nominally linear, as considered in [5]–[8], [14], [15], [31]. A straightforward approach is to maintain an outer bound on any unknown nonlinear terms in the dynamics and use it as a disturbance bound in any chosen robust MPC scheme [5]–[8], [14]. These methods avoid some of the difficulties associated with trajectory optimization for nonlinear dynamics by ignoring the actual values of the nonlinear terms at any point in the state space. That is, these methods do not exploit the learned structure in the a priori unknown dynamics, often rendering them overly conservative or fragile.

B. Contributions

We present an ARMPC method for systems with an additive unknown nonlinear dynamics term, subject to state and input constraints. Rather than construct an outer envelope for such terms, as is normative in ARMPC literature for linear systems

[5], [30], we develop theoretical guarantees for a broad class of *function approximators*, including set membership and least-squares methods for certain noise models. Our key idea is to decompose uncertain dynamics terms into a *matched* component that lies in a subspace that can be stably canceled by the control input, and an *unmatched* component that lies in an orthogonal complement to this subspace. We apply certainty equivalent adaptive control techniques to stably cancel the *matched* component from the dynamics and then apply robust MPC, considering the unmatched component as a bounded disturbance. Therefore, our method explicitly uses estimates of the unknown dynamics term throughout the state space for control, i.e., it takes advantage of the learned structure in the dynamics. We prove our method is recursively feasible and input-to-state stable. Moreover, we demonstrate on various simulated systems that our method reduces the conservatism and increases the feasible domain of the resulting robust MPC problem compared to typical adaptive robust MPC methods.

The performance of the adaptive control strategy, which learns a structured representation of the unknown dynamics term, relies on the quality of the features used in online learning. Thus, in addition to investigating standard techniques for feature construction, we introduce a Bayesian meta-learning algorithm [34] for feature learning. This method, which learns features that are broadly useful across tasks (or instantiations of unknown dynamics) produces useful features and well-calibrated priors, and satisfies our desiderata for learning algorithms. Beyond the central algorithmic contributions of this paper, we also show the utility of these meta-learning algorithms in adaptive and learning-based control.

C. Organization

In §II, we pose the general robust infinite-horizon optimal control problem central to this work. In §III, we discuss a standard robust MPC solution to this problem that is a core tool in the development of our approach. Moreover, we outline several tools that are used in the analysis of our approach. In §IV, we describe our approach, prove the stability of the method, and discuss different possible assumptions on the learning setting and the impact on the controller. In §V, we discuss two estimators that we use in our experiments – set-membership and linear regression. We also compare these estimators, and introduces meta-learning as a powerful black-box tool for feature learning and prior calibration. Finally, we present simulation results in §VI, and we conclude the paper and discuss directions for future work in §VII.

II. PROBLEM FORMULATION

We consider the robust control of nonlinear discrete-time systems of the form

$$x(t+1) = Ax(t) + Bu(t) + f(x(t)) + v(t), \quad (1)$$

where $x(t) \in \mathbb{R}^n$ is the system state, $u(t) \in \mathbb{R}^m$ is the control input, $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ are known constant matrices, and $v(t) \in \mathcal{V}$ is a disturbance in a known compact convex set \mathcal{V} containing the origin. In addition, an *unknown, nonlinear* dynamics term $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ acts on the system, representing the unmodelled influence of the environment on

the nominally linear dynamics of system (1). For example, $f(x)$ can model the effect that wind conditions have on the linearized dynamics of a drone. We assume the disturbances have zero mean and are independent and identically distributed (iid) according to some distribution $p(v)$, i.e., $v(t) \stackrel{\text{iid}}{\sim} p(v)$ and $\mathbb{E}[v(t)] = 0$ for all $t \geq 0$. Our goal is to regulate the system to the origin according to the robust optimal control problem

$$\begin{aligned} & \underset{x, u}{\text{minimize}} \quad \mathbb{E} \left[\sum_{t=0}^{\infty} h(x(t), u(t)) \right] \\ & \text{subject to} \quad x(t+1) = Ax(t) + Bu(t) + f(x(t)) + v(t), \\ & \quad \quad \quad u(t) \in \mathcal{U}, x(t) \in \mathcal{X}, v(t) \in \mathcal{V}, \forall t \in \mathbb{N}_{\geq 0} \end{aligned} \quad (2)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{U} \subseteq \mathbb{R}^m$ are compact convex sets containing the origin, and $h(x, u) = x^\top Qx + u^\top Ru$ is a quadratic stage cost parameterized by positive semi-definite matrix $Q \in \mathbb{S}_{\geq 0}^n$ and positive-definite matrix $R \in \mathbb{S}_{> 0}^m$. The problem (2) is computationally intractable to solve because the horizon is infinite and the nonlinear function f makes the problem non-convex. To approximately solve (2), we need additional assumptions on the unknown, nonlinear dynamics term f . In particular, to derive a controller that is robust to any possible value of $f(x)$, we need f to be bounded on \mathcal{X} . Moreover, to construct guarantees on the online estimation of f and establish properties of a controller using this estimate, we also need to assume some structure of f . For these reasons, we make the following assumption.

Assumption 1 (structure): The nonlinear dynamics term $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is linearly parameterizable, i.e.,

$$f(x) = W\phi(x), \quad \forall x \in \mathbb{R}^n, \quad (3)$$

where $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ is a *known* nonlinear feature map, and $W \in \mathbb{R}^{n \times d}$ is an *unknown* weight matrix. Moreover, $\|\phi(x)\| \leq 1$ for any $x \in \mathcal{X}$, where $\|\cdot\|$ is the Euclidean norm.

Representing a nonlinear function using a feature map is common both in adaptive control [24], [32] and contemporary machine learning [21], [34], as they can represent arbitrary functions if properly designed. Without loss of generality, we assume the upper norm bound on the features is one for simplicity. This parameterization admits function classes such as neural networks with scaled sigmoid outputs.

A. Matched and Unmatched Uncertainty

While it is common in adaptive control to assume the uncertain function f in (1) can be stably cancelled in its entirety [24], [33], we will generalize this approach to a setting where perfect cancellation is not possible. We use the following definition to distinguish between components of the uncertain dynamics f that can and cannot be cancelled.

Definition 1 (matched and unmatched uncertainty): The uncertain function $f(x)$ in (1) is a *matched uncertainty* if $f(x) \in \text{Range}(B)$ for all $x \in \mathcal{X}$. Conversely, if there exists an $x \in \mathcal{X}$ such that $f(x) \notin \text{Range}(B)$, then $f(x)$ is an *unmatched uncertainty*.

In this work, we assume the matrix B in (1) has full column rank, i.e., there are no redundant actuators; this guarantees that the Moore-Penrose pseudoinverse $B^\dagger := (B^\top B)^{-1} B^\top$ exists.

If $f(x)$ is a matched uncertainty, then the function $g(x) = B^\dagger f(x)$ satisfies $f(x) = Bg(x)$ for any $x \in \mathcal{X}$.

Controlling systems with matched uncertainty is a classical problem in the adaptive control literature, much of which relies on the observation that setting $u(t) = \bar{u}(t) - g(x(t))$ in (1) would cancel the nonlinear term to yield linear dynamics with respect to the nominal input $\bar{u}(t)$. *Certainty equivalent* controllers approximately cancel $g(x)$ with an estimate $\hat{g}(x)$ and can yield simple nonlinear adaptive laws that achieve asymptotic tracking performance for matched systems. Even though systems are often designed to be easy to control, unmatched uncertainty affects many practical systems of interest, such as underactuated robots (e.g., quadrotors and cars). We propose to decompose the uncertain function f into matched and unmatched components, apply certainty equivalent cancellation to the matched component, and curb the impact of the unmatched component with robust MPC. Applying part of the input to cancel matched uncertainty instantaneously prevents part of f from leaking into the dynamics, avoiding the need to react to large observed disturbances.

III. ROBUST MPC BACKGROUND

We now briefly review how we could approximately solve the optimal control problem (2) with existing robust MPC techniques that treat $d(t) := f(x(t)) + v(t)$ as a single bounded disturbance term, since $d(t)$ lies in the set

$$\mathcal{D} := \{f(x) + v \in \mathbb{R}^n \mid x \in \mathcal{X}, v \in \mathcal{V}\} \quad (4)$$

for all $t \in \mathbb{N}_{\geq 0}$. Indeed, \mathcal{D} is bounded under Assumption 1 and the boundedness of \mathcal{X} and \mathcal{V} .

A. Receding Horizon Control

We focus on receding-horizon robust MPC schemes, whereby an approximate version of (2) with finite horizon $N \in \mathbb{N}_{>0}$ is solved online with full state feedback. Rather than search over open-loop input sequences, which can incur issues with feasibility and stability, we search over closed-loop feedback policies [8], [36], [38]. In particular, we follow [38] in optimizing over time-varying, causal, affine disturbance feedback policies of the form

$$u_{t+k|t} = \bar{u}_{t+k|t} + \sum_{j=0}^{k-1} K_{k,j|t} d_{t+j|t}, \quad (5)$$

via the robust MPC problem

$$\begin{aligned} & \underset{\substack{\{K_{k,j|t}\}_{k=0, j=0}^{N-1, k-1}, \\ \{\bar{u}_{t+k|t}\}_{k=0}^{N-1}}}{\text{minimize}} & & V_N(\bar{x}_{t+N|t}) + \sum_{k=0}^{N-1} h(\bar{x}_{t+k|t}, \bar{u}_{t+k|t}) \\ & \text{subject to} & & \bar{x}_{t+k+1|t} = A\bar{x}_{t+k|t} + B\bar{u}_{t+k|t} \\ & & & x_{t+k+1|t} = Ax_{t+k|t} + Bu_{t+k|t} + d_{t+k|t} \\ & & & u_{t+k|t} = \bar{u}_{t+k|t} + \sum_{j=0}^{k-1} K_{k,j|t} d_{t+j|t} \\ & & & x_{t+k|t} \in \mathcal{X}, u_{t+k|t} \in \mathcal{U} \\ & & & \forall k \in \{0, 1, \dots, N-1\} \\ & & & \bar{x}_{t|t} = x(t), x_{t|t} = x(t), x_{t+N|t} \in \mathcal{O} \\ & & & \forall \{d_{t+k|t}\}_{k=0}^{N-1} \subset \mathcal{D} \end{aligned} \quad (6)$$

The problem (6) optimizes a time-varying feedback policy with a cost on the nominal trajectory (\bar{x}, \bar{u}) subject to state

and input constraints on the realized trajectory (x, u) . We use the subscript $t+k|t$ for quantities at the k -th step of the prediction horizon when (6) is solved online at time $t \in \mathbb{N}_{\geq 0}$. If the function $V_N : \mathcal{X} \rightarrow \mathbb{R}$, the terminal set \mathcal{O} , and the disturbance set \mathcal{D} are convex, then (6) is a convex problem; we refer readers to [38] for implementation details. One might also consider a formulation of (6) where the feedback gains are fixed, yielding a more basic tube MPC problem [36].

Solving (6) in a receding horizon fashion encodes the closed-loop feedback policy $u^* : \mathcal{X} \times \mathbb{N}_{\geq 0} \rightarrow \mathcal{U}$, where $u^*(x(t), t) = u_{t|t}^*$ is the first element of an optimal control input sequence from solving (6) with the initial condition $(x(t), t)$. This feedback policy is *robust* since the constraints in (6) are enforced for every possible N -step sequence of disturbances.

B. Invariant Sets

The choice of terminal ingredients V_N and \mathcal{O} in (6) is pivotal to guarantee that the closed-loop system formed by the dynamics (1) and the MPC policy satisfies $x(t) \in \mathcal{X}$, $u(t) \in \mathcal{U}$ for all $t \geq 0$, and is stable. In particular, to establish these recursive feasibility and stability guarantees, the terminal set \mathcal{O} must be invariant with respect to the underlying dynamics (1) under some policy that satisfies the input constraints on \mathcal{O} . Moreover, the terminal cost V_N must be a Lyapunov function associated with the stage cost on \mathcal{O} under the policy associated with \mathcal{O} . We review some invariant set notions below.

Definition 2 (invariant sets [29]): Consider the dynamical system $x(t+1) = q(x(t), v(t))$, where $v(t)$ is a disturbance signal that takes values in some set \mathcal{V} , subject to the state constraint set $\mathcal{X} \subseteq \mathbb{R}^n$. Then a *robust positive invariant (RPI) set* for the constrained system is any set $\mathcal{O} \subseteq \mathcal{X}$ satisfying

$$x(0) \in \mathcal{O} \implies x(t) \in \mathcal{O}, \forall t \in \mathbb{N}_{\geq 0}, \quad (7)$$

for any possible sequence $\{v(t) \mid t \in \mathbb{N}_{\geq 0}\} \subset \mathcal{V}$. The *maximal RPI set* $\mathcal{O}_\infty \subseteq \mathcal{X}$ is the RPI set satisfying $\mathcal{O} \subseteq \mathcal{O}_\infty$ for any other RPI set $\mathcal{O} \subseteq \mathcal{X}$.

Invariant sets are essential to predictive control design, as planning a trajectory into an RPI set associated with a fixed stabilizing feedback policy guarantees that there exists a robust MPC policy that satisfies the constraints for all time. Moreover, computing the maximal RPI set for a linear time-invariant system is algorithmically straightforward [29].

C. ISS Stability

Due to the disturbance term $v(t)$, the system (1) typically cannot be regulated to the origin even asymptotically. Therefore, we briefly review relevant results of input-to-state stability (ISS) theory, which is often used to analyze robust control algorithms [7], [37], [38]. We propose an adaptive approach that refines an estimate of the unknown function f online. As a result, the closed-loop system is time-varying. First, we review standard comparison function notation [39]. A function $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a class- \mathcal{K} function if it is continuous, strictly increasing, and $\alpha(0) = 0$. In addition, α is class- \mathcal{K}_∞ if it is class- \mathcal{K} and $\lim_{x \rightarrow \infty} \alpha(x) = \infty$. A function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{N}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is class- \mathcal{KL} if $\beta(\cdot, t)$ is class- \mathcal{K} for any fixed $t \geq 0$, $\beta(x, \cdot)$ is decreasing for any fixed $x \geq 0$, and

$\lim_{t \rightarrow \infty} \beta(x, t) = 0$ for any fixed $x \geq 0$. We now use these function classes to state the ISS definitions.

Definition 3 (input-to-state stable (ISS) [39]): The system $x(t+1) = q(t, x(t), v(t))$ with disturbance $v(t)$ is globally *input-to-state stable (ISS)* if there exists a class- \mathcal{KL} function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{N}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and a class- \mathcal{K} function $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\|x(t)\| \leq \beta(\|x(0)\|, t) + \gamma(\sup_{k \in \{0, 1, \dots, t\}} \|v(k)\|), \quad (8)$$

for all $x(0) \in \mathbb{R}^n$ and $t \in \mathbb{N}_{\geq 0}$.

In essence, a system is ISS if it is nominally asymptotically stable and the influence of the disturbance is bounded. This makes ISS a convenient framework to analyze the stability of systems subject to random disturbances. Similarly to nonlinear stability analysis for deterministic systems, we can show a system is ISS if there exists an ISS-Lyapunov function.

Definition 4 (ISS-Lyapunov function [39]): The function $V : \mathbb{N}_{\geq 0} \times \mathbb{R}^n \rightarrow \mathbb{R}$ is an *ISS-Lyapunov function* for the system $x(t+1) = q(t, x(t), v(t))$ if it is continuous in x , continuous at the origin for all $t \in \mathbb{N}_{\geq 0}$, and there exist three class- \mathcal{K}_∞ functions $\alpha_1, \alpha_2, \alpha_3$ and a class- \mathcal{K} function σ such that

$$\begin{aligned} \alpha_1(\|x(t)\|) &\leq V(t, x(t)) \leq \alpha_2(\|x(t)\|) \\ V(t+1, x(t+1)) - V(t, x(t)) &\leq -\alpha_3(\|x(t)\|) + \sigma(\|v(t)\|), \end{aligned} \quad (9)$$

for all $x(t) \in \mathbb{R}^n$.

Theorem 1 ([39]): A time-varying system is globally ISS if it admits an ISS-Lyapunov function.

The above definitions naturally extend to local ISS stability; for a detailed discussion, we refer readers to [37], [39], [40].

IV. ADAPTIVE ROBUST MPC

In this section, we first describe assumptions on and necessary features of the learning procedure in a way that is agnostic to the choice of learning algorithm. We then introduce our adaptive robust MPC approach, and prove stability of the combined learning and control framework.

A. Learning Desiderata

Since the nonlinear dynamics term f is unknown, our method takes a *certainty equivalent* approach by substituting an estimate \hat{f} that is refined online as more data becomes available. To guarantee the adaptive robust MPC framework satisfies state and input constraints for all time (i.e., safety), we make several assumptions on \hat{f} . We discuss two commonplace estimators that satisfy these assumptions later in §V.

We maintain the estimate

$$\hat{f}(x, t) = \widehat{W}(t)\phi(x) \quad (10)$$

of $f(x)$, where $\widehat{W}(t) \in \mathbb{R}^{n \times d}$ is our estimate of W at time t . To this end, we need bounds on our initial uncertainty, i.e., the difference between $f(x)$ and $\hat{f}(x, 0)$ for all x . For a general statistical estimator, this entails specifying a risk tolerance $\delta \in (0, 1)$ and computing confidence intervals on the estimate.

Assumption 2 (prior knowledge): Let $w_i, \hat{w}_i(t)$, and $\tilde{w}_i(t)$ be the i -th rows of $W, \widehat{W}(t)$, and $\widetilde{W}(t) := \widehat{W}(t) - W$, respectively, for $i \in \{1, 2, \dots, n\}$. At $t = 0$, we know an initial estimate $\widehat{W}(0) \in \mathbb{R}^{n \times d}$ and bounded sets $\{\mathcal{W}_i(0)\}_{i=1}^n$

with $\mathcal{W}_i(0) \subset \mathbb{R}^d$, such that $\widehat{W}(0) \in \mathcal{W}(0)$ with probability at least $1 - \delta$, where we define the sets

$$\mathcal{W}(t) := \{\widehat{W} \in \mathbb{R}^{n \times d} \mid \tilde{w}_i \in \mathcal{W}_i(t), \forall i \in \{1, \dots, n\}\}, \quad (11)$$

for all $t \in \mathbb{N}_{\geq 0}$.

Assumption 2 provides only an initial bound on the error of the estimate, which we explicitly label as the estimate at $t = 0$. Later, we will define $\widehat{\mathcal{W}}(t)$ for all $t \in \mathbb{N}_{\geq 0}$ when we adaptively update our estimate $\widehat{W}(t)$ and the bounds $\{\mathcal{W}_i(t)\}_{i=1}^n$ online. Our approach leverages the certainty equivalent ‘‘estimate and cancel’’ control laws pioneered in classical unconstrained adaptive control [24], [33]. As such, Assumption 1 and Assumption 2 are necessary to bound our approximation error. However, since we specify the risk tolerance δ , we cannot guarantee exact constraint satisfaction for all time. Instead, we slightly relax our definition of safety.

Definition 5: Under a risk tolerance of $\delta \in [0, 1]$, the system (1) is *safe* when

$$\text{Prob}(x(t) \in \mathcal{X}, u(t) \in \mathcal{U}, \forall t \geq 0) \geq 1 - \delta. \quad (12)$$

That is, the probability of a constraint violation should be no more than δ over the *entire* realized trajectory.

To guarantee closed-loop safety, we assume we have an online adaptation strategy that ensures the quality of the estimate $\widehat{W}(t)$ cannot get worse over time.

Assumption 3 (online learning): We have an *online parameter estimator* that maps an initial estimate $\widehat{W}(0)$, the associated $1 - \delta$ confidence interval $\mathcal{W}(0)$, and the trajectory history $\{x(k), u(k)\}_{k=0}^t$ to an *online estimate* $\widehat{W}(t)$ and confidence interval $\mathcal{W}(t)$ at time t , such that $\hat{w}_i(t) - w_i \in \mathcal{W}_i(t)$ for all time $t \in \mathbb{N}_{\geq 0}$ and $i \in \{1, 2, \dots, n\}$ with probability at least $1 - \delta$. We assume the confidence intervals on $\widehat{W}(t)$ are not growing with time, i.e., that

$$\mathcal{W}(t+1) \subseteq \mathcal{W}(t), \quad (13)$$

for all $t \in \mathbb{N}_{\geq 0}$.

Intuitively, Assumption 3 states that more data should not decrease the confidence in our estimate of W . We discuss two commonplace estimators that satisfy Assumption 3 in §V. Formulating separate confidence intervals for each row of $\widehat{W}(t)$ is a natural approach, as fitting $\widehat{W}(t)$ to historical data decomposes into n separate least-squares problems (one for each row) if the cross-covariance of $v(t)$ is zero.

Crucially, Assumption 3 allows us to treat the confidence intervals $\mathcal{W}(t)$ as exact bounds in the control design, since a robust controller that guarantees constraint satisfaction conditioned on the event that $\hat{w}_i(t) - w_i \in \mathcal{W}_i(t)$ for $i \in \{1, \dots, n\}$ and all time then satisfies (12). More formally, let SAFE denote the event that $x(t) \in \mathcal{X}$ and $u(t) \in \mathcal{U}$ for all $t \geq 0$. Then, apply the law of total probability to see that

$$\begin{aligned} &\text{Prob}(\text{SAFE}) \\ &\geq \text{Prob}(\text{SAFE} \mid \widehat{W}(t) - W \in \mathcal{W}(t), \forall t \geq 0) \\ &\quad \times \text{Prob}(\widehat{W}(t) - W \in \mathcal{W}(t), \forall t \geq 0) \quad , \quad (14) \\ &\geq \underbrace{\text{Prob}(\text{SAFE} \mid \widehat{W}(t) - W \in \mathcal{W}(t), \forall t \geq 0)}_{=1 \text{ using robust MPC in §IV-B}} (1 - \delta) \end{aligned}$$

since application of the estimators in §V that satisfy Assumption 3 ensures $\text{Prob}(\widehat{W}(t) - W \in \mathcal{W}(t), \forall t \geq 0) \geq 1 - \delta$.

Therefore, we treat the chance constraint (12) as a proxy for robust constraint satisfaction and construct our approach for the remainder of §IV conditioned on the event that $\hat{w}_i(t) - w_i \in \mathcal{W}_i(t)$ for all time and $i \in \{1, \dots, N\}$. This approach was also taken in [13], [16].

Assumption 3 does not require the estimated range of the unknown function f to shrink over time. Therefore, our estimation procedure differs from methods such as [6], [7], [14] that refine a non-increasing bound exclusively on the range of f without taking direct advantage of the structure in the nonlinear dynamics.

B. Certainty Equivalent Cancellation

We propose optimizing over feedback policies that cancel as much of the nonlinear term $f(x)$ as possible.

Definition 6: The set of *matching certainty equivalent (CE) policies* is the time-varying function class whose elements $\pi : \mathcal{X} \times \mathbb{N}_{\geq 0} \rightarrow \mathcal{U}$ are of the form

$$\pi(x(t), t) = u^*(x(t), t) - B^\dagger \hat{f}(x(t), t) \quad (15)$$

The matching CE policies simply project $\hat{f}(x(t), t)$ onto $\text{Range}(B)$, and cancel out as much of the disturbance as possible in the Euclidean norm sense, since

$$B^\dagger \hat{f}(x(t), t) = \arg \min_z \|Bz - \hat{f}(x(t), t)\|. \quad (16)$$

The matching CE law (15) results in the closed-loop dynamics

$$\begin{aligned} x(t+1) &= Ax(t) + B\pi(x(t), t) + f(x(t)) + v(t) \\ &= Ax(t) + Bu^*(x(t), t) + d(t) \end{aligned}, \quad (17)$$

where we define the compound disturbance term $d(t)$ as

$$\begin{aligned} d(t) &:= v(t) + f(x(t)) - BB^\dagger \hat{f}(x(t), t) \\ &= v(t) + BB^\dagger (f(x(t)) - \hat{f}(x(t), t)) \\ &\quad + (I - BB^\dagger) f(x(t)) \end{aligned} \quad (18)$$

We have written $d(t)$ above with three terms to highlight that it is driven by the process disturbance $v(t)$, the estimation error $f(x) - \hat{f}(x)$, and the imperfect matching using B^\dagger .

Remark 1: If we know that f is a *matched* uncertainty (i.e., that $f(x(t)) = Bg(x(t))$ for some function $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$), we can reduce the compound disturbance to

$$d(t) := v(t) + B(g(x(t)) - \hat{g}(x(t), t)), \quad (19)$$

since $B^\dagger B = I$. Therefore, the matching certainty equivalent controller (15) generalizes approaches for systems with matched uncertainty to those with unmatched uncertainty.

As mentioned in our learning desiderata, to guarantee safety as defined in (12), we construct a controller that guarantees constraint satisfaction when $\widehat{W}(t) - W \in \mathcal{W}(t)$ for all $t \geq 0$, which we will assume is the case for the remainder of this section. To design a robust matching CE policy, we need to:

- Guarantee that the closed-loop dynamics (17) do not violate state constraints.
- Tighten input constraints to account for the certainty equivalent cancellation in (15).

Therefore, we introduce two simple polytopic approximations that bound the support of the cancellation term in (15) and the terms that make up the compound disturbance $d(t)$.

Lemma 1: Consider online approximation of $f(x)$ with features satisfying Assumption 1, an estimator satisfying Assumption 3, and define the estimated support set as

$$\mathcal{F}(t) := \{z \in \mathbb{R}^n : |z_i| \leq \|\hat{w}_i(t)\| + 2 \max_{\tilde{w}_i \in \mathcal{W}_i(t)} \|\tilde{w}_i\|\}. \quad (20)$$

Then, for all $x \in \mathcal{X}$ and $t, k \in \mathbb{N}_{\geq 0}$ it holds that

$$\hat{f}(x, t+k), f(x) \in \mathcal{F}(t). \quad (21)$$

Proof: We show $\mathcal{F}(t)$ over-approximates the range of values $z = \hat{f}(x, t+k)$ can take for any $x \in \mathcal{X}$ and $k \geq 0$. Let $z = \hat{f}(x, t+k) = \widehat{W}(t+k)\phi(x)$. Then,

$$\begin{aligned} |z_i| &= |\hat{w}_i(t+k)^\top \phi(x)| \\ &\leq \|\hat{w}_i(t+k)\| \\ &\leq \|\hat{w}_i(t)\| + \|\hat{w}_i(t+k) - \hat{w}_i(t)\| \\ &\leq \|\hat{w}_i(t)\| + \|\hat{w}_i(t+k) - w_i\| + \|\hat{w}_i(t) - w_i\|. \end{aligned} \quad (22)$$

The non-increasing confidence interval property from Assumption 3 gives $\mathcal{W}_i(t+k) \subseteq \mathcal{W}_i(t)$ for $k \geq 0$, so

$$|z_i| \leq \|\hat{w}_i(t)\| + 2 \max_{\tilde{w}_i \in \mathcal{W}_i(t)} \|\tilde{w}_i\|, \quad (23)$$

which proves that $\hat{f}(x, t+k) \in \mathcal{F}(t)$ for all $k \geq 0$. In addition, let $y = f(x) = W\phi(x)$ for some $x \in \mathcal{X}$. Then,

$$\begin{aligned} |y_i| &= |w_i^\top \phi(x)| \\ &\leq \|w_i\| \\ &\leq \|\hat{w}_i(t)\| + \|\hat{w}_i(t) - w_i\| \\ &\leq \|\hat{w}_i(t)\| + \max_{\tilde{w}_i \in \mathcal{W}_i(t)} \|\tilde{w}_i\|. \end{aligned} \quad (24)$$

Hence, $f(x) \in \mathcal{F}(t)$. ■

The set $\mathcal{F}(t)$ in Lemma 1 contains all possible values that our online estimate can take for all future times. It is not straightforward to create a tighter approximation (i.e., eliminate the factor of 2) without additional assumptions. To see this, consider a constant unit norm ball confidence interval. In the worst case, the true parameter lies on the boundary of the ball around the current estimate. This means all future estimates may lie a Euclidean distance of 2 units away from the current estimate, yielding the bound in Lemma 1.

Remark 2: Representing the set $\mathcal{F}(t)$ defined in Lemma 1 requires finding the max-norm element of a convex set. For many convex confidence intervals this is straightforward to compute, although it is not always efficient. For ellipsoids, this amounts to an eigenvalue computation. For polytopes, this requires vertex enumeration, incurring exponential complexity in the dimension of the state space n .

Moreover, Lemma 1 does not require that $\mathcal{F}(t+1) \subseteq \mathcal{F}(t)$, so we provide the following corollary to help us create an approximation that is non-increasing in size.

Corollary 1: At time t , the sets $\{\mathcal{F}(i)\}_{i=0}^t$ are known, so

$$\hat{f}(x, t+k), f(x) \in \bigcap_{i=0}^t \mathcal{F}(i) =: \widehat{\mathcal{F}}(t) \quad (25)$$

for all $x \in \mathcal{X}$ and $k \in \mathbb{N}_{\geq 0}$, where we define $\widehat{\mathcal{F}}(t)$ as the set

$$\{z : |z_i| \leq \min_{j \in \{0, \dots, t\}} [\|\hat{w}_i(j)\| + 2 \max_{\tilde{w}_i \in \mathcal{W}_i(j)} \|\tilde{w}_i\|]\}. \quad (26)$$

Note $\widehat{\mathcal{F}}(t)$ can be computed recursively in time.

To construct a robust MPC problem to optimize the CE policy (15), we need to account for the compound disturbance $d(t)$. We do this with the following lemma.

Lemma 2: Assume the online parameter estimator satisfies Assumption 3 with features that satisfy Assumption 1 and define the approximation error support as the set

$$\mathcal{D}(t) := \{z \in \mathbb{R}^n \mid |z_i| \leq \max_{\tilde{w}_i \in \mathcal{W}_i(t)} \|\tilde{w}_i\|, \forall i \in \{1, \dots, n\}\}. \quad (27)$$

If we control the system (1) using the certainty equivalent control law (15), then at time t for all $k \in \mathbb{N}_{\geq 0}$, the compound disturbance $d(t+k)$ in the dynamics (17) is contained in the set $\widehat{\mathcal{D}}(t) \subseteq \widehat{\mathcal{D}}(t-1)$, defined as

$$\widehat{\mathcal{D}}(t) := (I - BB^\dagger)\widehat{\mathcal{F}}(t) \oplus BB^\dagger\mathcal{D}(t) \oplus \mathcal{V}. \quad (28)$$

Here \oplus indicates the Minkowski sum and a matrix-set multiplication indicates a linear transformation of the set's elements.

Proof: At any state $x \in \mathcal{X}$ and time $t+k$, let $z = \hat{f}(x, t+k) - f(x) = (\widehat{W}(t+k) - W)\phi(x) = \widetilde{W}\phi(x)$ for some $\widetilde{W} \in \mathcal{W}(t+k) \subseteq \mathcal{W}(t)$. Then

$$|z_i| = |\widetilde{w}_i^\top \phi(x_t)| \leq \|\widetilde{w}_i\| \leq \max_{\tilde{w}_i \in \mathcal{W}_i(t)} \|\tilde{w}_i\|.$$

So $\hat{f}(x, t+k) - f(x) \in \mathcal{D}(t)$ for all $k \in \mathbb{N}_{\geq 0}$. Since $\mathcal{W}_i(t) \subseteq \mathcal{W}_i(t-1)$ by Assumption 3, this implies $\mathcal{D}(t) \subseteq \mathcal{D}(t-1)$. Then, by Corollary 1, $f(x) \in \widehat{\mathcal{F}}(t)$. Therefore, $d(t)$ in the closed-loop dynamics (17) is contained in $\widehat{\mathcal{D}}(t)$ since $\mathcal{D}(t)$ is symmetric. In addition, since both $\widehat{\mathcal{F}}(t) \subseteq \widehat{\mathcal{F}}(t-1)$ and $\mathcal{D}(t) \subseteq \mathcal{D}(t-1)$, the support of $d(t)$ is nested over time, i.e., $\widehat{\mathcal{D}}(t) \subseteq \widehat{\mathcal{D}}(t-1)$. ■

Remark 3: If the estimation error is small, (28) shows that the compound disturbances $d(t)$ will only depend on the process noise and the component of the unknown function f that cannot be cancelled using the CE control law (15). Therefore, if the estimation error is small and a significant component of the uncertainty is matched, we generally expect our method to tolerate nonlinear uncertainty with larger magnitude than a conventional robust MPC scheme.

Remark 4: We could consider multiple variations on the bounds in Lemma 1 and Lemma 2 that would yield equivalent properties of the closed-loop system. For example, if a bound on the true range of f is known a priori, we may project \hat{f} into a known box enclosing the support of f .

Remark 5: Lemma 1 and Lemma 2 result in box constraints on the disturbances. Approximations of this form are often more favorable from a practical perspective compared to (operator) norm type bounds, since these box constraints are better able to preserve the relative magnitudes of state variables that represent physical quantities.

We now use the set $\mathcal{F}(t)$ from Lemma 1 and the set $\widehat{\mathcal{D}}(t)$

from Lemma 2 to modify the robust MPC problem (6) into

$$\begin{aligned} & \text{minimize} && V_N(\bar{x}_{t+N|t}) + \sum_{k=0}^{N-1} h(\bar{x}_{t+k|t}, \bar{u}_{t+k|t}) \\ & \text{subject to} && \bar{x}_{t+k+1|t} = A\bar{x}_{t+k|t} + B\bar{u}_{t+k|t} \\ & && x_{t+k+1|t} = Ax_{t+k|t} + Bu_{t+k|t} + d_{t+k|t} \\ & && u_{t+k|t} = \bar{u}_{t+k|t} + \sum_{j=0}^{k-1} K_{k,j|t} d_{t+j|t} \\ & && x_{k|t} \in \mathcal{X}, u_{t+k|t} \in \mathcal{U} \ominus B^\dagger \widehat{\mathcal{F}}(t) \\ & && \forall k \in \{0, 1, \dots, N-1\} \\ & && \bar{x}_{t|t} = x(t), x_{t|t} = x(t), x_{t+N|t} \in \mathcal{O}(t) \\ & && \forall \{d_{t+k|t}\}_{k=0}^{N-1} \subset \widehat{\mathcal{D}}(t) \end{aligned} \quad (29)$$

Compared to (6), in (29) we have tightened the input constraints to account for the matching term in the certainty equivalent policy (15). As in standard robust MPC, we assume we can compute a robust control invariant set $\mathcal{O}(t)$ and that we have access to a convex terminal cost function V_N .

Let $\{u_{t|t}^*, \dots, u_{t+N-1|t}^*\}$ be an optimal policy sequence for (29). Online, we solve (29) at each time step $t \in \mathbb{N}_{\geq 0}$ and choose the robust control term of the *certainty equivalent* control policy (15) as the receding horizon feedback law $u^* : \mathcal{X} \times \mathbb{N}_{\geq 0} \rightarrow \mathcal{U}$ such that

$$u^*(x(t), t) = u_{t|t}^*. \quad (30)$$

Assumption 4: The terminal cost $V_N : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is a continuous convex Lyapunov function for the nominal dynamics under a policy $u_N(x) = -Kx$. That is, there exists a class- \mathcal{K}_∞ function $\alpha_N(\|x\|) \geq h(x, u_N(x))$ for which

$$V_N((A - BK)x) - V_N(x) \leq -\alpha_N(\|x\|), \quad (31)$$

for all $x \in \mathbb{R}^n$.

Assumption 5: For the policy $u_N(x) = -Kx$ in Assumption 4, the terminal set $\mathcal{O}(t) \subseteq \mathcal{X}$ is a maximal robust positive invariant set for the closed-loop system $x(k+1) = (A - BK)x(k) + d(k)$ for $d(k) \in \widehat{\mathcal{D}}(t)$ subject to $x(k) \in \mathcal{X}$ and $u_N(x(k)) \in \mathcal{U} \ominus B^\dagger \widehat{\mathcal{F}}(t)$ for all $t \geq 0$.

Assumption 4 and Assumption 5 are standard and easily satisfied by taking $u_N(x) = -Kx$ and $V_N(x) = x^\top Px$ as the solution to an LQR problem with h as the stage cost. Then, $\mathcal{O}(t)$ can be computed efficiently using the standard algorithms in [29].

Lemmas 1 and 2 imply that $\mathcal{O}(t-1) \subseteq \mathcal{O}(t)$ since $\widehat{\mathcal{D}}(t) \subseteq \widehat{\mathcal{D}}(t-1)$ and $\widehat{\mathcal{F}}(t) \subseteq \widehat{\mathcal{F}}(t-1)$. Therefore, the terminal constraint becomes less conservative over time.

C. Stability

We prove the stability of our algorithm through a recursive feasibility and input-to-state stability argument.

Theorem 2: Consider the system (1), a parameter estimator that satisfies Assumption 3 with features that satisfy Assumption 1 in closed-loop feedback with the *matching certainty equivalent* control law (15),(30). If the tube MPC problem (29) is feasible at $t = 0$, then for all $t \geq 0$ we have that (29) is feasible and the closed-loop system (1),(15),(30) satisfies $x(t) \in \mathcal{X}$, and $\pi(x(t), t) \in \mathcal{U}$.

Proof: Suppose the optimal control problem (29) is feasible at time t , with solution $\{u_{t|t}^*(\cdot), \dots, u_{t+N-1|t}^*(\cdot)\}$. By

Lemma 1 and Corollary 1 we have that $\hat{f}(x, t+k) \in \widehat{\mathcal{F}}(t)$ for all $k \in \mathbb{N}_{\geq 0}$. Therefore, the time-varying CE control law

$$\pi_{t+k|t}(\cdot) = u_{t+k|t}^*(\cdot) - B^\dagger \hat{f}(x(t), t+k) \quad (32)$$

satisfies the input constraints for $t \in [t, t+N-1]$, since (29) then implies $u_{t+k|t}^*(\cdot) \in \mathcal{U} \ominus B^\dagger \widehat{\mathcal{F}}(t)$. Moreover, by Lemma 2 the disturbance support does not grow in time, i.e., $\widehat{\mathcal{D}}(t+1) \subseteq \widehat{\mathcal{D}}(t)$. Therefore, we have that under policy (32) the closed-loop trajectory formed by (1),(32) satisfies $x(t+k) \in \mathcal{X}$ for all $k \in [0, N]$ and that $x(t+N) \in \mathcal{O}(t)$. Hence, if we apply the CE policy (15),(30) at time t , then $x(t+1) \in \mathcal{X}$ and $\pi(x(t), t) \in \mathcal{U}$.

By Assumption 5, for any $x \in \mathcal{O}(t) \subseteq \mathcal{O}(t+1)$, applying the policy $u_N(x) \in \mathcal{U} \ominus B^\dagger \widehat{\mathcal{F}}(t)$ implies that $Ax + Bu_N(x) + d \in \mathcal{O}(t+1)$ for any $d \in \widehat{\mathcal{D}}(t)$. Therefore, Corollary 1 and Lemma 2 imply the policy sequence $\{u_{t+1|t}^*(\cdot), \dots, u_{t+N-1|t}^*(\cdot), u_N(\cdot)\}$ is feasible for the tube MPC problem (29) at time $t+1$. Therefore, if the MPC program (29) is feasible at time $t=0$, it is also feasible for all $t \geq 0$ and the closed-loop system formed by the matching CE law (1), (15), (30) must robustly satisfy state and input constraints by induction. ■

Theorem 3: Consider a system of the form in (1), a parameter estimator that satisfies Assumption 3 with features that satisfy Assumption 1 in closed-loop feedback with the *certainty equivalent* control law (15),(30). Let $\mathcal{X}_N \subseteq \mathcal{X}$ denote the set of states for which the tube MPC problem (29) is feasible. Then the closed-loop system is locally input-to-state stable with region of attraction \mathcal{X}_N .

Proof: Our proof closely resembles [8, Thm. 2]. We argue that the nominal system is stable by a standard MPC argument, and that the closed-loop system is ISS since the disturbances are bounded. Let $J_N^*(t, x(t))$ be the optimal value of (29) associated with the nominal prediction $\{\bar{x}_{t|t}^*, \dots, \bar{x}_{t+N|t}^*\}$ and feedback policies $\{u_{t|t}^*(\cdot), \dots, u_{t+N-1|t}^*(\cdot)\}$. Since we assume the stage cost is quadratic, there exist two class- \mathcal{K}_∞ functions α_1, α_2 such that for all $t \geq 0$, $\alpha_1(\|x\|) \leq J_N^*(t, x) \leq \alpha_2(\|x\|)$, a class- \mathcal{K}_∞ function α_3 such that $h(x, u) \geq \alpha_3(\|x\|)$, and $J_N^*(t, 0) = 0$ (see [37, Prop. 1], [8, Thm. 2]). As in the proof of Theorem 2, we have that if we apply the CE control law (15),(30) at time t , then the policies $\{u_{t+1|t}^*, \dots, u_{t+N-1|t}^*, u_N\}$ are a feasible solution for (29) at time $t+1$. Let $\bar{J}(t, x)$ be the cost associated with forward simulating the nominal system using the policies $\{u_{t+1|t}^*, \dots, u_{t+N-1|t}^*, u_N\}$ with x as initial condition. i.e., set $u_{t+N|t}^* = u_N(\cdot)$ and let $\bar{x}_{t+1|t+1} = x$, $\bar{x}_{k+1|t+1} = A\bar{x}_{k|t+1} + Bu_{k|t}^*(\bar{x}_{k|t+1})$ for $k \in \{t+1, \dots, t+N\}$ so that

$$\bar{J}(t, x) = \sum_{k=t+1}^{t+N} h(\bar{x}_{k|t+1}, u_{k|t}^*(\bar{x}_{k|t+1})) + V_N(\bar{x}_{t+N+1|t+1}).$$

This gives that $J_N^*(t+1, x(t+1)) \leq \bar{J}(t, x(t+1))$. Moreover, since the stage cost is quadratic and by Assumption 4, $\bar{J}(t, x)$ is uniformly continuous in x for all $t \geq 0$ on the state space since the inputs are constrained in a compact set. It follows that for $x_1, x_2 \in \mathcal{X}$, there exists a \mathcal{K}_∞ function α_J such that for all $t \geq 0$, $|\bar{J}(t, x_1) - \bar{J}(t, x_2)| \leq \alpha_J(\|x_1 - x_2\|)$ (see [37,

Lem. 1]). Therefore,

$$\begin{aligned} & J_N^*(t+1, x(t+1)) - J_N^*(t, x(t)) \\ & \leq \bar{J}(t, x(t+1)) - J_N^*(t, x(t)) \\ & = \bar{J}(t, x(t+1)) - \bar{J}(t, \bar{x}_{t+1|t}^*) + \bar{J}(t, \bar{x}_{t+1|t}^*) - J_N^*(t, x(t)) \\ & \leq |\bar{J}(t, x(t+1)) - \bar{J}(t, \bar{x}_{t+1|t}^*)| - h(x_t, \bar{u}_{t|t}^*(x(t))) \\ & \leq \alpha_J(\|d(t)\|) - \alpha_3(\|x(t)\|). \end{aligned}$$

So the system is ISS by Theorem 1. ■

Remark 6: The ISS result in Theorem 3 does not explicitly show that improvements in the confidence of the model lead to better performance of the controller, since we only assume the model confidence is non-decreasing in Assumption 3.

In adaptive control, stronger guarantees of performance improvement are typically made under *persistence of excitation* assumptions [24].

Remark 7: The only property of the feature map ϕ that we relied on to prove Lemmas 1, 2 was the boundedness assumption $\|\phi(x)\| \leq 1$. Therefore, Lemmas 1, 2 and Corollary 1 also hold when the feature map depends on an additional *exogenous* signal $z(t) \in \mathcal{Z} \subseteq \mathbb{R}^p$ beside the system state $x(t)$, as long as $\|\phi(x, z)\| \leq 1$ for all $x \in \mathcal{X}$ and $z \in \mathcal{Z}$. This implies that Theorems 2 and 3 also hold when the bounded feature map ϕ depends on an exogenous signal $z(t)$. Therefore, we can easily incorporate observable time-varying environmental effects that we know influence the dynamics, even when it is unclear how these exogenous signals evolve over time, reminiscent of approaches in parameter-varying control [41]. We take this approach on the example in §VI-E.

D. Iterative/Episodic Learning

Oftentimes, a controller is used to perform the same task repeatedly over many iterations (or episodes). We can apply the algorithm we developed in §IV-B to such an iterative learning (IL), or episodic, setting as long as the tube MPC problem (29) is feasible at the start of each iteration. In this subsection, we take advantage of the properties of the approach presented in §IV-B to reduce the amount of online computation required and update some of the MPC problem parameters only offline between episodes, generalizing our method to a *family* of algorithms suited to episodic operation.

Definition 7: In an *iterative* (or *episodic*) setting, the system starts from the fixed initial condition $x^{(j)}(0) = x(0) \in \mathcal{X}$ at $t=0$ and evolves for a finite duration $T^{(j)} \in \mathbb{N}_{>0}$ for each iteration $j \in \mathbb{N}_{\geq 0}$. We use the notation $x^{(j)}(t)$ to denote the state at time t of iteration j , and apply this notation to time-varying quantities in general, e.g., we use the notation $u^{(j)}(t)$ for the input into the system.

Updating the sets $\widehat{\mathcal{F}}$, $\widehat{\mathcal{D}}$, and in particular \mathcal{O} online can require significant computation. Therefore, it is attractive to only update these sets offline between iterations in an episodic setting, depending on available computational resources. We summarize the complete online method we developed in §IV-B as Algorithm 1.A and introduce two episodic variants as Algorithms 1.B-C that reduce online computation by keeping \mathcal{O} fixed throughout an iteration. In particular,

- **Alg. 1.B** updates $\widehat{\mathcal{F}}^{(j)}(t)$, $\widehat{\mathcal{D}}^{(j)}(t)$ at each timestep t online and keeps $\mathcal{O}^{(j)}(t)$ fixed as $\mathcal{O}^{(j)}(0)$ during the j -th episode.

Algorithm 1: Adaptive CE MPC with online or episodic updates

Require: initial estimate $\widehat{W}(0)$

Require: initial $1 - \delta$ confidence interval $\mathcal{W}(0)$

Require: terminal cost V_N and feedback gain K as in Assumption 4

- 1: Initialize $\widehat{\mathcal{F}}, \widehat{\mathcal{D}}, \mathcal{O}$ using (26), (28), and e.g., [29, Alg. 10.4].
 - 2: **for** $j = 0, 1, \dots$, episodes **do**
 - 3: Set initial condition $x^{(j)}(0) = x(0)$.
 - 4: **for** $t = 0, \dots, T^{(j)}$ **do**
 - 5: Compute u^* as (30) using (29).
 - 6: Apply the matching CE policy π (15).
 - 7: Observe the next state $x^{(j)}(t+1)$.
 - 8: Update estimate \hat{f} and confidence interval \mathcal{W} .
 - 9: **Alg. 1.A:** Update $\{\widehat{\mathcal{F}}, \widehat{\mathcal{D}}, \mathcal{O}\}$.
 - Alg. 1.B:** Update $\{\widehat{\mathcal{F}}, \widehat{\mathcal{D}}\}$.
 - Alg. 1.C:** Continue.
 - 10: **end for**
 - 11: **Alg. 1.A:** Continue.
 - Alg. 1.B:** Update \mathcal{O} .
 - Alg. 1.C:** Update $\{\widehat{\mathcal{F}}, \widehat{\mathcal{D}}, \mathcal{O}\}$ using $\{\hat{f}^{(j)}(\cdot, t), \mathcal{W}^{(j)}(t)\}_{t=0}^{T^{(j)}}$.
 - 12: **end for**
-

- **Alg. 1.C** keeps $\widehat{\mathcal{F}}^{(j)}(t)$, $\widehat{\mathcal{D}}^{(j)}(t)$, and $\mathcal{O}^{(j)}(t)$ fixed as $\widehat{\mathcal{F}}^{(j)}(0)$, $\widehat{\mathcal{D}}^{(j)}(0)$, and $\mathcal{O}^{(j)}(0)$ during the j -th episode.

It is straightforward to show that we retain the recursive feasibility and stability guarantees in Theorem 2 and Theorem 3 when we apply the episodic variants of Algorithm 1.A.

Lemma 3: Consider episodic control of the system (1) in closed-loop with Algorithm 1.B or Algorithm 1.C from the initial condition $x(0)$. If the tube MPC problem (29) is feasible at $t = 0$ at iteration $j = 0$, then for all $t \geq 0$ at iterations $j \geq 0$ it holds that (29) is feasible, the closed-loop system satisfies $x^{(j)}(t) \in \mathcal{X}$ and $\pi^{(j)}(x^{(j)}(t), t) \in \mathcal{U}$, and the optimal cost functions $J_N^{*,B,(j)}(t, x(t))$ and $J_N^{*,C,(j)}(t, x(t))$ for (29) are ISS-Lyapunov functions for the closed-loop systems formed by Algorithm 1.B and Algorithm 1.C respectively.

Proof: First, we note that the proofs of Theorem 2 and Theorem 3 only relied on the facts that $\widehat{\mathcal{D}}^{(j)}(t)$, $\widehat{\mathcal{F}}^{(j)}(t)$ do not grow over time during an iteration j , i.e., that $\widehat{\mathcal{F}}^{(j)}(t+1) \subseteq \widehat{\mathcal{F}}^{(j)}(t)$ and $\widehat{\mathcal{D}}^{(j)}(t+1) \subseteq \widehat{\mathcal{D}}^{(j)}(t)$, and that $\mathcal{O}^{(j)}(t)$ does not shrink, i.e., that $\mathcal{O}^{(j)}(t) \subseteq \mathcal{O}^{(j)}(t+1)$.

Therefore, if we keep $\widehat{\mathcal{F}}^{(j)}(t) := \widehat{\mathcal{F}}^{(j)}(0)$ and $\widehat{\mathcal{D}}^{(j)}(t) := \widehat{\mathcal{D}}^{(j)}(0)$ fixed over iteration j , consequently fixing $\mathcal{O}^{(j)}(t) := \mathcal{O}^{(j)}(0)$, Theorem 2 and Theorem 3 still hold, since Corollary 1 implies that $f(x^{(j)}(t))$, $\hat{f}^{(j)}(x^{(j)}(t), t) \in \widehat{\mathcal{F}}^{(j)}(0)$ and Lemma 2 implies that $d^{(j)}(t) \in \widehat{\mathcal{D}}^{(j)}(0)$ for all $t \geq 0$. Thus, Alg. 1.C is recursively feasible and ISS over an iteration j .

Moreover, note that $\mathcal{O}^{(j)}(0)$ is an RPI set for the closed-loop system in Assumption 5 associated with disturbances in $\widehat{\mathcal{D}}^{(j)}(0)$, state constraints \mathcal{X} , and input constraints $\mathcal{U} \oplus B^\dagger \widehat{\mathcal{F}}^{(j)}(0)$. Since $\widehat{\mathcal{F}}$ and $\widehat{\mathcal{D}}$ do not grow over time, it follows that $\mathcal{O}^{(j)}(0)$ is also RPI for the closed-loop system in Assumption 5 with disturbances in $\widehat{\mathcal{D}}^{(j)}(t)$ and input constraints $\mathcal{U} \oplus B^\dagger \widehat{\mathcal{F}}^{(j)}(t)$ for all $t \geq 0$. Hence, the proofs of Theorem 2 and Theorem 3 also apply to Algorithm 1.B.

Finally, note that updates to $\widehat{\mathcal{F}}$, $\widehat{\mathcal{D}}$, and \mathcal{O} can only increase the feasible domain of (29). Therefore, if (29) is feasible for Algorithms 1.B-C at $t = 0$ and $j = 0$, it will be also be feasible at $t = 0$ for all $j \geq 0$ since $x^{(j)}(0) = x(0)$ is constant. This proves the lemma. ■

Lemma 3 shows that we can safely reduce the amount of online computation required to apply our adaptive MPC in an episodic setting, even though we still adapt the model online in Algorithms 1.A-C. We can typically update the model estimate efficiently using recursive filters such as those we discuss in §V, though one could trivially keep the model fixed over an iteration as well.

Remark 8: In general, we should expect the computational benefits of applying Algorithms 1.B-C to come at the expense of conservatism, since updates to $\widehat{\mathcal{F}}$, $\widehat{\mathcal{D}}$, and \mathcal{O} can increase the size of the feasible set of problem (29). However, it is not straightforward to mathematically relate the realized closed-loop costs when we apply Algorithms 1.A-C to each other, since the model estimates depend on the trajectory histories induced by the applied controller.

We emphasize that under our learning desiderata, episodic application of our robust adaptive MPC in combination with an estimator that satisfies Assumption 3 ensures that we satisfy the safety guarantee (12) with probability $1 - \delta$ jointly for all episodes. For some estimators, such as the estimator we discuss in §V-B, one could consider re-initializing the initial estimate at each episode using all previously collected data to yield tighter confidence intervals, thereby reducing conservatism. However, by doing so, one can only guarantee that state and inputs are satisfied with probability at least $1 - \delta$ for each episode, as opposed to across all episodes.

V. ADAPTATION LAWS & LEARNING ALGORITHMS

In this section we describe two common online function approximation schemes—one statistical and one not—that satisfy the decaying confidence interval of Assumption 3 that we used to construct our ARMPC algorithm.

A. Set Membership Estimation

A common approach in the adaptive MPC literature is to estimate constant, or slowly changing, disturbances through *set membership* estimation [42], [43]. These estimators maintain a feasible parameter set that is refined as more data becomes available. The feasible parameter set contains all credible model parameters that explain previous observations, which means that the feasible parameter sets are nested over time. We consider learning the parameters of a nonlinear uncertainty model of the form in (3) directly using set-membership estimation. Under the prior knowledge Assumption 2, the initial feasible parameter set is given as $\Theta(0) = \{\widehat{W}(0)\} \oplus \mathcal{W}(0)$ and the feasible parameter set at time t is obtained as

$$\Theta(t) = \{W \in \Theta : x(k+1) - Ax(k) - Bu(k) - W\phi(x(k)) \in \mathcal{V}, \forall k \in \{0, \dots, t-1\}\}. \quad (33)$$

When $\Theta(0)$ and \mathcal{V} are a hyperboxes, this estimator maintains independent feasible sets for each row of W and can be updated recursively in time with polytopical set intersections by rewriting (33) in terms of the row-wise vectorization of

W . Clearly, $\Theta(t) \subseteq \Theta(t-1)$. As is common practice in the literature [7], we propose generating a point estimate of the parameters as the *Chebyshev center* of the feasible parameter set:

$$\widehat{W}(t) = \arg \min_{\widehat{W}} \max_{W \in \Theta(t)} \|\widehat{W} - W\|_F. \quad (34)$$

By definition, this approach minimizes the worst-case error of the point estimates and is typically straightforward to compute [44]. Denoting the Chebyshev radius for the feasible parameter set associated with the i -th row of W as $r_i(t) = \min_w \max_{w_i \in \Theta_i(t)} \|w - w_i\|_2$, we take the confidence interval on $\widehat{w}_i(t) - w_i$ as $\mathcal{W}_i(t) = \{\tilde{w}_i : \|\tilde{w}_i\|_2 \leq r_i(t)\}$.

By definition, since $\Theta(t) \subseteq \Theta(t-1)$, the Chebyshev radii must be decreasing over time: $r_i(t) \leq r_i(t-1)$. Therefore, a set-membership estimator with point estimates as the Chebyshev center satisfies Assumption 3 with a risk tolerance $\delta = 0$, thus guaranteeing the safety of the system with probability 1 according to Definition 5.

B. Recursive Bayesian Linear Regression (BLR)

In the case of Bayesian estimation, we can generate confidence intervals directly from the posterior distribution over parameters if we know the disturbance distribution. We outline this approach under a simple, standard assumption.

Assumption 6: We assume that each entry of the process noise is bounded $v(t) = [v_1(t), \dots, v_n(t)]^\top \in \mathcal{V} = \{v : |v| \leq \sigma_i\}$, and that each entry $v_i(t)$ is independent of the others. Hence, $v_i(t)$ is sub-Gaussian with variance proxy σ_i^2 .

Under Assumption 6, we can essentially treat the noise as both normally distributed for convenient analysis and provide safety guarantees for the algorithm proposed in §IV. If we place subjective priors over the rows of W of the form $w_i \sim \mathcal{N}(\widehat{w}_i(0), \sigma_i^2 \Lambda_i^{-1}(0))$, then the resulting posterior parameter distribution at time t is also Gaussian, $w_i \sim \mathcal{N}(\widehat{w}_i(t), \sigma_i^2 \Lambda_i^{-1}(t))$. We then use the mean of the posterior—also corresponding to the maximum a posteriori (MAP) estimate—as a point estimate for control: $\hat{f}(x, t) = \widehat{W}(t)\phi(x)$. Defining the measurement and prediction at time t as $y(t) := x(t+1) - Ax(t) - Bu(t)$ and $\hat{y}(t) := \widehat{W}(t)\phi(t)$, the MAP estimate for each row can then be updated with constant complexity in time using the recursive updates

$$\begin{aligned} \widehat{w}_i(t+1) &= \widehat{w}_i(t) - \frac{(\hat{y}_i(t) - y_i(t))\phi(t)^\top \Lambda_i^{-1}(t)}{1 + \phi(t)^\top \Lambda_i^{-1}(t)\phi(t)} \\ \Lambda_i^{-1}(t+1) &= \Lambda_i^{-1}(t) - \frac{\Lambda_i^{-1}(t)\phi(t)\phi(t)^\top \Lambda_i^{-1}(t)}{1 + \phi(t)^\top \Lambda_i^{-1}(t)\phi(t)} \end{aligned}, \quad (35)$$

where we write each entry of $\hat{y}(t)$ as $\hat{y}_i(t) = \widehat{w}_i(t)^\top \phi(t)$ and $\phi(t) := \phi(x(t))$. We can recover the frequentist ordinary least-squares estimator if we assume a flat prior [45], which requires the availability of some amount of prior data to yield the initial values $\widehat{W}(0)$, $\Lambda(0)$.

Taking a risk tolerance of $\delta \in (0, 1)$, we could naively define the confidence interval for the i -th row of $\widehat{W}(t)$ as

$$\mathcal{W}_i^{\text{naive}}(t) := \{\tilde{w}_i \in \mathbb{R}^n : \tilde{w}_i^\top \Lambda_i(t)\tilde{w}_i \leq \sigma_i^2 \chi_n^2(1 - \frac{\delta}{n})\}, \quad (36)$$

where $\chi_n^2(1 - \delta)$ is the $1 - \delta$ quantile of the chi-square distribution with n degrees of freedom. However, the confidence interval in (36) does not capture the fact that we want to certify the safety of the policy for all time with high probability. We cannot achieve this with a single confidence interval of a point estimate at time t , as (36) ignores the correlations between the model estimates over time. As we discussed in our learning desiderata for robust control, we instead desire confidence intervals such that

$$\widehat{W}(t) - W \in \mathcal{W}(t), \quad \forall t \in \mathbb{N}_{\geq 0}, \quad (37)$$

with probability at least $1 - \delta$.

Recent work applied a Martingale argument originating from the Bandits literature to generate such confidence intervals by scaling the naive confidence intervals with a time-varying parameter [13]. The resulting safety guarantees are subject to assumptions on the calibration of the prior, for which we refer the reader to [13, Assumption 3]. This assumption—that the prior contains the true parameter with high probability—is trivially satisfied for flat priors (i.e., when we have access to some data a priori). We also assume this assumption is satisfied for the non-flat priors that we use in this work. We may assume that these priors are constructed either via empirical Bayes (as we discuss in §V-D) or via expert knowledge.

Theorem 4: [13, Thm 1] For the recursive Bayesian linear filter (35), we have the estimation error $\widehat{w}_i(t) - w_i \in \mathcal{W}_i(t)$ for all $t \geq 0$ with probability at least $1 - \delta$, where

$$\mathcal{W}_i(t) := \{\tilde{w}_i : (\tilde{w}_i^\top \Lambda_i(t)\tilde{w}_i)^{\frac{1}{2}} \leq \sigma_i \beta_t(\delta/n)\}, \quad (38)$$

with $\beta_t(\delta)$ equal to

$$\sqrt{2 \log \left(\frac{\det(\Lambda_i(t))^{1/2}}{\delta \det(\Lambda_i(0))^{1/2}} \right)} + \sqrt{\frac{\lambda_{\max}(\Lambda_i(0))}{\lambda_{\min}(\Lambda_i(t))} \chi_n^2(1 - \delta)}. \quad (39)$$

The confidence intervals resulting from Theorem 4 unfortunately do not immediately satisfy the requirement that $\mathcal{W}(t+1) \subseteq \mathcal{W}(t)$ in Assumption 3 without a *persistence of excitation* or *active exploration* assumption as is made in [21]. A simple workaround is to update the estimate (35) fed into to the controller only when the associated confidence intervals (38) have shrunk, effectively disregarding new data until the system has been excited sufficiently. This approach was shown to perform well in practice [20]. Therefore, we can apply the BLR estimator to guarantee the safety of the system with any desired risk tolerance $\delta \in [0, 1]$ according to Definition 5. Still, future work should explore strategies to guarantee confidence intervals constructed using Theorem 4 (or other equivalent results) satisfy Assumption 3 more naturally.

C. Toy Estimation Example

We compare the performance of the set-membership estimator (34) with the Bayesian least squares estimator (35) on the toy problem

$$y(t) = [w_1, w_2] \begin{bmatrix} \sin(4x_1(t)) \\ \tanh(x_2(t)) \end{bmatrix} + v(t). \quad (40)$$

We set $w = [0.5, 0.5]^\top$, sample $v(t) \stackrel{\text{iid}}{\sim} \mathcal{U}[-0.4, 0.4]$ and generate training samples x with each entry $x_i(t) \stackrel{\text{iid}}{\sim} \mathcal{U}[-1, 1]$.




Fig. 2. Left: The trajectory of the parameter estimate found using set-membership with feasible parameter sets plotted in gray at 5 sample intervals. Right: The trajectory of the parameter estimate found using BLS with 95% confidence intervals plotted at 5 sample intervals.

For the set-membership estimator, we set the initial feasible parameter set $\Theta(0) = \{\tilde{w} : \|\tilde{w}\|_\infty \leq 1\}$. For the Bayesian least-squares estimator we select a zero-mean prior such that the 95% confidence interval $\mathcal{W}(0) = \{w : \|w\|_2 \leq 1\}$. Figure 2 shows the evolution of the parameter estimates after processing $t = 25$ samples. The set-membership estimator quickly converges to an accurate estimate as measured by the size of the feasible parameter set. In contrast, the BLS estimate is slower to converge with a looser confidence interval.

However, if we add a small bias of $\alpha = 0.05$ to the measurements and run the experiment again, we observe highly undesirable behavior from the set-membership estimator. As shown in Figure 3, the feasible parameter set shrinks to a set that does not include the true parameter. Moreover, in our experiment the feasible parameter set collapsed to the null set after $t = 13$ measurements, indicating there were no parameters that could explain the data any longer. In contrast, the BLS estimator performs similar to the well-specified example. These experiments highlight the fragility of the set-membership estimator, which we believe highlights the need for greater consideration of probabilistic estimators such as Bayesian regression in adaptive and learning MPC. Model misspecification in the form of imperfect features and outlier noise samples outside of the disturbance set \mathcal{V} can generally result in feasible parameter sets that vanish or are erroneous, resulting in the loss of safety guarantees of an adaptive MPC algorithm. In contrast, our example shows the BLS estimator is more robust to imperfections in the control design, making it a much more practical estimator for real-life applications.

D. Calibrated Priors using Meta-Learning

In our problem formulation and throughout §IV, we considered the uncertain function f in the dynamics (1) as a linear combination of known, fixed basis functions for all time and control iterations. Under this assumption (specifically, Assumption 1), we designed an adaptive robust MPC that provably guarantees the safety of the system when we use the estimators in §V-A-§V-B that satisfy Assumption 3. However, f commonly represents the unknown nonlinear effect of an uncertain environment, so a chosen feature representation ϕ may not satisfy Assumption 1 exactly in practice. Even if trajectory data in a specific environment is available, a learned feature representation might not capture the nonlinear influence of another environment.

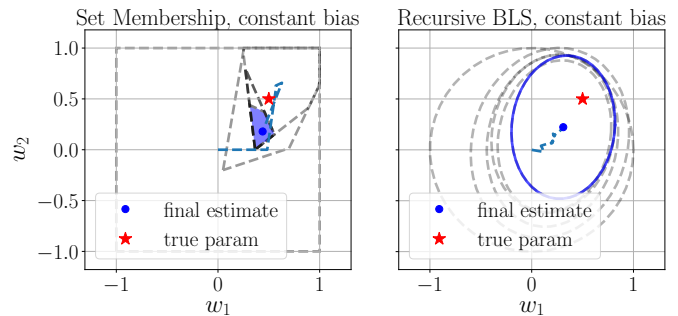


Fig. 3. Left: The trajectory of the parameter estimate found using set-membership with feasible parameter sets plotted in gray at 5 sample intervals. Right: The trajectory of the parameter estimate found using BLS with 95% confidence intervals plotted at 5 sample intervals.

Definition 8: A control task $\mathcal{T} = \{k, x(0), \epsilon\}$ consists of a number of iterations $k \in \mathbb{N}_{\geq 0}$, and an initial condition $x(0)$ and parameter $\epsilon \in \mathbb{R}^e$ that represent the *environment*, drawn from a distribution $\epsilon \sim \rho(\epsilon)$. The *environment* ϵ is constant over a control task and affects the dynamics (1) through the unknown function f , so we take the unknown function as $f(x) := f_\epsilon(x)$.

In applications, we typically either do not exactly know how the environment influences the dynamics or we cannot measure ϵ . Thus, a system designer is often left to choose basis functions and encode prior assumptions on the effect size arbitrarily, which may be difficult or inaccurate for complex systems when the controller is deployed in a wide range of operating conditions.

Therefore, we propose to cast the problem of learning the structure of the disturbance through the lens of meta-learning algorithms—methods that “learn to learn”—which have gained prominence in recent years for their ability to learn general representations that can rapidly adapt to new tasks [46]. These approaches provide an automatic method for identifying features and calibrating priors based on data from the system acting in various environments [34]. Moreover, a practitioner could generate such data from simulation, wherein the assumptions on the system dynamics are automatically mapped to numerically convenient basis functions [13].

In particular, we propose to leverage the ALPaCA meta-learning algorithm [34]. This learning system consists of two components: neural network features, which are fixed for the duration of a control task, and a last layer which is updated online. The model takes the form of a matrix last layer $W \in \mathbb{R}^{n \times d}$, applied to nonlinear neural network features, yielding nonlinear predictive model $\hat{f}(x) = W\phi_\theta(x)$, where θ denotes the neural network weights. The ALPaCA algorithm learns both the feature representation ϕ_θ and a prior distribution on the last layer $p(W)$ to approximate the distribution over environment uncertainties f_ϵ induced by $\epsilon \sim \rho(\epsilon)$ as the distribution over models $W\phi_\theta(x)$ induced by the prior $W \sim p(W)$. Following [13], we represent the prior with independent Gaussian distributions on each row of W .

We outline the ALPaCA algorithm’s mechanics under the assumption that each control task consists of a single iteration of length T to simplify notation. Let $\mathcal{H}^{(j)} = \{(x^{(j)}(0), u^{(j)}(0)), \dots, (x^{(j)}(T), u^{(j)}(T))\}$ be the trajectory data collected during task j . We assume access to a trajectory dataset $\mathcal{H}^{(0)}, \dots, \mathcal{H}^{(k)}$, each associated with a different

unknown instantiation $f^{(0)}(\cdot), \dots, f^{(k)}(\cdot)$ of the uncertain function f in (1) resulting from the changing environments $\epsilon^{(0)}, \dots, \epsilon^{(k)} \stackrel{\text{iid}}{\sim} \rho(\epsilon)$. In the *inner loop* learning algorithm, which is the learning that happens online within one environment, the posterior over W is computed based on an assumption of Gaussian noise, as discussed in §V-B. Then, the *offline outer learning algorithm* backpropagates through the Bayesian linear regression learning procedure to train the neural network features and learn the prior on W . As such, the ALPaCA model learns a set of features that are suited to the problem class expected to occur online, paired with a prior over these features that is calibrated to the environment distribution.

Feature-based function approximation is a well studied problem in classical adaptive control [24], [32], with methods that generally rely on a small number of hand-crafted basis functions. For applications where the dynamics are poorly understood, such as aircraft wake vortices [47], this makes it necessary to reason about the misspecification of the model even in unconstrained control [10], [32]. Another approach is to use a large number of random features that can approximate any function as the number of features goes to infinity [27]. This is undesirable in a constrained control setting, as the confidence intervals we require in §V scale poorly in the number of features. A meta-learning approach like ALPaCA [34] offers a solution to this apparent trade-off, learning an accurate and compact feature representation from data offline. Therefore, we do not consider misspecification of the features in this work, although future work could consider tightening the constraints in the tube MPC (29) further with an empirical bound on $\|\phi_\theta(x) - \phi(x)\|$ obtained, for example, by a bootstrapping procedure.

VI. EXPERIMENTS

In this section, we simulate our approach on several example systems to highlight the various benefits of our methods and illustrate their properties.

A. Adaptive Robust MPC Benchmark

To benchmark our adaptive, robust MPC (ARMPC) algorithm that *matches* as much of the uncertainty as possible, we compare it against the normative approach in ARMPC; we estimate the range of values that the uncertainty f can take and naively treat it as a disturbance using tube MPC as in [6], [7], [14], [31]. These methods improve their performance online by refining a non-increasing bound on the range of f . To the best of our knowledge, such algorithms have not been proposed for uncertain terms that satisfy Assumption 1 in the literature. Therefore, we apply the tube MPC strategy in [38] by solving (6) online using the disturbance set $\mathcal{D}'(t) = \hat{\mathcal{F}}(t) \oplus \mathcal{V}$, since Corollary 1 implies that both $f(x(t)) + v(t) \in \mathcal{D}'(t)$ and that $\mathcal{D}'(t) \subseteq \mathcal{D}'(t-1)$ for all $t \geq 0$. Analogous to our approach, we construct the terminal cost and maximally RPI set $\mathcal{O}(t)$ using $\mathcal{D}'(t)$ and the fixed (LQR) controller in Assumption 4. It is straightforward to verify the recursive feasibility and input-to-state stability of this adaptive benchmark via a symmetric argument to the proofs of Theorem 2, Theorem 3.




Fig. 4. Phase plots of the closed-loop trajectories of different systems. The realized system evolution is shown in black, the predicted nominal trajectories are depicted in red, as are the predicted reachable sets at $t = 0$. The predicted reachable sets at $t = 50$ are depicted in purple. The dark blue line indicates the state constraints. The terminal invariant $\mathcal{O}(50)$ is shown in cyan. Left: Adaptive CE MPC on the system with *matched* uncertainty (41). Middle: Benchmark Adaptive MPC on the *matched* system (41). Right: Adaptive CE MPC on the system with *unmatched* uncertainty (42).

B. Double Integrator with Matched Uncertainty

We illustrate the properties of our algorithm on a double-integrator system

$$x(t+1) = \begin{bmatrix} 1 & 0.2 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t) + f(x(t)) + v(t), \quad (41)$$

a typical example in the MPC literature that represents simplified second-order dynamics [29]. First, we consider the *matched* uncertainty $f(x) = [0, w_1]^\top \tanh([0, 1]x)$. We estimate the true parameter, $w_1 = 0.5$, online to improve performance. We take the disturbance as an isotropic Gaussian with $\sigma^2 = 5 \times 10^{-3}$ clipped at its 95% confidence intervals. In addition, the system is subject to the state and input constraints $(-4, -3) \preceq x \preceq (4, 3)$ and $-2 \leq u \leq 2$, respectively. We regulate the system to the origin from $x(0) = [2, 2]^\top$ while minimizing the quadratic cost function $h(x, u) = x^\top Qx + u^\top Ru$ over a horizon of length $N = 3$. We take $Q = I_2$, $R = 1$. We use the BLR estimator with a flat prior and collect $k = 45$ data points of the system evolution in a unit box near the origin to form an initial estimate of the model parameters and construct $\mathcal{O}(t)$ using Alg. 10.4 in [29].

We plot the closed-loop system evolution in Figure 4 (left, middle). Our adaptive certainty-equivalent MPC algorithm is able to effectively control the system. In contrast, the benchmark ARMPC can only react to the learned dynamics after f enters the system as a disturbance, resulting in considerably larger closed-loop oscillations and uncertainty on the predicted trajectory. In Figure 4, we also plot the reachable sets associated with the predicted trajectory at the first and last timesteps of the control task. The reachable sets show that our adaptive MPC resolves prediction uncertainty in the system since they shrink over time. Moreover, we emphasize that online learning does not noticeably improve the performance of the benchmark, as increasing the confidence in the model does not significantly reduce the estimated range of values that the nonlinear function takes.

Furthermore, we see that the shape of the terminal invariant of our adaptive MPC is qualitatively different from that of the benchmark in Figure 4 (left, middle). This is because in our method, we tighten the input constraints of the robust input u^* in the CE law (15) to account for the imperfect matching using \hat{f} . Our method therefore relinquishes some nominal control authority so that it can account for a smaller disturbance set when we are sufficiently confident in the model. To see this, note that when $\mathcal{W} = \{0\}$, i.e., when $\widehat{W} = W$, it holds that




Fig. 5. Left: Closed-loop realized trajectory cost for the matched system (41) as a function of w_1 . Exploding cost indicates infeasibility. The error bars indicate 2σ bounds. Right: Solid lines indicate size of feasible envelopes as a function of w_1 for a matched system. Dashed lines indicate the size of the feasible envelopes for the unmatched system (42) as a function of the magnitude of the unmatched dynamics w_1 with $w_2 = 0.5$.

$\widehat{\mathcal{D}} = (I - BB^\dagger)\widehat{\mathcal{F}} \oplus \mathcal{V} \subseteq \widehat{\mathcal{F}} \oplus \mathcal{V} = \mathcal{D}'$. The simulation shows that matching uncertainty improves the transient convergence behavior of the system.

In addition, we compare the asymptotic performance of our algorithm with the benchmark ARMPC as a function of w_1 , controlling the magnitude of the nonlinearity. To do this, we set the number of data points used to generate the initial model estimate to $k = 10^3$ and collect trajectory rollouts for various values of w_1 from the fixed initial condition. As Figure 5 (left) shows, our ARMPC is feasible for disturbances more than twice the magnitude of those the benchmark ARMPC can tolerate for the given initial condition. In addition, the realized control cost does not differ significantly from the benchmark for values of w_1 when both controllers are feasible.

Finally, we illustrate how our algorithm tolerates larger dynamic uncertainty by comparing the size of the feasible envelope (i.e., the set of initial conditions for which the MPC problem is feasible) as a function of w_1 . We set the number of data points to inform our prior to a modest $k = 50$, grid the state-space, and take the feasible region as the convex hull of the initial conditions for which the MPC problem is feasible. Then, we estimate the percentage of states $x_0 \in \mathcal{X}$ in the feasible envelope as the ratio of volumes between the feasible envelope and the state space \mathcal{X} , illustrated with solid lines in Figure 5 (right). Our CE ARMPC algorithm can tolerate much larger disturbances than the benchmark ARMPC. In these experiments, the feasible envelope of the benchmark becomes empty when the maximal robust invariant is null ($\mathcal{O}(t) = \emptyset$), indicating that there is no subset of \mathcal{X} in which the LQR policy associated with the stage cost h results in provably safe behavior [29]. Hence, Figure 5 highlights the fragility of existing ARMPC approaches under large disturbances.

C. Double Integrator with Unmatched Uncertainty

We now extend the simulations of the simple matched system to the unmatched case to understand the effect of additive nonlinear terms that cannot be canceled from the dynamics. We keep the nominal dynamics identical to (41) and take the nonlinear dynamics

$$f(x) = \frac{1}{\sqrt{2}}[w_1 \sin(4x_1), w_2 \tanh(x_2)]^\top, \quad (42)$$

where w_1 and w_2 are unknown parameters. Similar to the previous experiments, we initialize the model with $k = 45$ data points sampled around the origin. In this example, the certainty equivalent policy (15) can only compensate for the

second component of the nonlinear dynamics (42). We set $w_1 = 0.2$ and $w_2 = 0.3$. In Figure 4 (right), the size of the reachable sets increases if we simulate the system with the unmatched dynamics (42). The benchmark ARMPC was infeasible from this initial condition, showing that our method still outperforms the benchmark. Next, we fix $w_2 = 0.5$ and vary w_1 to understand the impact of an estimated, unmatched dynamics component. Figure 5 (dashed, right) shows that in our experiment, matching as much of the nonlinear dynamics as possible allows us to handle unmatched dynamics terms of about twice the magnitude as the benchmark. We conclude that our method is a more effective strategy even if the uncertainty is unmatched. Naturally, Figure 5 (right) also shows that the absolute benefit of our method diminishes as the proportion of the nonlinear dynamics $f(x(t))$ in $\text{Range}(B)$ becomes smaller.

D. Controlling a Planar Quadrotor

Next, we simulate a simplified example of a quadrotor in a windy environment. For this example, we examine the episodic setting discussed in §IV-D combined with an application of the meta-learning algorithms in §V-D. We consider a planar version of the quadrotor dynamics for simplicity [48], with 2D pose (p_x, p_y, θ) and front and rear thrust inputs u_f, u_r :

$$\begin{aligned} \ddot{p}_x &= -\frac{1}{m} \sin(\theta)(u_f + u_r) + \frac{1}{m} f_x(p_x, p_y, \theta) \\ \ddot{p}_y &= \frac{1}{m} \cos(\theta)(u_f - u_r) + \frac{1}{m} f_y(p_x, p_y, \theta) - g \\ \ddot{\theta} &= \frac{l}{I}(u_f - u_r) \end{aligned}$$

We take the state as $x = [p_x, p_y, \theta, \dot{p}_x, \dot{p}_y, \dot{\theta}]^\top$, linearize the dynamics around $\bar{x} = 0$, $\bar{u} = \frac{mg}{2}[1, 1]^\top$, discretize the simulation using Euler's method, and add process noise. The force field f induced by the wind varies spatially, modelling real-world scenarios such as down-wash from another quadrotor. We model the 2D wind disturbance force $f(x) = [f_x(x), f_y(x)]^\top$ as incident at a fixed angle with a maximum velocity V_w that drops of according to an inverse square exponential normal to the direction of incidence θ_w , i.e.,

$$v_w(x) = V_w \exp\left(-\frac{1}{2}(p_y \cos(\theta_w) - p_x \sin(\theta_w))^2\right), \quad (43)$$

resulting in a disturbance force of $\|f(x)\| = clv_w(x)^2$ along θ_w for length parameter $l = 0.4$ m and air resistance $c = 0.5 \text{ N} \cdot \text{s}^2/\text{m}^3$ (we consider the drone's velocity negligible). We only set constraints on the pose (p_x, p_y, θ) of the drone. Its linear and angular velocities are unconstrained. The quadrotor is an underactuated system, and therefore the discretized, linearized, simulation has unmatched dynamics terms. Still, similar to the illustration in Figure 6, a drone controller can always match disturbance forces along the y axis in the linearized simulation.

Spatially varying wind disturbances are challenging to model in practice, so we take a Bayesian approach and model the unknown wind disturbance using a feedforward neural network. To do this, we model the wind disturbance as a linear combination of $d = 5$ sigmoidal output activation functions




Fig. 6. Learned trajectories of our adaptive MPC (blue) compared to a naive tube MPC (red) in the xy -plane for a simulated planar quadrotor, with axes in meters. The icons show the orientation of the quadrotor over time. The contours and colorbar indicate the *learned* wind speed in m/s. Left: Wind comes straight from above. Right: Wind comes at $\theta_w = 22.5^\circ$.

(scaled by $1/\sqrt{d}$) of a 2 layer network with hidden ReLU activations.

We generate training data by sampling wind disturbance fields uniformly with wind speeds between 3 m/s and 5 m/s coming in at angles of incidence between -45° and 45° relative to the y -axis. We use an ALPaCA model [34] to learn an efficient Bayesian representation that can rapidly adapt to any specific conditions encountered in the wild.

After meta-learning, we control the system for several iterations in the same wind environment using Algorithm 1.C. We compare our adaptive robust MPC with a naive tube MPC similar to the controller in §III that disregards the wind disturbance in the control design, since the support of the estimated wind disturbance was too large for our benchmark ARMPC to be feasible in our experiments. Therefore, this simulation highlights the fragility of the benchmark approach, which primarily results from non-existence of a terminal invariant set under disturbances of considerable magnitude [29]. This phenomenon was also observed previously in applied work [30], requiring practical workarounds that lose safety guarantees.

We run our algorithm with the learned features for 5 iterations under the same wind conditions. As shown in Figure 6 (left), if the wind disturbance comes straight from above, the adaptive MPC learns to match the wind forces and reaches the origin quickly. In contrast, the naive tube MPC drifts significantly, as it only reacts to observed disturbances. In addition, if we set the angle of incidence of the wind as $\theta_w = 22.5^\circ$, Figure 6 (right) shows that our approach still achieves decent control performance. The certainty equivalent controller (15) cancels the y -component of the disturbance and converges to a small steady-state offset in the x direction, showcasing the intuitive behavior of the matching CE controller compared to a naive tube MPC. In contrast, the benchmark ARMPC algorithm could not guarantee safety for any of the tasks, and a naive unsafe tube MPC that does not consider the wind disturbance at all performs poorly.

Finally, we benchmark the efficacy of Bayesian meta-learning algorithms on this problem. We compare the performance of our meta-learned priors against two frequentist counterparts: A frequentist version of ALPaCA that only meta-learns features using a mean-square error (MSE) metric as considered in [28], and a naive model that does not consider the fact that the uncertain function changes between contexts in the training data, directly minimizing the aggregate MSE.




Fig. 7. Average closed-loop control cost of the learned adaptive controllers for the simplified drone experiment per control iteration. Error bars indicate 50% quantiles over 10 experiments. Blue: Controller using Bayesian meta-learning. Red: Controller using frequentist meta-learned features. Green: Controller using features that are not meta-learned.

To engineer a prior such that our safety guarantees hold for these frequentist models, we fix a zero mean prior with a fixed covariance and pre-train the frequentist models on $k = 200$ samples of the test environment—the vertical disturbance in Figure 6 (left)—such that the bounds on the unknown function match an a priori known bound (as in Remark 4, we then also use these bounds in the control design). Our Bayesian model does not receive any additional pre-training data before starting the task.

As shown in Figure 7, the Bayesian meta-learned features and prior result in a model that is able to adapt rapidly to the control task, achieving high performance within a single iteration. This showcases that the ALPaCA meta-learning algorithm can effectively learn a compact feature representation that can quickly learn in new contexts with high confidence. In contrast, the engineered priors on the frequentist models need several trajectory rollouts to drive the cost down.

E. Longitudinal Cruise Control

Finally, we consider the design of a longitudinal cruise controller for an autonomous vehicle (AV). The cruise controller needs to track a fixed reference velocity v_{ref} and can control the torque τ applied to the wheels of the vehicle, resulting in the second order dynamics

$$\ddot{p} = \frac{1}{m}(R\tau - k\dot{p} + F_{\text{dist}}(p, \dot{p})), \quad (44)$$

where p is the position of the car, $m = 10^3$ kg is its mass, $R = 1/3$ m is the wheel radius, $k = 20$ N·s/m is the friction coefficient, and $F_{\text{dist}}(p, \dot{p})$ is an external disturbance force [49]. In this example, we consider the disturbances caused by hilly terrain that the AV passes by on a daily commute, and write the disturbance as induced by an unknown constant incline of θ_k on the k 'th road segment $p \in [p_k^1, p_k^2]$ as

$$F_{\text{dist}}(p, \dot{p}) = -\frac{mg}{2} \sum_k \sin(\theta_k) \underbrace{(\tanh(p - p_k^1) + \tanh(p_k^2 - p))}_{=: 2\sqrt{k}\phi_k(p, \dot{p})} \quad (45)$$

and ignore the air-resistance term presented in [49] for simplicity. We set $\tau = \bar{\tau} - \frac{m}{R}u$, with u as our pseudo-acceleration input into the system around the reference input $\bar{\tau} = \frac{k}{R}v_{\text{ref}}$, yielding the reference tracking dynamics $\dot{x} = v_{\text{ref}} - \dot{p}$ as

$$\dot{x} = -\frac{k}{m}x + u - \frac{1}{m}F_{\text{dist}}(p, v_{\text{ref}} - x).$$




Fig. 8. Left: Velocity reference tracking performance over time in seconds. Right: Pseudo-acceleration command over time in seconds. Black: state and input constraints. Red: benchmark ARMPC. Blue: Our CE ARMPC. Red dashed: Reference. Orange: Indicates time spent on inclined road segments.

We discretize the dynamics with timestep $dt = .1$ s and inject isotropic Gaussian noise with $\sigma^2 = 10^{-3}$ into the position and velocity dynamics. In the spirit of Remark 7, we treat the position of the vehicle as an *exogenous* signal, regulating and constraining only the velocity and inputs of the vehicle. We collect a 40 s trajectory of training data to initialize a BLR prior over the features defined in (45). In Figure 8, we compare the performance of our CE ARMPC with the benchmark, and see that the tracking performance of the controller is improved by learning that the uncertainty in the dynamics is matched.

Oscillations in the speed of the vehicle cause passenger discomfort and increase the likelihood of traffic jams. Therefore, we quantify the benefit of our method over the benchmark by comparing the cumulative sum of squared accelerations on the vehicle over the route, a common measure of ride quality similar to ISO standards [50]. The lower this quantity is, the better the ride quality. Figure 9 shows that our approach improves the ride quality by about 30%, a significant improvement. However, Figure 9 also shows that this discrepancy can partially be attributed to the fact that our ARMPC accelerates less aggressively from the initial velocity of 85 km/h in the first 3 seconds of the simulation. The initial difference in acceleration is caused by the fact that we tightened the MPC’s input constraints in (29) to account for the matching input in the CE policy (15). This is visible in Figure 8 (right), where the control input applied by our method saturates at a lower constant than the benchmark’s inputs in the first second of the simulation. Still, after the initial acceleration, our method clearly outperforms the benchmark on the hilly sections of the route, as Figure 9 shows the relative benefit of our method over the benchmark grows throughout the simulation.

VII. DISCUSSION AND CONCLUSIONS

The simulations in §VI show that our method achieves substantial performance improvements compared to existing ARMPC approaches, even when significant components of the nonlinear dynamics are unmatched and cannot be cancelled by a certainty equivalent control policy. We conclude that by extending certainty equivalent control laws from classical adaptive control, we can reduce the conservatism of robust MPC approaches. As a result, we saw in our simulations that our method can tolerate more significant nonlinear terms in the dynamics. In addition, our experiments show that applying the Bayesian meta-learning algorithm ALPaCA [34] allows us to learn a feature based representation and a Bayesian prior on the last-layer weights that serve as a sensible engineering solution




Fig. 9. Cumulative sum of squared accelerations experienced by the car in time (seconds), a measure of ride-quality. Blue: Our CE ARMPC. Red: Benchmark ARMPC

to satisfy the initial assumptions we make to guarantee the safety of our controller. Moreover, we saw that these models can rapidly adapt to the environmental conditions encountered during deployment.

Since our control algorithm allows for adaptation laws based on statistical estimation techniques that are more robust to outliers than set-membership estimation, future work should extend our simulations to hardware experiments. In addition, we did not consider a setting in which the uncertain function can change during the control task. Future work can extend our analysis to slowly changing environments, commonly considered using exponential forgetting in adaptive control [24]. In addition, many systems of interest are inherently nonlinear, therefore, future work should consider applying our method in a constrained nonlinear MPC setting. Furthermore, as we have assumed in this work, exact measurement of the state is often not possible in applications. Extending learning-based MPC methods for nonlinear systems to the output feedback setting remains a highly complex and largely open problem, which we hope to investigate in future work.

REFERENCES

- [1] R. Sinha, J. Harrison, S. M. Richards, and M. Pavone. Adaptive robust model predictive control with matched and unmatched uncertainty. Under review. Available at <https://arxiv.org/abs/2104.08261>.
- [2] Carl Edward Rasmussen. *Gaussian processes in machine learning*. MIT press, 2006.
- [3] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [4] S. M. Richards, N. Azizan, J.-J. E. Slotine, and M. Pavone. Adaptive-control-oriented meta-learning for nonlinear systems. In *RSS*, 2021.
- [5] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin. Provably safe and robust learning-based model predictive control. *Automatica*, 2013.
- [6] M. Bujarbaruah, X. Zhang, U. Rosolia, and F. Borrelli. Adaptive MPC for Iterative Tasks. In *CDC*, 2018.
- [7] M. Bujarbaruah, S. H. Nair, and F. Borrelli. A Semi-Definite Programming Approach to Robust Adaptive MPC under State Dependent Uncertainty. In *ECC*, 2020.
- [8] M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli. Adaptive stochastic MPC under time varying uncertainty. *TAC*, 2021. In press.
- [9] G. Chowdhary and E. Johnson. Recursively updated least squares based modification term for adaptive control. In *ACC*, 2010.
- [10] G. Joshi, J. Virdi, and G. Chowdhary. Asynchronous Deep Model Reference Adaptive Control. *arXiv:2011.02920*, 2020.
- [11] L. Hewing, J. Kabzan, and M. N. Zeilinger. Cautious model predictive control using Gaussian process regression. *TAC*, 2017.
- [12] Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):253–279, 2019.
- [13] T. Lew, A. Sharma, J. Harrison, and M. Pavone. Safe Model-Based Meta-Reinforcement Learning: A Sequential Exploration-Exploitation Framework. *arXiv:2008.11700*, 2020.
- [14] R. Soloperto, M. A. Müller, S. Trimpe, and F. Allgöwer. Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty. *IFAC NMPC*, 2018.

- [15] S. Di Cairano. Indirect adaptive model predictive control for linear systems with polytopic uncertainty. In *ACC*, 2016.
- [16] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *NeurIPS*, 2018.
- [17] S. Dean, S. Tu, N. Matni, and B. Recht. Safely learning to control the constrained linear quadratic regulator. In *ACC*, 2019.
- [18] D. D. Fan, J. Nguyen, R. Thakker, N. Alatur, A.-a. Agha-mohammadi, and E. A. Theodorou. Bayesian Learning-Based Adaptive Control for Safety Critical Systems. In *ICRA*, 2020.
- [19] J. Harrison, A. Sharma, R. Calandra, and M Pavone. Control adaptation via meta-learning dynamics. In *NeurIPS Workshop on Meta-Learning*, 2018.
- [20] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause. Learning-Based Model Predictive Control for Safe Exploration. In *CDC*, 2018.
- [21] H. Mania, M. I. Jordan, and B Recht. Active learning for nonlinear system identification with guarantees, 2020.
- [22] P.K. Mishra, M.V. Gasparino, A. Velsasquez, and G. Chowdhary. Deep model predictive control with stability guarantees, 2021.
- [23] K. J. Åström and B. Wittenmark. *Adaptive control*. Courier, 2013.
- [24] J.-J. E. Slotine and W. Li. *Applied Nonlinear Control*. Prentice Hall, 1991.
- [25] P. Ioannou and J. Sun. *Robust Adaptive Control*. Dover Publications, 2012.
- [26] G. Chowdhary, H. A. Kingravi, J. P. How, and P. A. Vela. Bayesian nonparametric adaptive control using Gaussian processes. *IEEE Transactions on Neural Networks and Learning Systems*, 2014.
- [27] N. M. Boffi, S. Tu, and J.-J. E. Slotine. Regret bounds for adaptive nonlinear control. Available at <https://arxiv.org/abs/2011.13101>, 2020.
- [28] M. O’Connell, G. Shi, X. Shi, and S.-J. Chung. Meta-learning-based robust adaptive flight control under uncertain wind conditions. Available at <https://arxiv.org/abs/2103.01932>, 2021.
- [29] F. Borrelli, A. Bemporad, and M. Morari. *Predictive control for linear and hybrid systems*. Cambridge Univ. Press, 2017.
- [30] Anil Aswani, Patrick Bouffard, and Claire Tomlin. Extensions of learning-based model predictive control for real-time application to a quadrotor helicopter. In *ACC*, 2012.
- [31] J. Köhler, E. Andina, R. Soloperto, M. Müller, and F. Allgöwer. Linear robust adaptive model predictive control: Computational complexity and conservatism. In *CDC*, 2019.
- [32] P. Ioannou and B. Fidan. *Adaptive Control Tutorial*. SIAM, 2006.
- [33] E. Lavretsky and K. Wise. *Robust and Adaptive Control With Aerospace Applications*. Springer, 2013.
- [34] J. Harrison, A. Sharma, and M. Pavone. Meta-learning Priors for Efficient Online Bayesian Regression. In *WAFR*, 2018.
- [35] P. O. M. Scokaert and D. Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *TAC*, 1998.
- [36] D. Q. Mayne, M. M. Seron, and S. V. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 2005.
- [37] D. Limon, T. Alamo, D. M. Raimondo, D. Muñoz de la Peña, J. M. Bravo, A. Ferramosca, and E. F. Camacho. Input-to-state stability: A unifying framework for robust model predictive control. In *Nonlinear Model Predictive Control: Towards New Challenging Applications*, volume 384 of *Lecture Notes in Control and Information Sciences*, pages 1–26. Springer, 2009.
- [38] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 2006.
- [39] H. Li, A. Liu, and L. Zhang. Input-to-state stability of time-varying nonlinear discrete-time systems via indefinite difference Lyapunov functions. *ISA Transactions*, 2018.
- [40] Z.-P. Jiang and Y. Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 2001.
- [41] Gary Balas. Linear, parameter-varying control and its application to aerospace systems. *ICAS congress proceedings*, 2002.
- [42] Mario Milanese and Antonio Vicino. Optimal estimation theory for dynamic systems with set membership uncertainty: An overview. *Automatica*, 1991.
- [43] Mario Milanese, John Norton, Hélène Piet-Lahanier, and Éric Walter. *Bounding approaches to system identification*. Springer, 2013.
- [44] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge Univ. Press, 2004.
- [45] M. R. Deisenroth, A. A. Faisal, and C. S. Ong. *Mathematics for machine learning*. Cambridge University Press, 2020.
- [46] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [47] E. Lavretsky. Adaptive control: Introduction, overview, and applications. *ACC Robust and Adaptive Control Workshop*, 2008.
- [48] R. Tedrake. Underactuated robotics: Algorithms for walking, running, swimming, flying, and manipulation. Available at <http://underactuated.mit.edu>, 2021.
- [49] Rajesh Rajamani. *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [50] Lars Svensson and Jenny Eriksson. Tuning for ride quality in autonomous vehicle: Application to linear quadratic path planning algorithm, 2015.



Rohan Sinha is a Ph.D. candidate in the Autonomous Systems Lab at Stanford University. He received a B.S. in Mechanical Engineering and B.A. in Computer Science in 2020, both from the University of California, Berkeley. Rohan’s research interests lie at the intersection of control theory, machine learning, and applied robotics. Currently, his research focuses on developing learning-based control algorithms with safety guarantees.



James Harrison is a Ph.D. candidate in the Autonomous Systems Lab at Stanford University. He received an M.S. degree from Stanford University in 2018 and a B.Eng. degree from McGill University in 2015, both in mechanical engineering. His research interests include few-shot, adaptive, and open-world learning, Bayesian deep learning, and applications in safe robot autonomy, decision-making, and control.



Spencer M. Richards is a Ph.D. candidate in the Autonomous Systems Lab at Stanford University. He received an M.Sc. degree in Robotics, Systems, and Control from ETH Zürich in 2018, and a B.A.Sc. degree in Engineering Science from the University of Toronto in 2016. His research interests lie at the intersection of control theory and machine learning for robotics, where he works to blend ideas such as adaptive control with meta learning, and stability theory with reinforcement learning.



Marco Pavone is an Associate Professor of Aeronautics and Astronautics at Stanford University, where he is the Director of the Autonomous Systems Laboratory. Before joining Stanford, he was a Research Technologist within the Robotics Section at the NASA Jet Propulsion Laboratory. He received a Ph.D. degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology in 2010. His main research interests are in the development of methodologies for the analysis, design, and control of autonomous systems, with an emphasis on self-driving cars, autonomous aerospace vehicles, and future mobility systems. He is a recipient of a number of awards, including a Presidential Early Career Award for Scientists and Engineers, an ONR YIP Award, an NSF CAREER Award, and a NASA Early Career Faculty Award. He was identified by the American Society for Engineering Education (ASEE) as one of America’s 20 most highly promising investigators under the age of 40. He is currently serving as an Associate Editor for the *IEEE Control Systems Magazine*.