Robustness and Consistency in Linear Quadratic Control with Untrusted Predictions

TONGXIN LI, California Institute of Technology, United States

RUIXIAO YANG, Tsinghua University, China

GUANNAN QU, Carnegie Mellon University, United States

GUANYA SHI, California Institute of Technology, United States

CHENKAI YU, Columbia University, United States

ADAM WIERMAN, California Institute of Technology, United States

STEVEN LOW, California Institute of Technology, United States

We study the problem of learning-augmented predictive linear quadratic control. Our goal is to design a controller that balances consistency, which measures the competitive ratio when predictions are accurate, and robustness, which bounds the competitive ratio when predictions are inaccurate. We propose a novel λ -confident controller and prove that it maintains a competitive ratio upper bound of $1 + \min\{O(\lambda^2 \varepsilon) + O(1 - \lambda)^2, O(1) + O(\lambda^2)\}$ where $\lambda \in [0, 1]$ is a trust parameter set based on the confidence in the predictions, and ε is the prediction error. Further, we design a self-tuning policy that adaptively learns the trust parameter λ with a competitive ratio that depends on ε and the variation of system perturbations and predictions.

CCS Concepts: • **Theory of computation** → **Online learning algorithms**; *Regret bounds*.

Additional Key Words and Phrases: online control; model predictive control; online learning; competitive analysis

ACM Reference Format:

Tongxin Li, Ruixiao Yang, Guannan Qu, Guanya Shi, Chenkai Yu, Adam Wierman, and Steven Low. 2022. Robustness and Consistency in Linear Quadratic Control with Untrusted Predictions. *Proc. ACM Meas. Anal. Comput. Syst.* 6, 1, Article 18 (March 2022), 35 pages. <https://doi.org/10.1145/3508038>

1 INTRODUCTION

We study a classical online linear quadratic control problem where the controller has access to untrusted predictions/advice during each round, potentially from a black-box AI tool.

One consequence of the success of machine learning is that accurate predictions are available for many online decision and control problems. For example, the development of deep learning has enabled generic prediction models in various domains, e.g. weather, demand forecast, and

This work is supported by the National Science Foundation, under grants ECCS1931662, CCF 1637598, ECCS 1619352, CPS 1739355, AitF-1637598, CNS-1518941, PIMCO and Amazon Web Services. Tongxin Li and Ruixiao Yang contributed equally to the paper.

Authors' addresses: Tongxin Li, California Institute of Technology, Pasadena, United States, tongxin@caltech.edu; Ruixiao Yang, Tsinghua University, Beijing, China, yangruixiao@foxmail.com; Guannan Qu, Carnegie Mellon University, Pittsburgh, United States, gqu@andrew.cmu.edu; Guanya Shi, California Institute of Technology, Pasadena, United States, gshi@caltech.edu; Chenkai Yu, Columbia University, New York City, United States, cyu26@gsb.columbia.edu; Adam Wierman, California Institute of Technology, Pasadena, United States, adamw@caltech.edu; Steven Low, California Institute of Technology, Pasadena, United States, slow@caltech.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

2476-1249/2022/3-ART18

<https://doi.org/10.1145/3508038>

user behaviors. Such predictions are powerful because future information plays a significant role in optimizing the current control decision. The availability of accurate future predictions can potentially lead to order-of-magnitude performance improvement in decision and control problems, where one can simply plug-in the predictions and achieve near optimal performance when compared to the best control actions in hindsight, a.k.a., *consistency*. However, an important caveat is that the predictions are helpful only when they are accurate, which is not guaranteed in many scenarios. Since many predictions are obtained from black box AI models like neural networks, there is no uncertainty quantification and it is unclear whether the predictions are accurate. In the case when the predictions are not accurate, the consequences can be catastrophic, leading to unbounded worst-case performance, e.g., an unbounded competitive ratio. The possibility of such worst-case unbounded competitive ratio prevents the use of ML predictions in safety-critical applications that are adverse to potential risks.

The use of predictions described above is a sharp contrast to the approaches developed by the online algorithm community, where the algorithms have access to no future prediction, yet can be *robust* to all future variations and achieve a finite competitive ratio. While such algorithms miss out on the improvements possible when accurate predictions are available, their robustness properties are necessary in safety-critical settings. Therefore, a natural question arises:

Can such adversarial guarantees be provided for control policies that use black-box AI predictions?

To provide adversarial guarantees necessarily means not precisely following the black-box AI predictions. Thus, there must be a trade-off between the performance in the typical case (consistency) and the quality of the adversarial guarantee (robustness). Trade-offs between consistency and robustness have received considerable attention in recent years in the online algorithms community, starting with the work of [20], but our work represents the first work in the context of control.

Contributions. In this paper, we answer the question above in the affirmative, in the context of linear quadratic control, providing an novel algorithm that trades off consistency and robustness to provide adversarial guarantees on the use of untrusted predictions.

Our first result provides a novel online control algorithm, termed λ -*confident control*, that provides a competitive ratio of $1 + \min\{O(\lambda^2\epsilon) + O(1-\lambda)^2, O(1) + O(\lambda^2)\}$, where $\lambda \in [0, 1]$ is a *trust parameter* set based on the confidence in the predictions, and ϵ is the prediction error (Theorem 3.2). When the predictions are accurate ($\epsilon \approx 0$), setting λ close to 1 will obtain a competitive ratio close to 1, and hence the power of the predictions is fully utilized; on the other hand, when the predictions are inaccurate (ϵ very large), setting $\lambda \approx 0$ will still guarantee a constant competitive ratio, meaning the algorithm will still have good robustness guarantees when the predictions turn out to be bad. Therefore, our approach can get the best of both worlds, effectively using black-box predictions but still guaranteeing robustness.

The above discussion highlights that the optimal choice of λ depends on the prediction error, which may not be known a priori. Therefore, we further provide an adaptive, self-tuning learning policy (Algorithm 3) that selects λ so as to learn the optimal parameter for the actual prediction error; thus selecting the optimal balance between robustness and consistency. Our main result proves that the self-tuning policy maintains a competitive ratio bound that is always bounded regardless of the prediction error ϵ (Theorem 4.1). This result is informally presented below.

THEOREM (INFORMAL). *Under our model assumptions, there is a self-tuning online control algorithm that selects some $\lambda_t \in [0, 1]$ for all $t = 0, \dots, T - 1$ and achieves a competitive ratio*

$$CR(\epsilon) \leq 1 + \frac{O(\epsilon)}{\Theta(1) + \Theta(\epsilon)} + O(\mu_{\text{Var}})$$

as a function of the prediction error ϵ where μ_{Var} measures the variation of perturbations and predictions.

This result provides a worst-case performance bound for the use of untrusted predictions, e.g., the predictions from a black-box AI tool, regardless of the accuracy of the predictions. The second term in the competitive ratio upper bound indicates a nontrivial non-linear dependency of $CR(\epsilon)$ and the prediction error ϵ , matching our experimental results shown in Section 5. The third term measures the variation of perturbations and predictions. Such a term is common in regret analysis based on the “Follow The Leader” (FTL) approach [14, 15]. For example, the regret analysis of the Follow the Optimal Steady State (FOSS) method in [19] contains a similar “path length” term that captures the variation of the state trajectory.

Proving our main result is complex due to the fact that, different from classical online learning models, the cost function in our problem depends on previous actions via a linear dynamical system (see (1)). The time coupling can even be exponentially large if the dynamical system is unstable. To tackle this time-coupling structure, we develop a new proof technique that relates the regret and competitive ratio with the convergence rate of the trust parameter.

Finally, in Section 5 we demonstrate the effectiveness of our self-tuning approach using three examples: a robotic tracking problem, an adaptive battery-buffered EV charging problem and the Cart-Pole problem. For the robotic tracking and adaptive battery-buffered EV charging cases, we illustrate that the competitive ratio of the self-tuning policy performs nearly as well as the lower envelope formed by picking multiple trust parameters optimally offline. We also validate the practicality of our self-tuning policy by showing that it not only works well for linear quadratic control problems; it also performs well in the nonlinear Cart-Pole problem.

Related Work. Our work contributes to the growing literature on learning-augmented online algorithm design. There has been significant interest in the goal of trading-off consistency and robustness in order to ensure worst-case performance bounds for black-box AI tools in online problems. As discussed earlier, prediction based algorithms can achieve consistency, while online algorithms can have robustness. These two classes of algorithms can be viewed as two extremes, and a number of works attempt to develop algorithms that balance between consistency and robustness in settings like online caching [20], ski-rental [4, 7, 22, 27], online set cover [7], secretary and online matching [6], and metric task systems [5]. For example, in the ski rental problem, [22] proposes an algorithm that achieves $1 + \lambda$ consistency and $1 + \frac{1}{\lambda}$ robustness for a tuning parameter $\lambda \in (0, 1)$. Compared to these works, our setting is fundamentally more challenging because of the existence of dynamics in the control problem couples all decision points, and a mistake at one time can be magnified and propagated to all future time steps.

Our work is also closely related to a broad literature on regret and competitive ratio analysis for Linear Quadratic Control (LQC) and Linear Quadratic Regulator (LQR) systems with predictions. In [29], LQR regret analysis for Model Predictive Control (MPC) is given, assuming accurate predictions of perturbations. Inaccurate predictions are considered in [28], with competitive results provided. It is proven in [28, 29] that the action generated by MPC can be explicitly written as the action of optimal linear control plus a linear combination of inaccurate predictions. The competitive analysis of the consistent and robust control scheme in this work makes use of this fact in the analysis of the more challenging case of untrusted predictions. Other related regret and competitive ratio results for MPC include [18, 24, 30].

While our work is the first to study learning-augmented control via the lens of robustness and consistency, there are two classical communities in control that are related to the goals of our work: robust control and adaptive control.

Robust control is a large area that concerns the design of controllers with performance guarantees that are robust against model uncertainty or adversarial disturbances [13]. Tools of robust control include H_∞ synthesis [12, 31] and robust MPC [8]. Like the robust control literature, our work also considers robustness, but our main focus is on balancing between robustness and consistency in a

predictive control setting. Consistency is not a focus of the robust control literature. Further, we focus on the metrics of competitive ratio and regret, which is different from the typical performance measures in the robust control literature, which focus on measures such as system norms [31].

The design of our self-tuning control in Section 4 falls into the category of adaptive control. There is a rich body of literature studying Lyapunov stability and asymptotic convergence in adaptive control theory [26]. Recently, there has been increasing interest in studying adaptive control with non-asymptotic metrics from learning theory. Typical results guarantee convergence in finite time horizons using measures such as regret [1, 2, 11, 25], dynamic regret [19, 29, 30], and competitive ratio [24, 28]. Different from these works, this paper deploys an adaptive policy with the goal of balancing robustness and consistency. Additionally, such results do not focus on incorporation of untrusted predictions.

2 PRELIMINARIES

We consider a Linear Quadratic Control (LQC) model. Throughout this paper, $\|\cdot\|$ denotes the ℓ_2 -norm for vectors and the matrix norm induced by the ℓ_2 -norm. Denote by $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^m$ the system state and action at each time t . We consider a linear dynamic system with adversarial perturbations,

$$x_{t+1} = Ax_t + Bu_t + w_t, \text{ for } t = 0, \dots, T-1, \quad (1)$$

where $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$, and $w_t \in \mathbb{R}^n$ denotes some unknown perturbation chosen adversarially. We make the standard assumption that the pair (A, B) is stabilizable. Without loss of generality, we also assume the system is initialized with some fixed $x_0 \in \mathbb{R}^n$. The goal of control is to minimize the following quadratic costs given matrices A, B, Q, R :

$$J := \sum_{t=0}^{T-1} (x_t^\top Q x_t + u_t^\top R u_t) + x_T^\top P x_T,$$

where $Q, R > 0$ are positive definite matrices, and P is the solution of the following discrete algebraic Riccati equation (DARE), which must exist because (A, B) is stabilizable and $Q, R > 0$ [13].

$$P = Q + A^\top P A - A^\top P B (R + B^\top P B)^{-1} B^\top P A.$$

Given P , we can define $K := (R + B^\top P B)^{-1} B^\top P A$ as the optimal LQC controller in the case of no disturbance ($w_t = 0$). Further, let $F := A - BK$ be the closed-loop system matrix when using $u_t = -Kx_t$ as the controller. By [13], F must have a spectral radius $\rho(F)$ less than 1. Therefore, Gelfand's formula implies that there must exist a constant $C > 0$, $\rho \in (0, 1)$ s.t. $\|F^t\| \leq C\rho^t, \forall t \geq 0$.

Our model is a classical control model [16] and has broad applicability across various engineering fields. In the following, we introduce ML/AI predictions into the classical model and study the trade-off between consistency and robustness in this classical model for the first time.

2.1 Untrusted Predictions

Our focus is on predictive control and we assume that, at the beginning of the control process, a sequence of predictions of the disturbances $(\hat{w}_0, \dots, \hat{w}_{T-1})$ is given to the decision maker. At time t , the decision maker observes x_t, w_{t-1} and picks a decision u_t . Then, the environment picks w_t , and the system transitions to the next step according to (1). We emphasize that, at time t , the decision maker has no access to $(w_t, \dots, w_T)$ and their values may be different from the predictions $(\hat{w}_t, \dots, \hat{w}_T)$. Also, note that w_t can be adversarially chosen at each time t , adaptively.

The assumption that a sequence of predictions available is $(\hat{w}_t, \dots, \hat{w}_T)$ is not as strong as it may first appear, nor as strong as other similar assumptions made in literature, e.g., [19, 29], because we allow for prediction error. If there are no predictions or only a subset of predictions are available,

we can simply set the unknown predictions to be zero and this does not affect our theoretical results and algorithms.

In our model, there are two types of uncertainty. The first is caused by the *perturbations* because the future perturbations $(w_t, \dots, w_{T-1})$ are unknown to the controller at time t . The second is the *prediction error* due to the mismatch $e_t := \widehat{w}_t - w_t$ between the perturbation w_t and the prediction $\widehat{w}_t$ at each time. Formally, we define the prediction error as

$$\varepsilon(F, P, e_0, \dots, e_{T-1}) := \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P e_\tau \right\|^2. \quad (2)$$

Notice that the prediction error is not defined as a form of classical mean squared error for our problem. The reason is because the mismatch e_t at each time has different impact on the system. Writing the prediction error as in (2) simplifies our analysis. In fact, if we define u_t and $\widehat{u}_t$ as two actions given by an optimal linear controller (formally defined in Section 2.3.2) as if the true perturbations are $w_0, \dots, w_{T-1}$ and $\widehat{w}_0, \dots, \widehat{w}_{T-1}$, respectively, then it can be verified that $\varepsilon = \sum_{t=0}^{T-1} \|u_t - \widehat{u}_t\|^2$, which is the accumulated action mismatch for an optimal linear controllers provided with different estimates of perturbations. In Section 5, using experiments, we show that the competitive ratios (with a fixed “trust parameter”, defined in 3.2) grow linearly in the prediction error ε defined in (2).

Finally, we assume that the perturbations $(w_0, \dots, w_{T-1})$ and predictions $(\widehat{w}_0, \dots, \widehat{w}_{T-1})$ are uniformly bounded, i.e., there exist $\bar{w} > 0$ and $\widehat{w} > 0$ such that $\|w_t\| \leq \bar{w}$ and $\|\widehat{w}_t\| \leq \widehat{w}$ for all $0 \leq t \leq T-1$. In summary, Figure 1 demonstrates the system model considered in this paper.

Fig. 1. System Model.

2.2 Defining Consistency and Robustness

As discussed in the introduction, while predictions can be helpful, inaccurate predictions can lead to unbounded competitive ratio. Our goal is to utilize predictions to achieve good performance (consistency) while still providing adversarial worst-case guarantees (robustness). In this subsection, we formally define the notions of consistency and robustness we study. These notions have received increasing attention recently in the area of online algorithms with untrusted advice, e.g., [4–7, 22, 27].

We use the competitive ratio to measure the performance of an online control policy and quantify its robustness and consistency. Specifically, let OPT be the offline optimal cost when all the disturbances $(w_t)_{t=0}^T$ are known, and ALG be the cost achieved by an online algorithm. Throughout this paper we assume $\text{OPT} > 0$. We define the competitive ratio for a given bound on the prediction error ε , as follows.

Definition 2.1. The **competitive ratio** for a given prediction error ε , $\text{CR}(\varepsilon)$, is defined as the smallest constant $C \geq 1$ such that $\text{ALG} \leq C \cdot \text{OPT}$ for fixed A, B, Q, R and any adversarially and adaptively chosen perturbations $(w_0, \dots, w_{T-1})$ and predictions $(\widehat{w}_0, \dots, \widehat{w}_{T-1})$.

Building on the definition of competitive ratio, we define robustness and consistency as follows.

Definition 2.2. *An online algorithm is said to be γ -robust if, for any prediction error $\varepsilon > 0$, the competitive ratio satisfies $CR(\varepsilon) \leq \gamma$, and an algorithm is said to be β -consistent if the competitive ratio satisfies $CR(0) \leq \beta$.*

2.3 Background: Existing Algorithms

Before proceeding to our algorithm and its analysis, we first introduce two extreme algorithm choices that have been studied previously: a myopic policy that we refer to as 1-confident control, which places full trust in the predictions, and a pure online strategy that we refer to as 0-confident control, which places no trust in the predictions. These represent algorithms that can achieve consistency and robustness *individually*, but cannot achieve consistency and robustness *simultaneously*. The key challenge of this paper is to understand how to integrate ideas such as what follows into an algorithm that achieves consistency and robustness simultaneously.

2.3.1 A Consistent Algorithm: 1-Confident Control. A simple way to achieve consistency is to put full faith in the untrusted predictions. In particular, if the algorithm trusts the untrusted predictions and follows them, the performance will always be optimal if the predictions are accurate. We refer to this as the 1-confident policy, which is defined by a finite-time optimal control problem that trusts that $(\hat{w}_0, \dots, \hat{w}_{T-1})$ are the true disturbances. Formally, at time step t , the actions $(u_t, \dots, u_T)$ are computed via

$$\arg \min_{(u_t, \dots, u_{T-1})} \left(\sum_{\tau=t}^{T-1} (x_\tau^\top Q x_\tau + u_\tau^\top R u_\tau) + x_T^\top P x_T \right) \quad \text{s.t. (1) for all } \tau = t, \dots, T-1. \quad (3)$$

With the obtained solution $(u_t, \dots, u_{T-1})$, the control action u_t at time t is fixed to be u_t and the other actions $(u_{t+1}, \dots, u_{T-1})$ are discarded.

We highlight the following result (Theorem 3.2 in [29]) that provides an explicit expression of the algorithm in (3), which can be viewed as a form of Model Predictive Control (MPC).

THEOREM 2.1 (THEOREM 3.2 IN [29]). *With predictions $(\hat{w}_0, \dots, \hat{w}_{T-1})$ fixed, the solution u_t of the algorithm in (3) can be expressed as*

$$u_t = -(R + B^\top P B)^{-1} B^\top \left(P A x_t + \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \hat{w}_\tau \right) \quad (4)$$

where $F := A - B(R + B^\top P B)^{-1} B^\top P A = A - B K$.

It is clear that this controller (3) (or equivalently (4)) achieves 1-consistency because, when the prediction errors are 0, the control action from (3) (and the state trajectory) will be exactly the same as the offline optimal. However, this approach is not robust, and one can show that prediction errors can lead to unbounded competitive ratios. In the next subsection, we introduce a robust (but non consistent) controller.

2.3.2 A Robust Algorithm: 0-Confident Control. On the other extreme, a natural way to be robust is to ignore the untrusted predictions entirely, i.e., place no confidence in the predictions. The 0-confident policy does exactly this. It places no trust in the predictions and synthesizes the controller by assuming $w_t = 0$. Formally, the policy is given by

$$u_t = -K x_t. \quad (5)$$

This recovers the optimal pure online policy in classical linear control theory [3]. As shown by [29], this controller has a constant competitive ratio and therefore is $O(1)$ -robust. However, this

Algorithm 1: Threshold-based Control

```

Initialize  $\delta = 0$ 
for  $t = 0, \dots, T - 1$  do
    if  $\delta < \sigma$  then
         $u_t = -(R + B^\top PB)^{-1} B^\top \left( P A x_t + \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right)$ 
    else
        Compute  $u_t$  with the best myopic online algorithm  $\mathcal{A}_{\text{Online}}$  without predictions
    end
    Update  $x_{t+1} = A x_t + B u_t + w_t$  and  $\delta \leftarrow \delta + \|\widehat{w}_t - w_t\|$ 
end

```

approach is not consistent as it does not utilize the predictions at all. In the next section, we discuss our proposed approach, which achieves both consistency and robustness.

3 CONSISTENT AND ROBUST CONTROL

The goal of this paper is to develop a controller that performs near-optimally when predictions are accurate (consistency) and meanwhile is robust when the prediction error is large. As discussed in the previous section, a myopic, 1-confident controller that puts full trust into the predictions is consistent, but not robust. On the other hand, any purely online 0-confident policy that ignores predictions is robust but not consistent.

The algorithms we present the trade off between these extremes by including a “confidence/trust level” for the predictions. The algorithm design challenge is to determine the right way to balance these extremes. In the first (warmup) algorithm, the policy starts out confident in the predictions, but when a threshold of error is observed, the policy loses confidence and begins to ignore predictions. This simple threshold-based policy highlights that it is possible for a policy to be both robust and consistent. However, the result also highlights the weakness of the standard notions of robustness and consistency since the policy cannot make use of intermediate quality predictions and only performs well in the extreme cases when predictions are either perfect or poor.

Thus, we move to considering a different approach, which we term *λ -confident control*. This algorithm selects a confidence level λ that serves as a weight for a linear combination between purely myopic 1-confident control and purely online 0-confident control. Our main result shows that this policy provides a smooth trade-off between robustness and consistency and, further, in Section 4, we show that the confidence level λ can be learned online adaptively so as to achieve consistency and robustness without exogenously specifying a trust level.

3.1 Warmup: Threshold-based control

We begin by presenting a simple threshold-based algorithm that can be both robust and consistent, though it does not perform well for predictions of intermediate quality. This distinction highlights that looking beyond the classical narrow definitions of robustness and consistency is important when evaluating algorithms.

The threshold-based algorithm is described in Algorithm 1. It works by trusting predictions (using 1-confident control update (4)) until a certain error threshold $\sigma > 0$ is crossed and then ignoring predictions (using an online algorithm $\mathcal{A}_{\text{Online}}$ that attains a (minimal) competitive ratio $C_{\min}$ ¹ for all online algorithms that do not use predictions). The following result shows that, with a

¹Note that $C_{\min}$ is guaranteed to exist, as setting $\lambda = 0$ in Theorem 3.2 gives a constant $1 + \|H\|/\lambda_{\min}(G)$ competitive ratio bound for the 0-confident control update (5), therefore $1 \leq C_{\min} \leq 1 + \|H\|/\lambda_{\min}(G)$.

Algorithm 2: λ -confident Control

```

for  $t = 0, \dots, T - 1$  do
    Take  $u_t = -(R + B^\top PB)^{-1} B^\top \left( PAx_t + \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P\widehat{w}_\tau \right)$ 
    Update  $x_{t+1} = Ax_t + Bu_t + w_t$ 
end

```

small enough threshold, this algorithm is both robust and consistent because, if predictions are perfect it trusts them entirely, but if there is an error, it immediately begins to ignore predictions and matches the 0-confident controller performance, which is optimal. A proof can be found in Appendix D.

THEOREM 3.1. *There exists a threshold parameter $\sigma > 0$ such that Algorithm 1 is 1-consistent and $(C_{\min} + o(1))$ -robust, where $C_{\min}$ is the minimal competitive ratio of any pure online algorithm.*

The term $o(1)$ in Theorem 3.1 converges to 0 as $T \rightarrow \infty$. While Algorithm 1 is optimally robust and consistent, it is unsatisfying because it does not improve over the online algorithm unless predictions are perfect since in the proof, we set the threshold parameter $\sigma > 0$ arbitrarily small to make the algorithm robust and 1-consistent and the definition of consistency and robustness only captures the behavior of the competitive ratio $CR(\varepsilon)$ for either $\varepsilon = 0$ or ε is large. As a result, in the remainder of the paper we look beyond the extreme cases and prove results that apply for arbitrary prediction error quality. In particular, we prove competitive ratio bounds that hold for arbitrary ε , of which consistency and robustness are then special cases.

3.2 λ -confident control

We now present our main results, which focus on a policy that, like Algorithm 1, looks to find a balance between the two extreme cases of 1-confident and 0-confident control. However, instead of using a threshold to decide when to swap between them, the λ -confident controller considers a linear combination of the two.

Specifically, the policy presented in Algorithm 2 works as follows. Given a *trust parameter* $0 \leq \lambda \leq 1$, it implements a linear combination of (4) and (5). Intuitively, the selection of λ allows a trade-off between consistency and robustness based on the extent to which the predictions are trusted. Our main result shows a competitive ratio bound that is consistent with this intuition. A proof is given in Appendix B.

THEOREM 3.2. *Under our model assumptions, with a fixed trust parameter $\lambda > 0$, the λ -confident control in Algorithm 2 has a worst-case competitive ratio of at most*

$$CR(\varepsilon) \leq 1 + \|H\| \min \left\{ \left(\frac{\lambda^2}{\text{OPT}} \varepsilon + \frac{(1-\lambda)^2}{C} \right), \left(\frac{1}{C} + \frac{\lambda^2}{\text{OPT}} \overline{W} \right) \right\} \quad (6)$$

where $H := B(R + B^\top PB)^{-1}B^\top$, OPT denotes the optimal cost, $C > 0$ is a constant that depends on A, B, Q, R and

$$\begin{aligned} \varepsilon(F, P, e_0, \dots, e_{T-1}) &:= \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_\tau - \widehat{w}_\tau) \right\|^2, \\ \overline{W}(F, P, \widehat{w}_0, \dots, \widehat{w}_{T-1}) &:= \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2. \end{aligned} \quad (7)$$

From this result we see that λ -confident control is guaranteed to be $(1 + \|H\| \frac{(1-\lambda)^2}{C})$ -consistent and $(1 + \|H\|(\frac{1}{C} + \frac{\lambda^2}{\text{OPT}} \overline{W}))$ -robust. This highlights a trade-off between consistency and robustness such that if a large λ is used (i.e., predictions are trusted), then consistency decreases to 1, while the robustness increases unboundedly. In contrast, when a small λ is used (i.e., predictions are distrusted), the robustness of the policy converges to the optimal value, but the consistency does not improve on the robustness value. Due to the time-coupling structure in the control system, the mismatches $e_t = \widehat{w}_t - w_t$ at different times contribute unequally to the system. As a result, the prediction error ε in (2) and (7) is defined as a weighed quadratic sum of $(e_0, \dots, e_{T-1})$. Moreover, the term OPT in (6) is common in the robustness and consistency analysis of online algorithms, such as [5, 7, 20, 22].

4 SELF-TUNING λ -CONFIDENT CONTROL

While the λ -confident control finds a balance between consistency and robustness, selecting the optimal λ parameter requires exogenous knowledge of the quality of the predictions ε , which is often not possible. For example, black-box AI tools typically do not allow uncertainty quantification. In this section, we develop a self-tuning λ -confident control approach that learns to tune λ in an online manner. We provide an upper bound on the regret of the self-tuning λ -confident control, compared with using the best possible λ in hindsight, and a competitive ratio for the complete self-tuning algorithm. These results provide the first worst-case guarantees for the integration of black-box AI tools into linear quadratic control.

Our policy is described in Algorithm 3 and is a “follow the leader” approach [14, 15]. At each time $t = 0, \dots, T-1$, it selects a λ_t in order to minimize the gap between ALG and OPT in the previous t rounds and chooses an action using the trust parameter λ_t . Then the state x_t is updated to x_{t+1} using the linear system dynamic in (1) and this process repeats. Note that the denominator of λ_t is zero if and only if $\eta(\widehat{w}; s, t-1) = 0$ for all s . To make λ_t well-defined, we set $\lambda = 1$ for this case.

The key to the algorithm is the update rule for λ_t . Given previously observed perturbations and predictions, the goal of the algorithm is to find a greedy λ_t that minimizes the gap between the algorithmic and optimal costs so that $\lambda_t := \min_{\lambda} \sum_{s=0}^{t-1} \psi_s^\top H \psi_s$ where $\psi_s := \sum_{\tau=s}^{t-1} (F^\top)^{\tau-s} P (w_\tau - \lambda \widehat{w}_\tau)$. This can be equivalently written as

$$\lambda_t = \arg \min_{\lambda} \sum_{s=0}^{t-1} \left[\left(\sum_{\tau=s}^{t-1} (F^\top)^{\tau-s} P (w_\tau - \lambda \widehat{w}_\tau) \right)^\top H \left(\sum_{\tau=s}^{t-1} (F^\top)^{\tau-s} P (w_\tau - \lambda \widehat{w}_\tau) \right) \right], \quad (8)$$

which is a quadratic function of λ . Rearranging the terms in (8) yields the choice of λ_t in the self-tuning control scheme.

Algorithm 3 is efficient since, in each time step, updating the η values only requires adding one more term. This means that the total computational complexity of λ_t is $O(T^2 n^\alpha)$, where $\alpha < 2.373$, which is polynomial in both the time horizon length T and state dimension n . According to the

Algorithm 3: Self-Tuning λ -Confident Control

```

for  $t = 0, \dots, T - 1$  do
  if  $t = 0$  or  $t = 1$  then
    | Initialize and choose  $\lambda_0$ 
  end
  else
    | Compute a trust parameter  $\lambda_t$ 
    
$$\lambda_t = \frac{\sum_{s=0}^{t-1} (\eta(\mathbf{w}; s, t-1))^\top H(\eta(\widehat{\mathbf{w}}; s, t-1))}{\sum_{s=0}^{t-1} (\eta(\widehat{\mathbf{w}}; s, t-1))^\top H(\eta(\widehat{\mathbf{w}}; s, t-1))}$$

    
$$\text{where } \eta(\mathbf{w}; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P \mathbf{w}_\tau$$

  end
  | Generate an action  $u_t$  using  $\lambda_t$ -confident control (Algorithm 2)
  | Update  $x_{t+1} = Ax_t + Bu_t + w_t$ 
end

```

expression of λ_t in Algorithm 3, at each time t , the terms $\eta(\mathbf{w}; s, t-2)$ and $\eta(\widehat{\mathbf{w}}; s, t-2)$ can be pre-computed for all $s = 0, \dots, t-1$. Therefore, the recursive formula $\eta(\mathbf{w}; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P \mathbf{w}_\tau = \eta(\mathbf{w}; s, t-1) + (F^\top)^{t-s} P \mathbf{w}_t$ implies the update rule of the terms $\{\eta(\mathbf{w}; s, t-1) : s = 0, \dots, t-1\}$ in the expression of λ_t . This gives that, at each time t , it takes no more than $O(Tn^\alpha)$ steps to compute λ_t where $\alpha < 2.373$ and $O(n^\alpha)$ is the computational complexity of matrix multiplication.

4.1 Convergence

We now move to the analysis of Algorithm 3. First, we study the convergence of λ_t , which depends on the variation of the predictions $\widehat{\mathbf{w}} := (\widehat{\mathbf{w}}_0, \dots, \widehat{\mathbf{w}}_{T-1})$ and the true perturbations $\mathbf{w} := (\mathbf{w}_0, \dots, \mathbf{w}_{T-1})$, where we use a boldface letter to represent a sequence of vectors. Specifically, our results are in terms of the variation of the predictions and perturbations, which we define as follows. The *self-variation* $\mu_{\text{VAR}}(\mathbf{y})$ of a sequence $\mathbf{y} := (y_0, \dots, y_{T-1})$ is defined as

$$\mu_{\text{VAR}}(\mathbf{y}) := \sum_{s=1}^{T-1} \max_{\tau=0, \dots, s-1} \|y_\tau - y_{\tau+T-s}\|.$$

The goal of the self-tuning algorithm is to converge to the optimal trust parameter λ^* for the problem instance. To specify this formally, let $\text{ALG}(\lambda_0, \dots, \lambda_{T-1})$ be the algorithmic cost with adaptively chosen trust parameters $\lambda_0, \dots, \lambda_{T-1}$ and denote by $\text{ALG}(\lambda)$ the cost with a fixed trust parameter λ . Then, λ^* is defined as $\lambda^* := \min_{\lambda \in \mathbb{R}} \text{ALG}(\lambda)$. Further, let $W(t) := \sum_{s=0}^t \eta(\widehat{\mathbf{w}}; s, t)^\top H \eta(\widehat{\mathbf{w}}; s, t)$.

We can now state a bound on the convergence rate of λ_t to λ^* under Algorithm 3. The bound highlights that if the variation of the system perturbations and predictions is small, then the trust parameter λ_t converges quickly to λ^* . A proof can be found in Appendix C.1.

LEMMA 1. Assume $W(T) = \Omega(T)$ and $\lambda_t \in [0, 1]$ for all $t = 0, \dots, T-1$. Under our model assumptions, the adaptively chosen trust parameters $(\lambda_0, \dots, \lambda_T)$ by self-tuning control satisfy that for any $1 < t \leq T$,

$$|\lambda_t - \lambda^*| = O((\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})) / t).$$

4.2 Regret and Competitiveness

Building on the convergence analysis, we now prove bounds on the regret and competitive ratio of Algorithm 3. These are the main results of the paper and represent the performance of an algorithm that adaptively determines the optimal trade-off between robustness and consistency.

Regret. We first study the regret as compared with the best, fixed trust parameter in hindsight, i.e., λ^* , whose corresponding worst-case competitive ratio satisfies the upper bound given in Theorem 3.2.

Denote by $\text{Regret} := \text{ALG}(\lambda_0, \dots, \lambda_{T-1}) - \text{ALG}(\lambda^*)$ the *regret* we consider where $(\lambda_0, \dots, \lambda_{T-1})$ are the trust parameters selected by the self-tuning control scheme. Our main result is the following variation-based regret bound, which is proven in Appendix C.2. Define the denominator of λ_t in Algorithm 2 by $W(t) := \sum_{s=0}^t \eta(\hat{\mathbf{w}}; s, t)^\top H \eta(\hat{\mathbf{w}}; s, t)$.

LEMMA 2. Assume $W(t) = \Omega(T)$ and $\lambda_t \in [0, 1]$ for all $t = 0, \dots, T-1$. Under our model assumptions, for any $\mathbf{w}$ and $\hat{\mathbf{w}}$, the regret of Algorithm 3 is bounded by

$$\text{Regret} = O\left((\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\hat{\mathbf{w}}))^2\right).$$

Note that the baseline we evaluate against in $\text{Regret} = \text{ALG}(\lambda_0, \dots, \lambda_{T-1}) - \text{ALG}(\lambda^*)$ is stronger than baselines in previous *static regret* analysis for LQR, such as [2, 10] where online controllers are compared with a linear control policy $u_t = -Kx_t$ with a strongly stable K . The baseline policy considered in our regret analysis is the λ -confident scheme (Algorithm 2) with

$$\begin{aligned} u_t &= -(R + B^\top PB)^{-1} B^\top \left(P A x_t + \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \hat{\mathbf{w}}_\tau \right) \\ &= -K x_t - \lambda (R + B^\top PB)^{-1} B^\top \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \hat{\mathbf{w}}_\tau, \end{aligned}$$

which contains the class of strongly stable linear controllers as a special case. Moreover, the regret bound in Lemma 2 holds for any predictions $\hat{\mathbf{w}}_0, \dots, \hat{\mathbf{w}}_{T-1}$. Taking $\hat{\mathbf{w}}_t = \mathbf{w}_t$ for all $t = 0, \dots, T-1$, our regret directly compares $\text{ALG}(\lambda_0, \dots, \lambda_{T-1})$ with the optimal cost OPT, and therefore, our regret also involves the *dynamic regret* considered in [19, 29, 30] for LQR as a special case. The regret bound in Lemma 2 depends on the variation of perturbations and predictions. Note that such a term commonly exists in regret analysis based on the “follow the leader” approach [14, 15]. For example, the regret analysis of the follow the optimal steady state (FOSS) method in [19] contains a similar “path length” term that captures the state variation. There is a variation budget of the predictions or prediction errors in theorem 1 of [30]. In many robotics applications (e.g., the trajectory tracking and EV charging experiments in this paper shown in Section 5), each $\mathbf{w}_t$ is from some desired smooth trajectory to track.

To interpret this lemma, suppose the sequences of perturbations and predictions satisfy:

$$\begin{aligned} \|\hat{\mathbf{w}}_\tau - \hat{\mathbf{w}}_{\tau+T-s}\| &\leq \rho(s), \\ \|\mathbf{w}_\tau - \mathbf{w}_{\tau+T-s}\| &\leq \rho(s), \text{ for any } s \geq 0, 0 \leq \tau \leq s. \end{aligned}$$

These bounds correspond to an assumption of smooth variation in the disturbances and the predictions. Note that it is natural for the disturbances to vary smoothly in applications such as tracking problems where the disturbances correspond to the trajectory and in such situations one would expect the predictions to also vary smoothly. For example, machine learning algorithms are often regularized to provide smooth predictions.

Given these smoothness bounds, we have that

$$\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}) \leq \sum_{s=0}^{T-1} 2\rho(s).$$

Note that, as long as $\sum_{s=0}^{T-1} \rho(s) = o(T)$, the regret bound is sub-linear in T . To understand how this bound may look in particular applications, suppose we have $\rho(s) = O(1/s)$. In this case, regret is poly-logarithmic, i.e., $\text{Regret} = O((\log T)^2)$. If $\rho(s)$ is exponential the regret is even smaller, i.e., if $\rho(s) = O(r^s)$ for some $0 < r < 1$ then $\text{Regret} = O(1)$.

Competitive Ratio. We are now ready to present our main result, which provides an upper bound on the competitive ratio of self-tuning control (Algorithm 3). Recall that, in Lemma 2, we bound the regret $\text{Regret} := \text{ALG}(\lambda_0, \dots, \lambda_{T-1}) - \text{ALG}(\lambda^*)$ and, in Theorem 3.2, a competitive ratio bound is provided for the λ -confident control scheme, including $\text{ALG}(\lambda^*)/\text{OPT}$. Therefore, combining Lemma 2 and Theorem 3.2 leads to a novel competitive ratio bound for the self-tuning scheme (Algorithm 3). Note that compared with Theorem 3.2, which also provides a competitive ratio bound for λ -confident control, Theorem 4.1 below considers a competitive ratio bound for the self-tuning scheme in Algorithm 3 where, at each time t , a trust parameter λ_t is determined by online learning and may be time-varying.

THEOREM 4.1. *Assume $W(T) = \Omega(T)$ and $\lambda_t \in [0, 1]$ for all $t = 0, \dots, T-1$. Under our model assumptions, the competitive ratio of Algorithm 3 is bounded by*

$$\text{CR}(\varepsilon) \leq 1 + \|H\| \frac{\varepsilon}{\text{OPT} + C\varepsilon} + O\left(\frac{(\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2}{\text{OPT}}\right)$$

where H, C OPT and ε are defined in Theorem 3.2.

In contrast to the regret bound, Theorem 4.1 states an upper bound on the competitive ratio $\text{CR}(\varepsilon)$ defined in Section 2.2, which indicates that $\text{CR}(\varepsilon)$ scales as $1 + O(\varepsilon)/(\Theta(1) + \Theta(\varepsilon))$ as a function of ε . As a comparison, the λ -confident control in Algorithm 2 has a competitive ratio upper bound that is linear in the prediction error ε (Theorem 3.2). This improved dependency highlights the importance of learning the trust parameter adaptively.

Our experimental results in the next section verify the implications of Theorem 3.2 and Theorem 4.1. Specifically, the simulated competitive ratio of the self-tuning control (Algorithm 3) is a non-linear envelope of the simulated competitive ratios for λ -confident control with fixed trust parameters and as a function of prediction error ε , it matches the implied competitive ratio upper bound $1 + O(\varepsilon)/(\Theta(1) + \Theta(\varepsilon))$.

Theorem 4.1 is proven by combining Lemma 2 with Theorem 3.2, which bounds the competitive ratios for fixed trust parameters.

PROOF OF THEOREM 4.1. Denote by $\text{ALG}(\lambda_0, \dots, \lambda_{T-1})$ the algorithmic cost of the self-tuning control scheme. We have

$$\frac{\text{ALG}(\lambda_0, \dots, \lambda_{T-1})}{\text{OPT}} \leq \frac{|\text{ALG}(\lambda_0, \dots, \lambda_{T-1}) - \text{ALG}(\lambda^*)|}{\text{OPT}} + \frac{\text{ALG}(\lambda^*)}{\text{OPT}}. \quad (9)$$

Using Theorem 3.2,

$$\begin{aligned} \frac{\text{ALG}(\lambda^*)}{\text{OPT}} &\leq 1 + \|H\| \min \left\{ \min_{\lambda} \left(\frac{\lambda^2}{\text{OPT}} \varepsilon + \frac{(1-\lambda)^2}{C} \right), \min_{\lambda} \left(\frac{1}{C} + \frac{\lambda^2}{\text{OPT}} \overline{W} \right) \right\} \\ &= 1 + \|H\| \min \left\{ \frac{\varepsilon}{\text{OPT} + \varepsilon C}, \frac{1}{C} \right\} = 1 + \|H\| \frac{\varepsilon}{\text{OPT} + \varepsilon C}. \end{aligned} \quad (10)$$

Moreover, the regret bound in Lemma 2 implies

$$\frac{|\text{ALG}(\lambda_0, \dots, \lambda_{T-1}) - \text{ALG}(\lambda^*)|}{\text{OPT}} = O\left(\frac{(\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2}{\text{OPT}}\right),$$

combing which with (10), (9) gives the results. $\square$

5 CASE STUDIES

We now illustrate our main results using numerical examples and case studies to highlight the impact of the trust parameter λ in λ -confident control and demonstrate the ability of the self-tuning control algorithm to learn the appropriate trust parameter λ . We consider three applications. The first is a robot tracking example where a robot is asked to follow locations of an unknown trajectory and the desired location is only revealed the time immediately before the robot makes a decision to modify its velocity. Predictions about the trajectory are available. However, the predictions can be untrustworthy so that they may contain large errors. The second is an adaptive battery-buffered electric vehicle (EV) charging problem where a battery-buffered charging station adaptively supplies energy demands of arriving EVs while maintaining the state of charge of the batteries as close to a nominal level as possible. Our third application considers a non-linear control problem – the Cart-Pole problem. Our λ -confident and self-tuning control schemes use a linearized model while the algorithms are tested with the non-linear environment. We use the third application to demonstrate the practicality of our algorithms by showing that they do not only work for LQC problems, but also non-linear systems.

To illustrate the impact of randomness in prediction errors in our case studies, the three applications all use different forms of random error models. For each selected distribution of $\widehat{\mathbf{w}} - \mathbf{w}$, we repeat the experiments multiple times and report the worst case with the highest algorithmic cost; see Appendix E for details.

5.1 Application 1: Robot tracking

Problem description. The first example we consider is a two-dimensional robot tracking application [19, 28]. There is a robot controller following a fixed but unknown cloud-shaped trajectory (see Figures 2 and 3), which is

$$y_t := \begin{bmatrix} 2 \cos(\pi t/30) + \cos(\pi t/5) \\ 2 \sin(\pi t/30) + \sin(\pi t/5) \end{bmatrix}, \quad t = 0, \dots, T-1.$$

The robot controller's location at time $t+1$, denoted by $p_{t+1} \in \mathbb{R}^2$, depends on its previous location and its velocity $v_t \in \mathbb{R}^2$ such that $p_{t+1} = p_t + 0.2v_t$ and at each time $t+1$, the controller is able to apply an adjustment u_t to modify its velocity such that $v_{t+1} = v_t + 0.2u_t$. Together, letting $x_t := p_t - y_t$, this system can be recast in the canonical form in (1) as

$$\begin{bmatrix} x_{t+1} \\ v_{t+1} \end{bmatrix} = A \begin{bmatrix} x_t \\ v_t \end{bmatrix} + Bu_t + w_t, \quad \text{with}$$

$$A := \begin{bmatrix} 1 & 0 & 0.2 & 0 \\ 0 & 1 & 0 & 0.2 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B := \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0.2 & 0 \\ 0 & 0.2 \end{bmatrix}, \quad \text{and } w_t := Ay_t - y_{t+1}.$$

Fig. 2. Tracking trajectories and trust parameters $(\lambda_0, \dots, \lambda_{T-1})$ of the self-tuning control scheme (Left: low binomial prediction error with $c = 10^{-2}$; middle: medium prediction error with $c = 10^{-1}$; right: high prediction error with $c = 1$ where c is a tuning parameter defined in Appendix E). The x-axis and y-axis in the top 6 figures are locations of the robot. The y-axis in the bottom 3 figures denotes the value of the trust parameter.

To track the trajectory, the controller sets

$$Q := \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \text{ and } R := \begin{bmatrix} 10^{-2} & 0 \\ 0 & 10^{-2} \end{bmatrix}.$$

Experimental results. In our first experiment, we demonstrate the convergence of the self-tuning scheme in Algorithm 3. To mimic the worst-case error, a random prediction error $e_t = \hat{w}_t - w_t$ at each time t is used. We then sample prediction error and implement our algorithm with several error instances and choose the one the worst competitive ratio. The details of settings can be found in Appendix E. To better simulate the task of tracking a trajectory and make it easier to observe the tracking accuracy, we ignore the cost of increasing velocity by setting R as a zero matrix for Figure 2 and Figure 3.

In Figure 2, we observe that the tracking trajectory generated by the self-tuning scheme converges to the unknown trajectory $(y_1, \dots, y_T)$, regardless of the level of prediction error. We plot the tracking trajectories every 60 time steps with a scaling parameter (defined in Appendix E) $c = 10^{-2}$ (left), $c = 10^{-1}$ (mid) and $c = 1$ (right) respectively. In all cases, we observe convergence of the

Fig. 3. Tracking trajectories and trust parameters ($\lambda_0, \dots, \lambda_{T-1}$) of the self-tuning control scheme (Left: low Gaussian prediction error with variance $\sigma^2 = 10^{-2}$; middle: medium Gaussian prediction error with variance $\sigma^2 = 10^{-1}$; right: high Gaussian prediction error with variance $\sigma^2 = 1$). The x-axis and y-axis in the top 6 figures are locations of the robot. The y-axis in the bottom 3 figures denotes the value of the trust parameter.

Fig. 4. Impact of trust parameters and performance of self-tuning control for robot tracking.

trust parameters. Moreover, for a wide range of prediction error levels, without knowing the prediction error level in advance, the scheme is able to automatically switch its mode and become both consistent and robust by choosing an appropriate trust parameter λ_t to accurately track the unknown trajectory. In Figure 3, we observe similar behavior when the prediction error is generated from Gaussian distributions.

Next, we demonstrate the performance of self-tuning control and the impact of trust parameters. In Figure 4, we depict the competitive ratios of the λ -confident control algorithm described in Section 3.2 with varying trust parameters, together with the competitive ratios of the self-tuning control scheme described in Algorithm 3. The label of the x -axis is the prediction error ε (normalized by 10^3), defined in (7). We divide our results into two parts. The left sub-figure in Figure 4 considers a low-error regime where we observe that the competitive ratio of the self-tuning policy performs closely as the lower envelope formed by picking multiple trust parameters optimally offline. The right sub-figure in Figure 4 shows the performance of self-tuning for the case when the prediction error is high. For the high-error regime, the competitive ratio of the self-tuning control policy is close to those with the best fixed trust parameter.

5.2 Application 2: Adaptive battery-buffered EV charging

Fig. 5. Adaptive battery-buffered EV charging modelled as a linear quadratic control problem.

Problem description. We consider an adaptive battery-buffered Electric Vehicle (EV) charging problem. There is a charging station with N chargers, with each charger connected to a battery energy storage system. Let x_t be a vector in $\mathbb{R}_+^N$, whose entries represent the State of Charge (SoC) of the batteries at time t . The charging controller decides a charging schedule u_t in $\mathbb{R}_+^N$ where each entry in u_t is the energy to be charged to the i -th battery from external power supply at time t . The canonical form of the system can be represented by

$$x_{t+1} = Ax_t + Bu_t - w_t,$$

where A is an $N \times N$ matrix denotes the *degradation* of battery charging levels and B is an $N \times N$ diagonal matrix whose diagonal entry $0 \leq B_i \leq 1$ represents the charging efficiency coefficient. In our experiments, without loss of generality, we assume A and B are identity matrices. The perturbation w_t is defined as a length- N vector, whose entry $w_t(i) = E$ when at time t an EV arrives at charger i and demands energy $E > 0$; otherwise $w_t(i) = 0$. Therefore the perturbations $(w_0, \dots, w_{T-1})$ depend on the arrival of EVs and their energy demands. The charging controller can only make a charging decision u_t at time t before knowing w_t (as well as $w_{t+1}, \dots, w_{T-1}$) and the EVs that arrive at time t (as well as future EV arrivals). The goal of the adaptive battery-buffered EV charging problem is to maintain the battery SoC as close to a nominal value $\bar{x}$ as possible. Therefore, the charging controller would like to minimize $\sum_{t=0}^{T-1} (x_t - \bar{x})^\top Q (x_t - \bar{x}) + u_t^\top R u_t$, equivalently, $\sum_{t=0}^{T-1} x_t^\top Q x_t + u_t^\top R u_t$ where Q can be some positive-definite matrix and R encodes the costs of external power supply. In our experiments, we set Q as an identity matrix and $R = 0.1 \times Q$.

Experimental results. We show the performance of self-tuning control and the impact of trust parameters for adaptive EV charging in Figure 7 and 8. In Figure 7 we consider a synthetic case

Fig. 6. An example of the daily charging demands in ACN-Data [17] on Nov 1st, 2018.

Fig. 7. Impact of trust parameters and performance of self-tuning control for adaptive battery-buffered EV charging with synthetic EV charging data.

Fig. 8. Impact of trust parameters and performance of self-tuning control for adaptive battery-buffered EV charging with realistic daily EV charging data [17].

when EVs with 5 kWh battery capacity arrive at a constant rate 0.2, e.g., 1 EV arrives every 5 time slots. The results are divided into two parts. In Figure 8, we use daily data (ACN-Data) that contain EVs' energy demands, arrival times and departure times collected from a real-world adaptive EV

charging network [17]. We select a daily charging record on Nov 1st, 2018, depicted in Figure 6. The left sub-figure considers a magnified low-error regime and the right sub-figure shows the performance of self-tuning for the case when the prediction error is high. For both regimes, the competitive ratios of the self-tuning control policy perform nearly as well as the lower envelope formed by picking multiple trust parameters optimally offline. We see in both Figure 4 and Figure 7 that with fixed trust parameters the competitive ratio is linear in ε , matching what Theorem 3.2 indicates (in the sense of order in ε). Moreover, for the self-tuning scheme, in both Figure 4 and Figure 7, we observe a competitive ratio $1 + O(\varepsilon)/(O(1) + O(\varepsilon))$, which matches the competitive ratio bound given in Theorem 4.1 in order sense (in ε).

5.3 Application 3: Cart-Pole

Problem description. The third set of experiments we consider is the classic Cart-Pole problem illustrated in Figure 9. The goal of a controller is to stabilize the pole in the upright position.

This is a widely studied nonlinear system. Neglecting friction, the dynamical equations of the Cart-Pole problem are

Fig. 9. The Cart-Pole model in Application 3.

$$\ddot{\theta} = \frac{g \sin \theta + \cos \theta \left(\frac{-u - ml\dot{\theta}^2 \sin \theta}{m+M} \right)}{l \left(\frac{4}{3} - \frac{m \cos^2 \theta}{m+M} \right)}, \quad (11)$$

$$\ddot{y} = \frac{u + ml \left(\dot{\theta}^2 \sin \theta - \ddot{\theta} \cos \theta \right)}{m + M} \quad (12)$$

where u is the input force; θ is the angle between the pole and the vertical line; y is the location of the pole; g is the gravitational acceleration; l is the pole length; m is the pole mass; and M is the cart mass. Taking $\sin \theta \approx \theta$ and $\cos \theta \approx 1$ and ignoring higher order terms, the dynamics of the Cart-Pole problem can be linearized as

$$\frac{d}{dt} \begin{bmatrix} \dot{y} \\ \ddot{y} \\ \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & -\frac{m l g}{\eta(m+M)} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & \frac{g}{\eta} & 0 \end{bmatrix} \frac{d}{dt} \begin{bmatrix} y \\ \dot{y} \\ \theta \\ \dot{\theta} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{(m+M)\eta + ml}{(m+M)^2\eta} \\ 0 \\ -\frac{1}{(m+M)\eta} \end{bmatrix} u_t + w_t,$$

where in the above $\eta := l \left(\frac{4}{3} - \frac{m}{m+M} \right)$ and, in our experiments, we set the cart mass $M = 10.0kg$, pole mass $m = 1.0kg$, pole length $l = 10.0m$ and gravitational acceleration $g = 9.8m/s^2$. We set $Q = I$ and $R = 10^{-3}$ and each w_t is a fixed external force defined as

$$60 \times \left[0, \frac{(m+M)\eta + ml}{(m+M)^2\eta}, 0, -\frac{1}{(m+M)\eta} \right]^\top.$$

Experimental results. We show the performance of the self-tuning control (Algorithm 3) and the impact of trust parameters for the Cart-Pole problem in Figure 10, together with the λ -confident control scheme in Algorithm 2 for several fixed trust parameters λ . The algorithms are tested using the true nonlinear dynamical equations in (11)-(12).

In Figure 10, we change the variance σ^2 of the prediction noise $e_t = \hat{w}_t - w_t$ at each time t and plot the average episodic rewards in the OpenAI Gym environment [9]. Different from the worst-case settings in the previous two applications, we run episodes multiples times and show plot the mean rewards. The height of the shadow area in Figure 10 represents the standard deviation of the rewards. The detailed hyper-parameters are given in Section E. Our results show that, despite the

Fig. 10. Impact of trust parameters and performance of self-tuning control for the Cart-Pole problem.

fact that the problem is *nonlinear*, the self-tuning control algorithm using a linearized model is still able to automatically adjust the trust parameter λ_t and achieves both consistency and robustness, regardless of the prediction error. In particular, it is close to the best algorithms for small prediction error while also staying among the best when prediction error is large.

6 CONCLUDING REMARKS

In this paper we detail an approach that allows the use of black-box AI tools in a way that ensures worst-case performance bounds for linear quadratic control. Further, we demonstrate the effectiveness of our approach in multiple applications.

There are many potential future directions that build on this work. First, we consider a linear quadratic control problem in this paper, and an important extension will be to analyze the robustness and consistency of non-linear control systems. Second, our regret bound (Lemma 2) and competitive results (Theorem 4.1) are not tight when the variation of perturbations or predictions is high, therefore it is interesting to explore the idea of “follow-the-regularized-leader” [21, 23] and understand if adding an extra regularizer in the update rule of λ for self-tuning control can improve the convergence and/or the regret. Finally, characterizing a tight trade-off between robustness and consistency for linear quadratic control is of particular interest. For example, the results in [22, 27] together imply a tight robustness and consistency trade-off for the ski-rental problem. It would be interesting to explore if it is possible to do the same for linear quadratic control.

REFERENCES

- [1] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26. JMLR Workshop and Conference Proceedings, 2011.
- [2] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.
- [3] Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.
- [4] Spyros Angelopoulos, Christoph Dürr, Shendan Jin, Shahin Kamali, and Marc Renault. Online computation with untrusted advice. *arXiv preprint arXiv:1905.05655*, 2019.
- [5] Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. Online metric algorithms with untrusted predictions. In *International Conference on Machine Learning*, pages 345–355. PMLR, 2020.

- [6] Antonios Antoniadis, Themis Gouleakis, Pieter Kleer, and Pavel Kolev. Secretary and online matching problems with machine learned advice. *arXiv preprint arXiv:2006.01026*, 2020.
- [7] Etienne Bamas, Andreas Maggiori, and Ola Svensson. The primal-dual method for learning augmented algorithms. *arXiv preprint arXiv:2010.11632*, 2020.
- [8] Alberto Bemporad and Manfred Morari. Robust model predictive control: A survey. In *Robustness in identification and control*, pages 207–226. Springer, 1999.
- [9] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [10] Alon Cohen, Avinatan Hasidim, Tomer Koren, Nevena Lazic, Yishay Mansour, and Kunal Talwar. Online linear quadratic control. In *International Conference on Machine Learning*, pages 1029–1038. PMLR, 2018.
- [11] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019.
- [12] John Doyle, Keith Glover, Pramod Khargonekar, and Bruce Francis. State-space solutions to standard h_2 and h_∞ control problems. In *1988 American Control Conference*, pages 1691–1696. IEEE, 1988.
- [13] Geir E Dullerud and Fernando Paganini. *A course in robust control theory: a convex approach*, volume 36. Springer Science & Business Media, 2013.
- [14] James Hannan. 4. approximation to rayes risk in repeated play. In *Contributions to the Theory of Games (AM-39), Volume III*, pages 97–140. Princeton University Press, 2016.
- [15] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [16] Islam SM Khalil. *Robust and optimal control*.
- [17] Zachary J Lee, Tongxin Li, and Steven H Low. Acn-data: Analysis and applications of an open ev charging dataset. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pages 139–149, 2019.
- [18] Tongxin Li, Yue Chen, Bo Sun, Adam Wierman, and Steven H Low. Information aggregation for constrained online control. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 5(2):1–35, 2021.
- [19] Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: algorithms and regret analysis. In *NeurIPS*, pages 14858–14870, 2019.
- [20] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3296–3305. PMLR, 2018.
- [21] Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and l1 regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 525–533. JMLR Workshop and Conference Proceedings, 2011.
- [22] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. In *Advances in Neural Information Processing Systems*, pages 9661–9670, 2018.
- [23] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and trends in Machine Learning*, 4(2):107–194, 2011.
- [24] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. In *Advances in Neural Information Processing Systems*, volume 33, pages 20636–20647. Curran Associates, Inc., 2020.
- [25] Max Simchowitz and Dylan Foster. Naive exploration is optimal for online lqr. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020.
- [26] Jean-Jacques E Slotine, Weiping Li, et al. *Applied nonlinear control*, volume 199. Prentice hall Englewood Cliffs, NJ, 1991.
- [27] Alexander Wei and Fred Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. *arXiv preprint arXiv:2010.11443*, 2020.
- [28] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Competitive control with delayed imperfect information. *arXiv preprint arXiv:2010.11637*, 2020.
- [29] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. The power of predictions in online control. *Advances in Neural Information Processing Systems*, 33, 2020.
- [30] Runyu Zhang, Yingying Li, and Na Li. On the regret analysis of online lqr control with predictions. *arXiv preprint arXiv:2102.01309*, 2021.
- [31] Kemin Zhou, John Comstock Doyle, Keith Glover, et al. *Robust and optimal control*, volume 40.

A USEFUL LEMMAS

Before proceeding to the proofs of our main results, we present some useful lemmas. We first present a lemma below from [28] that characterizes the difference between the optimal and the algorithmic costs.

LEMMA 3 (LEMMA 10 IN [28]). *For any $\psi_t \in \mathbb{R}^n$, if at each time $t = 0, \dots, T-1$,*

$$u_t = -(R + B^\top PB)^{-1} B^\top \left(P A x_t + \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \psi_t \right),$$

then the gap between the optimal cost OPT and the algorithm cost ALG induced by selecting control actions $(u_1, \dots, u_T)$ equals to

$$\text{ALG} - \text{OPT} = \sum_{t=0}^{T-1} \psi_t^\top H \psi_t \quad (13)$$

where $H := B(R + B^\top PB)^{-1} B^\top$ and $F := A - HPA$.

The next lemma describes the form of the optimal trust parameter.

LEMMA 4. *The optimal trust parameter λ^* that minimizes $\text{ALG}(\lambda) - \text{OPT}$ is $\lambda^* = \lambda_T$.*

PROOF OF LEMMA 4. The optimal trust parameter λ^* is

$$\lambda^* := \min_{\lambda} \sum_{s=0}^{T-1} \left[\left(\sum_{\tau=s}^{t-1} (F^\top)^{\tau-s} P (w_\tau - \lambda \widehat{w}_\tau) \right)^\top H \left(\sum_{\tau=s}^{T-1} (F^\top)^{\tau-s} P (w_\tau - \lambda \widehat{w}_\tau) \right) \right], \quad (14)$$

implying that $\lambda^* = \lambda_T$. □

Next, we note that the static regret depends on the convergence of λ_t .

LEMMA 5. *The static regret satisfies*

$$\text{Regret} \leq \|H\| \sum_{t=0}^{T-1} \left\| |\lambda_t - \lambda_T| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2.$$

PROOF OF LEMMA 5. Let $\text{ALG}((\lambda_0, \dots, \lambda_{T-1}))$ and $\text{ALG}(\lambda_T)$ denote the corresponding algorithm costs for using trust parameters $(\lambda_0, \dots, \lambda_{T-1})$ and a fixed optimal trust parameter λ_T in hindsight correspondingly. It follows that

$$\text{ALG}((\lambda_0, \dots, \lambda_{T-1})) - \text{ALG}(\lambda_T) \leq \|H\| \sum_{t=0}^{T-1} \left\| |\lambda_t - \lambda_T| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2. \quad (15)$$

□

LEMMA 6. *Suppose two real sequences $(V_1, \dots, V_T)$ and $(W_1, \dots, W_T)$ with $W_t > 0$ for all $1 \leq t \leq T$, converge to V_T and $W_T > 0$ such that for any integer $1 \leq t \leq T$, $|V_t - V_T| \leq C_1/t$ and $|W_t - W_T| \leq C_2/t$ for some constants $C_1, C_2 > 0$. Then the sequence $\left(\frac{V_1}{W_1}, \dots, \frac{V_T}{W_T} \right)$ converges to $\frac{V_T}{W_T}$ such that for any $1 \leq t \leq T$,*

$$\left| \frac{V_t}{W_t} - \frac{V_T}{W_T} \right| \leq \frac{1}{t} \left(\frac{C_1 \alpha_t + C_2}{|W_T|} \right).$$

where $\alpha_t := \max\{V_t/W_t\}$.

PROOF OF LEMMA 6. Based on the assumption, for any $1 \leq t \leq T$, we have that

$$\begin{aligned} \left| \frac{V_t}{W_t} - \frac{V_T}{W_T} \right| &= \left| \frac{V_t W_T - V_T W_t}{W_t W_T} \right| = \left| \frac{V_t W_T - V_t W_t + V_t W_t - V_T W_t}{W_t W_T} \right| \\ &\leq \left| \frac{V_t (W_T - W_t)}{W_t W_T} \right| + \left| \frac{W_t (V_t - V_T)}{W_t W_T} \right| \\ &\leq \frac{1}{t} \left(\frac{C_1 |V_t|}{|W_t W_T|} + \frac{C_2}{|W_T|} \right). \end{aligned}$$

Since $W_t \neq 0$ for all $1 \leq t \leq T$ and $W_T \neq 0$, the lemma follows. $\square$

LEMMA 7. Suppose a sequence $(A_0, \dots, A_{T-1})$ satisfies that for any integer $0 \leq s \leq T-1$, $|A_s - A_T| \leq \rho(s)$. Then, for any $0 \leq s \leq T$, $\left| \frac{1}{t} \left(\sum_{s=0}^t A_s \right) - A_T \right| \leq \frac{1}{t} \sum_{s=0}^{T-1} \rho(s)$.

PROOF OF LEMMA 7. Based on the assumption,

$$\left| \frac{1}{t} \sum_{s=0}^t A_s - A_T \right| = \frac{1}{t} \left| \sum_{s=0}^t (A_s - A_T) \right| \leq \frac{1}{t} \sum_{s=0}^t |A_s - A_t| \leq \frac{1}{t} \sum_{s=0}^{T-1} \rho(s).$$

$\square$

B COMPETITIVE ANALYSIS

Throughout, for notational convenience, we write

$$W(t) := \sum_{s=0}^t \eta(\widehat{w}; s, t)^\top H \eta(\widehat{w}; s, t), \quad \text{and} \quad V(t) := \sum_{s=0}^t \eta(w; s, t)^\top H \eta(\widehat{w}; s, t)$$

where

$$\eta(w; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P w_\tau, \quad \text{and} \quad \eta(\widehat{w}; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P \widehat{w}_\tau.$$

We first prove the following theorem.

THEOREM B.1. With a fixed trust parameter $\lambda > 0$, the λ -confident control in Algorithm 2 has a worst-case competitive ratio of at most

$$\text{CR}(\varepsilon) \leq 1 + \|H\| \min \left\{ \left(\frac{\lambda^2}{\text{OPT}} \varepsilon + \frac{(1-\lambda)^2}{C} \right), \left(\frac{1}{C} + \frac{\lambda^2}{\text{OPT}} \overline{W} \right) \right\}$$

where $H := B(R + B^\top P B)^{-1} B^\top$, OPT denotes the optimal cost, $C > 0$ is a constant that depends on A, B, Q, R and

$$\begin{aligned} \varepsilon(F, P, e_0, \dots, e_{T-1}) &:= \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_\tau - \widehat{w}_\tau) \right\|^2, \\ \overline{W}(F, P, \widehat{w}_0, \dots, \widehat{w}_{T-1}) &:= \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2. \end{aligned}$$

B.1 Proof of Theorem 3.2

Denote by ALG the cost induced by taking actions $(u_0, \dots, u_{T-1})$ in Algorithm 2 and OPT the optimal total cost. Note that we assume $\text{OPT} > 0$. Lemma 3 implies that

$$\text{ALG} - \text{OPT} = \sum_{t=0}^{T-1} \left(\sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_t - \lambda \widehat{w}_\tau) \right)^\top H \left(\sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_t - \lambda \widehat{w}_\tau) \right). \quad (16)$$

Therefore, with a sequence of actions $(u_1, \dots, u_T)$ generated by the λ -confident control scheme, (16) leads to

$$\begin{aligned} \text{ALG} - \text{OPT} &\leq \|H\| \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2 \\ &= \|H\| \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_\tau + e_\tau) \right\|^2 \\ &= \|H\| \sum_{t=0}^{T-1} \left\| (1-\lambda) \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P e_\tau \right\|^2 \\ &\leq \|H\| \left((1-\lambda)^2 \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau \right\|^2 + \lambda^2 \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P e_\tau \right\|^2 \right) \end{aligned}$$

where $e_t := \widehat{w}_t - w_T$ for all $t = 0, \dots, T-1$. Moreover, denoting by x_t^* and u_t^* the offline optimal state and action at time t , the optimal cost satisfies

$$\begin{aligned} \text{OPT} &= \sum_{t=0}^{T-1} (x_t^*)^\top Q x_t^* + (u_t^*)^\top R u_t^* + (x_T^*)^\top P x_T^* \\ &\geq \sum_{t=0}^{T-1} \lambda_{\min}(Q) \|x_t^*\|^2 + \lambda_{\min}(R) \|u_t^*\|^2 + \lambda_{\min}(P) \|x_T^*\|^2 \end{aligned} \quad (17)$$

$$\begin{aligned} &\geq 2D_0 \sum_{t=0}^{T-1} (\|Ax_t^*\|^2 + \|Bu_t^*\|^2) + \frac{1}{2} \sum_{t=0}^{T-1} \lambda_{\min}(Q) \|x_t^*\|^2 + \lambda_{\min}(P) \|x_T^*\|^2 \\ &\geq D_0 \sum_{t=0}^{T-1} \|Ax_t^* + Bu_t^*\|^2 + \frac{1}{2} \sum_{t=0}^{T-1} \lambda_{\min}(Q) \|x_t^*\|^2 + \lambda_{\min}(P) \|x_T^*\|^2 \\ &= D_0 \sum_{t=0}^{T-1} \|x_{t+1}^* - w_t\|^2 + \frac{1}{2} \sum_{t=0}^{T-1} \lambda_{\min}(Q) \|x_t^*\|^2 + \lambda_{\min}(P) \|x_T^*\|^2 \\ &\geq \frac{D_0}{2} \sum_{t=0}^{T-1} \|w_t\|^2 + \left(\frac{\lambda_{\min}(Q)}{2} - D_0 \right) \sum_{t=0}^{T-1} \|x_t^*\|^2 + (\lambda_{\min}(P) - C) \|x_T^*\|^2 \end{aligned} \quad (18)$$

for some constant $0 < D_0 < \min\{\lambda_{\min}(P), \lambda_{\min}(Q)/2\}$ that depends on Q, R and K where in (17), $\lambda_{\min}(Q)$, $\lambda_{\min}(R)$ and $\lambda_{\min}(P)$ are the smallest eigenvalues of positive definite matrices Q, R and P , respectively. Then,

$$\sum_{t=0}^{T-1} \|\psi_t\|^2 = \sum_{t=0}^{T-1} \left\| \sum_{\tau=0}^{T-t-1} (F^\top)^\tau P w_{t+\tau} \right\|^2 \leq 2\|P\| \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-t-1} \|F^\top\|^\tau \|w_{t+\tau}\|^2$$

where $\psi_t := \sum_{\tau=0}^{T-t-1} (F^\top)^\tau P w_{t+\tau}$. Note that $F = A - BK$ and we define $\rho := \frac{1+\rho(F)}{2} < 1$ where $\rho(F)$ denotes the spectral radius of F . From Gelfand's formula, there exists a constant $D_1 \geq 0$ such that $\|F^t\| \leq D_1 \rho^t$ for all $t \geq 0$. Therefore,

$$\begin{aligned}
 \sum_{t=0}^{T-1} \|\psi_t\|^2 &= \sum_{t=0}^{T-1} \left\| \sum_{\tau=0}^{T-t-1} (F^\top)^\tau P w_{t+\tau} \right\|^2 \\
 &\leq D_1^2 \|P\|^2 \sum_{t=0}^{T-1} \left(\sum_{\tau=0}^{T-t-1} \rho^\tau \|w_{t+\tau}\| \right)^2 \\
 &= D_1^2 \|P\|^2 \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-t-1} \sum_{\ell=0}^{T-t-1} \rho^\tau \rho^\ell \|w_{t+\tau}\| \|w_{t+\ell}\| \\
 &\leq \frac{D_1^2}{2} \|P\|^2 \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-t-1} \sum_{\ell=0}^{T-t-1} \rho^\tau \rho^\ell (\|w_{t+\tau}\|^2 + \|w_{t+\ell}\|^2)
 \end{aligned} \tag{19}$$

Continuing from (19),

$$\begin{aligned}
 \sum_{t=0}^{T-1} \|\psi_t\|^2 &\leq \frac{D_1^2}{2} \|P\|^2 \left(\sum_{\ell=0}^{T-1} \rho^\ell \right) \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-t-1} \rho^\tau \|w_{t+\tau}\|^2 \\
 &\quad + \frac{D_1^2}{2} \|P\|^2 \left(\sum_{\tau=0}^{T-1} \rho^\tau \right) \sum_{t=0}^{T-1} \sum_{\ell=0}^{T-t-1} \rho^\ell \|w_{t+\ell}\|^2
 \end{aligned} \tag{20}$$

$$\begin{aligned}
 &\leq \frac{D_1^2}{1-\rho} \|P\|^2 \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-t-1} \rho^\tau \|w_{t+\tau}\|^2 \\
 &\leq \frac{D_1^2}{1-\rho} \|P\|^2 \sum_{t=0}^{T-1} \sum_{\tau=0}^{T-1} \rho^\tau \|w_{(t+\tau) \bmod T}\|^2 \\
 &= \frac{D_1^2}{1-\rho} \|P\|^2 \left(\sum_{\tau=0}^{T-1} \rho^\tau \right) \left(\sum_{t=0}^{T-1} \|w_t\|^2 \right) \\
 &\leq \frac{D_1^2}{(1-\rho)^2} \|P\|^2 \sum_{t=0}^{T-1} \|w_t\|^2.
 \end{aligned} \tag{21}$$

Putting (21) into (18), we obtain

$$\text{OPT} \geq \frac{D_0(1-\rho)^2}{D_1^2 \|P\|^2} \sum_{t=0}^{T-1} \|\psi_t\|^2,$$

which implies that

$$\frac{\text{ALG} - \text{OPT}}{\text{OPT}} \leq \|H\| \left(\frac{\lambda^2}{\text{OPT}} \varepsilon + \frac{(1-\lambda)^2}{C} \right)$$

where $C := \frac{D_0(1-\rho)^2}{D_1^2 \|P\|^2}$ and

$$\varepsilon := \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P (w_\tau - \widehat{w}_\tau) \right\|^2.$$

To obtain the second bound, noting that

$$\begin{aligned} \text{ALG} - \text{OPT} &\leq \|H\| \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \lambda \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2 \\ &\leq \|H\| \sum_{t=0}^{T-1} \left(\left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau \right\|^2 + \lambda^2 \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2 \right). \end{aligned}$$

Noting that $W := \sum_{t=0}^{T-1} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{w}_\tau \right\|^2$, therefore,

$$\frac{\text{ALG} - \text{OPT}}{\text{OPT}} \leq \|H\| \left(\frac{1}{C} + \frac{\lambda^2}{\text{OPT}} W \right).$$

for some constant $C > 0$ that depends on A, B, Q and R .

C REGRET ANALYSIS OF SELF-TUNING CONTROL

Throughout, for notational convenience, we write

$$W(t) := \sum_{s=0}^t \eta(\widehat{w}; s, t)^\top H \eta(\widehat{w}; s, t), \quad \text{and} \quad V(t) := \sum_{s=0}^t \eta(w; s, t)^\top H \eta(\widehat{w}; s, t)$$

where

$$\eta(w; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P w_\tau, \quad \text{and} \quad \eta(\widehat{w}; s, t) := \sum_{\tau=s}^t (F^\top)^{\tau-s} P \widehat{w}_\tau.$$

C.1 Proof of Lemma 1

In this section, we show the proof of Lemma 2 and Lemma 1. We begin with rewriting $\lambda_t - \lambda_T$ as below.

$$\lambda_t - \lambda_T = \frac{V(t-1)}{W(t-1)} - \frac{V(T-1)}{W(T-1)} = \frac{\frac{V(t-1)}{t-1}}{\frac{W(t-1)}{(t-1)}} - \frac{\frac{V(T-1)}{T-1}}{\frac{W(T-1)}{T-1}}. \quad (22)$$

Applying Lemma 6, it suffices to prove that for any $1 \leq t \leq T$, $\left| \frac{1}{T} V(T) - \frac{1}{t} V(t) \right| \leq \frac{C_1}{t}$ and $\left| \frac{1}{T} W(T) - \frac{1}{t} W(t) \right| \leq \frac{C_2}{t}$ for some constants $C_1 > 0$ and $C_2 > 0$. In the sequel, we show the bound on $\left| \frac{1}{T} V(T) - \frac{1}{t} V(t) \right|$ and the bound on $\left| \frac{1}{T} W(T) - \frac{1}{t} W(t) \right|$ follows using the same argument. Continuing from (22),

$$\begin{aligned} \left| \frac{1}{T} V(T) - \frac{1}{t} V(t) \right| &\leq \underbrace{\left| \frac{1}{T} \sum_{s=0}^T \eta(w; s, T)^\top H \eta(\widehat{w}; s, T) - \frac{1}{t} \sum_{s=0}^t \eta(w; s, T)^\top H \eta(\widehat{w}; s, T) \right|}_{=:(a)} \\ &\quad + \underbrace{\left| \frac{1}{t} \sum_{s=0}^t \eta(w; s, t)^\top H \eta(\widehat{w}; s, t) - \frac{1}{t} \sum_{s=0}^t \eta(w; s, T)^\top H \eta(\widehat{w}; s, T) \right|}_{=:(b)}. \quad (23) \end{aligned}$$

In the following, we deal with the terms (a) and (b) separately.

C.1.1 Upper bound on (a). To bound the term (a) in (23), we notice that (a) can be regarded as a difference between two algebraic means. Rewriting the first mean in (a), we get

$$\begin{aligned} \sum_{s=0}^T \eta(w; s, T)^\top H \eta(\widehat{w}; s, T) &= \sum_{s=0}^T \left(\sum_{\tau=s}^T (F^\top)^{\tau-s} P w_\tau \right)^\top H \left(\sum_{\tau=s}^T (F^\top)^{\tau-s} P \widehat{w}_\tau \right) \\ &= \sum_{s=0}^T \left(\sum_{\tau=0}^{T-s} (F^\top)^\tau P w_{\tau+s} \right)^\top H \left(\sum_{\tau=0}^{T-s} (F^\top)^\tau P \widehat{w}_{\tau+s} \right) \\ &= \sum_{s=0}^T \bar{\eta}(w; s, T)^\top H \bar{\eta}(\widehat{w}; s, T) \end{aligned}$$

where for notational convenience, for $s \leq T$ we have defined two series

$$\bar{\eta}(\widehat{w}; s, T) := \sum_{\tau=0}^s (F^\top)^\tau P \widehat{w}_{\tau+T-s}, \quad \text{and} \quad \bar{\eta}(w; s, T) := \sum_{\tau=0}^s (F^\top)^\tau P w_{\tau+T-s}.$$

We state a lemma below, which states that the sequence $(\bar{\eta}(\widehat{w}; 0, T), \dots, \bar{\eta}(\widehat{w}; T, T))$ satisfies the assumption in Lemma 7.

LEMMA 8. *Given an integer s with $0 \leq s \leq T$, we have*

$$\begin{aligned} &|\bar{\eta}(w; T, T)^\top H \bar{\eta}(\widehat{w}; T, T) - \bar{\eta}(w; s, T)^\top H \bar{\eta}(\widehat{w}; s, T)| \\ &\leq 2 \|H\| \left(\frac{C \|P\|}{1 - \rho} \right)^2 \left(2\rho^{s+1} \bar{w} \widehat{w} + \max_{\tau} \|\widehat{w}_\tau - \widehat{w}_{\tau+T-s}\| + \max_{\tau} \|w_\tau - w_{\tau+T-s}\| \right). \end{aligned}$$

PROOF OF LEMMA 8. With $s \leq T$, according to the definitions of $\bar{\eta}(w; s, T)$, $\bar{\eta}(w; T, T)$, $\bar{\eta}(\widehat{w}; s, T)$ and $\bar{\eta}(\widehat{w}; T, T)$, we obtain

$$\begin{aligned} \bar{\eta}(w; T, T) &= \bar{\eta}(w; s, T) + \sum_{\tau=s+1}^T (F^\top)^\tau P w_\tau + \sum_{\tau=0}^s (F^\top)^\tau P (w_\tau - w_{\tau+T-s}), \\ \bar{\eta}(\widehat{w}; T, T) &= \bar{\eta}(\widehat{w}; s, T) + \sum_{\tau=s+1}^T (F^\top)^\tau P \widehat{w}_\tau + \sum_{\tau=0}^s (F^\top)^\tau P (\widehat{w}_\tau - \widehat{w}_{\tau+T-s}), \end{aligned}$$

implying that

$$\begin{aligned} &\bar{\eta}(w; T, T)^\top H \bar{\eta}(\widehat{w}; T, T) - \bar{\eta}(w; s, T)^\top H \bar{\eta}(\widehat{w}; s, T) \\ &= \bar{\eta}(w; s, T)^\top H \xi_2 + \xi_1^\top H \bar{\eta}(\widehat{w}; s, T) + \xi_1^\top H \xi_2 \end{aligned} \tag{24}$$

where

$$\begin{aligned} \xi_1 &:= \sum_{\tau=s+1}^T (F^\top)^\tau P w_\tau + \sum_{\tau=0}^s (F^\top)^\tau P (w_\tau - w_{\tau+T-s}), \\ \xi_2 &:= \sum_{\tau=s+1}^T (F^\top)^\tau P \widehat{w}_\tau + \sum_{\tau=0}^s (F^\top)^\tau P (\widehat{w}_\tau - \widehat{w}_{\tau+T-s}) \end{aligned}$$

By our model assumption, $\|w_t\| \leq \omega$ and $\|\widehat{w}_t\| \leq \bar{w}$ for all $t = 0, \dots, T-1$. Then, there exists some $e > 0$ such that the prediction error $e_t = \widehat{w}_t - w_t$ satisfies $e_t \leq e$ for all $t = 0, \dots, T-1$. Note that $F = A - BK$ and we define $\rho := \frac{1+\rho(F)}{2} < 1$ where $\rho(F)$ denotes the spectral radius of F . From

Gelfand's formula, there exists a constant $C \geq 0$ such that $\|F^t\| \leq C\rho^t$ for all $t \geq 0$. The following holds for $\bar{\eta}(\hat{\mathbf{w}}; s, T)$ and $\bar{\eta}(\mathbf{w}; s, T)$:

$$\|\bar{\eta}(\hat{\mathbf{w}}; s, T)\| \leq \sum_{\tau=0}^s \|F^\tau\| \|P\| \bar{\mathbf{w}} \leq C \frac{1 - \rho^{s+1}}{1 - \rho} \|P\| \bar{\mathbf{w}} \leq \frac{C}{1 - \rho} \|P\| \bar{\mathbf{w}}, \quad (25)$$

$$\|\bar{\eta}(\mathbf{w}; s, T)\| \leq \sum_{\tau=0}^s \|F^\tau\| \|P\| \bar{\mathbf{w}} = C \frac{1 - \rho^{s+1}}{1 - \rho} \|P\| \bar{\mathbf{w}} \leq \frac{C}{1 - \rho} \|P\| \bar{\mathbf{w}}. \quad (26)$$

Moreover,

$$\|\xi_1\| \leq \sum_{\tau=s+1}^T \|F^\tau\| \|P\| \bar{\mathbf{w}} + \sum_{\tau=0}^s \|F^\tau\| \|P\| \|\mathbf{w}_\tau - \mathbf{w}_{\tau+T-s}\| \quad (27)$$

$$\leq \frac{C\|P\|}{1 - \rho} \left(\bar{\mathbf{w}}\rho^{s+1} + \max_{\tau} \|\mathbf{w}_\tau - \mathbf{w}_{\tau+T-s}\| \right) \quad (28)$$

$$\|\xi_2\| \leq \sum_{\tau=s+1}^T \|F^\tau\| \|P\| \hat{\mathbf{w}} + \sum_{\tau=0}^s \|F^\tau\| \|P\| \|\hat{\mathbf{w}}_\tau - \hat{\mathbf{w}}_{\tau+T-s}\| \quad (29)$$

$$\leq \frac{C\|P\|}{1 - \rho} \left(\hat{\mathbf{w}}\rho^{s+1} + \max_{\tau} \|\hat{\mathbf{w}}_\tau - \hat{\mathbf{w}}_{\tau+T-s}\| \right). \quad (30)$$

Combining (25)-(30) with (24),

$$\begin{aligned} & \left| \bar{\eta}(\mathbf{w}; T, T)^\top H \bar{\eta}(\hat{\mathbf{w}}; T, T) - \bar{\eta}(\mathbf{w}; s, T)^\top H \bar{\eta}(\hat{\mathbf{w}}; s, T) \right| \\ & \leq 2\|H\| \left(\frac{C\|P\|}{1 - \rho} \right)^2 \left(2\rho^{s+1} \bar{\mathbf{w}}\hat{\mathbf{w}} + \max_{\tau} \|\hat{\mathbf{w}}_\tau - \hat{\mathbf{w}}_{\tau+T-s}\| + \max_{\tau} \|\mathbf{w}_\tau - \mathbf{w}_{\tau+T-s}\| \right). \end{aligned}$$

□

Therefore, applying Lemma 7, we conclude that

$$\begin{aligned} (a) & := \left| \frac{1}{T} \sum_{s=0}^T \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, T) - \frac{1}{t} \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, T) \right| \\ & \leq \frac{4}{t} \|H\| \rho \left(\frac{C\|P\|}{(1 - \rho)^{3/2}} \right)^2 \bar{\mathbf{w}}\hat{\mathbf{w}} + \frac{2}{t} \|H\| \left(\frac{C\|P\|}{1 - \rho} \right)^2 (\mu_{\text{VAR}}(\hat{\mathbf{w}}) + \mu_{\text{VAR}}(\mathbf{w})) \end{aligned} \quad (31)$$

where $\mu_{\text{VAR}}(\mathbf{x}) := \sum_{s=0}^T \max_{\tau} \|\mathbf{x}_\tau - \mathbf{x}_{\tau+T-s}\|$ denotes the self-variation of a sequence $\mathbf{x}$.

C.1.2 Upper bound on (b). Next, we provide a bound on (b) in (23). For (b), we have

$$\begin{aligned} (b) & := \frac{1}{t} \left| \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\hat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, T) \right| \\ & \leq \frac{1}{t} \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, T) \right| \\ & \quad + \frac{1}{t} \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\hat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\hat{\mathbf{w}}; s, t) \right|. \end{aligned} \quad (32)$$

Noting that $\eta(\widehat{\mathbf{w}}; s, T) - \eta(\widehat{\mathbf{w}}; s, t) = \sum_{\tau=t+1}^T (F^\top)^{\tau-s} P \widehat{\mathbf{w}}_\tau$, we obtain

$$\begin{aligned} & \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, T) \right| \\ &= \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H (\eta(\widehat{\mathbf{w}}; s, t) - \eta(\widehat{\mathbf{w}}; s, T)) \right| \\ &= \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \left(\sum_{\tau=t+1}^T (F^\top)^{\tau-s} P \widehat{\mathbf{w}}_\tau \right) \right| \end{aligned} \quad (33)$$

and similarly,

$$\begin{aligned} & \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\widehat{\mathbf{w}}; s, t) \right| \\ &= \left| \sum_{s=0}^t (\eta(\mathbf{w}; s, T) - \eta(\mathbf{w}; s, t))^\top H \eta(\widehat{\mathbf{w}}; s, t) \right| \\ &= \left| \sum_{s=0}^t \left(\sum_{\tau=t+1}^T (F^\top)^{\tau-s} P \mathbf{w}_\tau \right)^\top H \eta(\widehat{\mathbf{w}}; s, t) \right|. \end{aligned} \quad (34)$$

By our assumption, $\|\mathbf{w}_t\| \leq \bar{\mathbf{w}}$ and $\|\widehat{\mathbf{w}}_t\| \leq \widehat{\mathbf{w}}$ for all $t = 0, \dots, T-1$. Therefore, for any $s \leq t$:

$$\left\| \sum_{\tau=t+1}^T (F^\top)^{\tau-s} P \widehat{\mathbf{w}}_\tau \right\| \leq \frac{C \rho^{t-s+1} \|P\| \widehat{\mathbf{w}}}{1 - \rho} \quad (35)$$

and

$$\|\eta(\mathbf{w}; s, T)\| = \left\| \sum_{\tau=s}^T (F^\top)^{\tau-s} P \mathbf{w}_\tau \right\| \leq \frac{C \|P\| \bar{\mathbf{w}}}{1 - \rho}. \quad (36)$$

Plugging (35) and (36) into (33),

$$\begin{aligned} & \left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, T) \right| \\ &\leq 2C \|H\| \left(\frac{\|P\|}{1 - \rho} \right)^2 \bar{\mathbf{w}} \widehat{\mathbf{w}} \sum_{s=0}^t \|F^{t-s+1}\| \\ &\leq 2 \|H\| \left(\frac{C \|P\|}{1 - \rho} \right)^2 \frac{\rho (1 - \rho^t)}{1 - \rho} \bar{\mathbf{w}} \widehat{\mathbf{w}} \\ &\leq 2 \|H\| \left(\frac{C \|P\|}{(1 - \rho)^{3/2}} \right)^2 \rho \bar{\mathbf{w}} \widehat{\mathbf{w}}. \end{aligned} \quad (37)$$

Using the same argument, the following bound holds for (34):

$$\left| \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\widehat{\mathbf{w}}; s, t) \right| \leq 2 \|H\| \left(\frac{C \|P\|}{(1 - \rho)^{3/2}} \right)^2 \rho \bar{\mathbf{w}} \widehat{\mathbf{w}}. \quad (38)$$

Combining (37) and (38) and using (32),

$$\begin{aligned} (b) &:= \frac{1}{t} \left| \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\widehat{\mathbf{w}}; s, t) - \sum_{s=0}^t \eta(\mathbf{w}; s, T)^\top H \eta(\widehat{\mathbf{w}}; s, T) \right| \\ &\leq \frac{4}{t} \|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \rho \overline{\mathbf{w}} \widehat{\mathbf{w}}. \end{aligned} \quad (39)$$

Finally, together, (31) and (39) imply the following:

$$\begin{aligned} \left| \frac{1}{T} V(T) - \frac{1}{t} V(t) \right| &\leq \frac{8}{t} \|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \rho \overline{\mathbf{w}} \widehat{\mathbf{w}} \\ &\quad + \frac{2}{t} \|H\| \left(\frac{C\|P\|}{1-\rho} \right)^2 (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})). \end{aligned} \quad (40)$$

The same argument also guarantees that

$$\begin{aligned} \left| \frac{1}{T} W(T) - \frac{1}{t} W(t) \right| &\leq \frac{8}{t} \|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \rho \widehat{\mathbf{w}}^2 \\ &\quad + \frac{4}{t} \|H\| \left(\frac{C\|P\|}{1-\rho} \right)^2 \mu_{\text{VAR}}(\widehat{\mathbf{w}}). \end{aligned} \quad (41)$$

The following lemma together with (40) and (41) justify the conditions needed to apply Lemma 6.

LEMMA 9. *For any integer $1 \leq t \leq T$,*

$$\frac{V(t)}{t} \leq 2\|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \overline{\mathbf{w}} \widehat{\mathbf{w}}$$

where $C > 0$ is some constant satisfying $\|F^t\| \leq C\rho^t$ for all $t \geq 0$.

PROOF OF LEMMA 9. We have

$$\begin{aligned} \frac{V(t)}{t} &= \frac{1}{t} \sum_{s=0}^t \eta(\mathbf{w}; s, t)^\top H \eta(\widehat{\mathbf{w}}; s, t) \\ &\leq \frac{\|H\|}{t} \sum_{s=0}^t \left\| \sum_{\tau=0}^{t-1-s} (F^\top)^\tau P \mathbf{w}_{\tau+s} \right\| \left\| \sum_{\tau=0}^{t-1-s} (F^\top)^\tau P \widehat{\mathbf{w}}_{\tau+s} \right\| \\ &\leq \frac{\|H\|}{t} \left(\frac{C\|P\|}{1-\rho} \right)^2 \sum_{s=0}^t (1-\rho^{t-s}) \overline{\mathbf{w}} \widehat{\mathbf{w}} \\ &= \frac{\|H\|}{t} \left(\frac{C\|P\|}{1-\rho} \right)^2 \left(t + \frac{1-\rho^{t+1}}{1-\rho} \right) \overline{\mathbf{w}} \widehat{\mathbf{w}} \\ &\leq 2\|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \overline{\mathbf{w}} \widehat{\mathbf{w}}. \end{aligned}$$

□

First, based on our assumption, $\lambda_t = V(t)/W(t) = V_t/W_t \leq 1$. Moreover, $W(T)/T = \Omega(1)$. Therefore, using (40), (41), Lemma 6 and Lemma 9, (22) implies that for any $1 < t \leq T$,

$$\begin{aligned} |\lambda_t - \lambda_T| &\leq \frac{1}{t-1} \frac{\|H\| \left(\frac{C\|P\|}{1-\rho} \right)^2}{W(T)/T} \cdot \left(\frac{8\rho\widehat{\mathbf{w}}\overline{\mathbf{w}}}{1-\rho} + 2(\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})) \right) \\ &\quad + \frac{2\|H\| \left(\frac{C\|P\|}{(1-\rho)^{3/2}} \right)^2 \overline{\mathbf{w}}\widehat{\mathbf{w}}}{W(T)/T} \left(\frac{8\rho\widehat{\mathbf{w}}^2}{1-\rho} + 4\mu_{\text{VAR}}(\widehat{\mathbf{w}}) \right) \\ &= O\left(\frac{\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})}{t} \right). \end{aligned}$$

C.2 Proof of Lemma 2

Using Lemma 8,

$$|\lambda_t - \lambda_T| \leq \frac{C}{t} (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})), \quad \text{where } C_1 > 0 \text{ is some constant.}$$

Applying Lemma 5, and noting that

$$\left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{\mathbf{w}}_\tau \right\| \leq C \frac{1-\rho^{T-t}}{1-\rho} \|P\| \widehat{\mathbf{w}},$$

(15) implies

$$\begin{aligned} \text{Regret} &\leq C_1^2 \|H\| \sum_{t=1}^{T-1} \left\| \frac{\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}})}{t} \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{\mathbf{w}}_\tau \right\|^2 + C_0 \\ &= C_1^2 \|H\| (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2 \sum_{t=1}^{T-1} \frac{1}{t^2} \left\| \sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \widehat{\mathbf{w}}_\tau \right\|^2 + C_0 \\ &\leq C_1^2 \|H\| (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2 \left(\frac{C\|P\|}{1-\rho} \widehat{\mathbf{w}} \right)^2 \sum_{t=2}^{T-1} \frac{1}{t^2} + C_0 \\ &\leq \frac{C_1^2 \pi^2}{6} \|H\| (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2 \left(\frac{C\|P\|}{1-\rho} \widehat{\mathbf{w}} \right)^2 + C_0 \end{aligned} \tag{42}$$

where

$$C_0 := \|H\| |\lambda_T - \lambda_0| \left\| \sum_{\tau=0}^{T-1} (F^\top)^\tau P \widehat{\mathbf{w}}_\tau \right\|^2 \leq \|H\| |\lambda_T - \lambda_0| \left(\frac{C\|P\|}{1-\rho} \widehat{\mathbf{w}} \right)^2.$$

Moreover, for any $t = 1, \dots, T$, $|\lambda_t| \leq 1$, whence,

$$C_0 \leq 2\|H\| \left(\frac{C\|P\|}{1-\rho} \widehat{\mathbf{w}} \right)^2.$$

Therefore, continuing from (42),

$$\begin{aligned} \text{Regret} &\leq \|H\| \left(\frac{C\|P\|}{1-\rho} \widehat{\mathbf{w}} \right)^2 \left(\frac{C_1 \pi^2}{6} (\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2 + 2 \right) \\ &= O\left((\mu_{\text{VAR}}(\mathbf{w}) + \mu_{\text{VAR}}(\widehat{\mathbf{w}}))^2 \right). \end{aligned}$$

D PROOF OF THEOREM 3.1

First, note that the total cost is given by $J = \sum_{t=0}^{T-1} x_t^\top Q x_t + u_t^\top R u_t + x_T^\top P x_T$. Since we can choose a threshold $\sigma > 0$ arbitrarily small, the error must exceed a threshold σ . Without loss of generality, we suppose the accumulated error δ exceeds the threshold σ at time $s \geq 0$ and assume the predictions $\hat{w}_t$, $0 < t < s - 1$ are accurate.

Throughout, we define $J_1 := \sum_{t=1}^{s-1} x_t^\top Q x_t + u_t^\top R u_t$ and $J_2 := \sum_{t=s}^{T-1} x_t^\top Q x_t + u_t^\top R u_t$ and use diacritical letters $\hat{J}, \hat{x}$ and $\hat{u}$ to denote the corresponding cost, action and state of the threshold algorithm (Algorithm 1). We consider the best online algorithm (with no predictions available) that minimizes its corresponding competitive ratio and use diacritical letters $\tilde{J}, \tilde{x}$ and $\tilde{u}$ to denote the corresponding cost, action and state. The competitive ratio of the best online algorithm is denoted by $C_{\min}$.

D.1 Upper Bound on $\hat{J}_1$

We first provide an upper bound on $\hat{J}_1$, the first portion of the total cost. For $1 \leq t < s$, the threshold-based algorithm gives

$$\begin{aligned} \hat{u}_t &= -K\hat{x}_t - (R + B^\top P B)^{-1} B^\top \left(\sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P \hat{w}_\tau \right) \\ &= -K\hat{x}_t - (R + B^\top P B)^{-1} B^\top \left(\sum_{\tau=t}^{T-1} (F^\top)^{\tau-t} P w_\tau - \eta_t \right). \end{aligned}$$

Lemma 10 in [28] implies

$$J_1 = \text{ALG}(0 : T) - \text{ALG}(s : T)$$

where

$$\begin{aligned} \text{ALG}(0 : T) &= \sum_{t=0}^{T-1} \left(w_t^\top P w_t + 2w_t^\top \sum_{i=1}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad - \sum_{t=0}^{T-1} \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right)^\top H \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad + \sum_{t=0}^{T-1} \eta_t^\top H \eta_t + x_0^\top P x_0 + 2x_0^\top \sum_{i=0}^{T-1} (F^\top)^{i+1} P w_i, \end{aligned} \quad (43)$$

and

$$\begin{aligned} \text{ALG}(s : T) &:= \sum_{t=0}^{T-s-1} \left(w_{t+s}^\top P w_{t+s} + 2w_{t+s}^\top \sum_{i=1}^{T-s-t-1} (F^\top)^i P w_{t+s+i} \right) \\ &\quad - \sum_{t=0}^{T-s-1} \left(\sum_{i=0}^{T-s-t-1} (F^\top)^i P w_{t+s+i} \right)^\top H \left(\sum_{i=0}^{T-s-t-1} (F^\top)^i P w_{t+s+i} \right) \\ &\quad + \sum_{t=0}^{T-s-1} \eta_{t+s}^\top H \eta_{t+s} + x_s^\top P x_s + 2x_s^\top \sum_{i=0}^{T-s-1} (F^\top)^{i+1} P w_{i+s}. \end{aligned} \quad (44)$$

Rewriting (44),

$$\begin{aligned} \text{ALG}(s : T) &:= \sum_{t=s}^{T-1} \left(w_t^\top P w_t + 2w_t^\top \sum_{i=1}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad - \sum_{t=s}^{T-1} \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right)^\top H \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad + \sum_{t=s}^{T-1} \eta_t^\top H \eta_t + x_s^\top P x_s + 2x_s^\top \sum_{i=s}^{T-1} (F^\top)^{i+1-s} P w_i. \end{aligned} \quad (45)$$

Therefore, combining (43) and (45),

$$\begin{aligned} J_1 &= \sum_{t=0}^{s-1} \left(w_t^\top P w_t + 2w_t^\top \sum_{i=1}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad - \sum_{t=0}^{s-1} \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right)^\top H \left(\sum_{i=0}^{T-t-1} (F^\top)^i P w_{t+i} \right) \\ &\quad + \sum_{t=0}^{s-1} \eta_t^\top H \eta_t + x_0^\top P x_0 + 2x_0^\top \sum_{i=0}^{T-1} (F^\top)^{i+1} P w_i - x_s^\top P x_s - 2x_s^\top \sum_{i=s}^{T-1} (F^\top)^{i+1-s} P w_i. \end{aligned}$$

Denote by $\Delta J_1 := |J_1 - \widehat{J}_1|$. We obtain

$$\begin{aligned} \Delta J_1 &= \sum_{t=0}^{s-1} \eta_t^\top H \eta_t + x_s^\top P x_s - \widehat{x}_s^\top P \widehat{x}_s + 2(x_s - \widehat{x}_s)^\top \sum_{i=s}^{T-1} (F^\top)^{i+1-s} P w_i \\ &= \sum_{t=0}^{s-1} \eta_s^\top F^{s-t} H ((F^\top)^{s-t} \eta_s) + x_s^\top P x_s - \widehat{x}_s^\top P \widehat{x}_s + 2(x_s - \widehat{x}_s)^\top \sum_{i=s}^{T-1} (F^\top)^{i+1-s} P w_i \\ &\leq \frac{c \|H\|}{1 - \rho^2} \frac{c^2 \|P\|^2 R^2}{(1 - \rho)^2} + 2 \|P\| \|x_s\| \|x_s - \widehat{x}_s\| + \|x_s - \widehat{x}_s\|^2 + 2 \|x_s - \widehat{x}_s\| \frac{c \|P\| \rho}{1 - \rho}. \end{aligned}$$

Since the following is true:

$$\begin{aligned} x_s - \widehat{x}_s &= A(x_{s-1} - \widehat{x}_{s-1}) + B(u_{s-1} - \widehat{u}_{s-1}) \\ &= (A - BK)(x_{s-1} - \widehat{x}_{s-1}) + B(R + B^\top PB)^{-1} B^\top \eta_{s-1} \\ &= \sum_{t=0}^{s-1} (F^\top)^{s-t-1} B(R + B^\top PB)^{-1} B^\top \eta_t, \end{aligned}$$

we have

$$\|x_s - \widehat{x}_s\| \leq \frac{c^2 \|B(R + B^\top PB)^{-1} B^\top\| R}{(1 - \rho)^2}.$$

If $\|x_s\| = O(1)$, then $\Delta J_1 = O(1)$, else

$$\frac{\Delta J_1}{J_1} \leq \frac{O(1) \cdot \|x_s\| + O(1)}{x_s^\top Q x_s} \rightarrow 0.$$

Therefore, as a conclusion, $\widehat{J}_1$ can be bounded from above by

$$\widehat{J}_1 \leq J_1 + O(1). \quad (46)$$

D.2 Upper Bound on $\widehat{J}_2$

For section D.1, we know that $\|x_s - \widehat{x}_s\| = O(1)$. Let $\widetilde{J}_2$ denote the cost by running 1-confident algorithm from $\widehat{x}_s$ with correct prediction, and $\widetilde{x}_t$ denote the state we get in the procedure. Then

$$\|x_t - \widetilde{x}_t\| = \|(A - BK)(x_{t-1} - \widetilde{x}_{t-1})\| = \|F^{t-s}(x_s - \widetilde{x}_s)\| = \|F^{t-s}(x_s - \widehat{x}_s)\|.$$

Therefore,

$$\begin{aligned} |J_2 - \widetilde{J}_2| &\leq \left| \sum_{t=s}^{T-1} (\widetilde{x}_t - x_t)^\top Q x_t + x_t^\top Q (\widetilde{x}_t - x_t) + (\widetilde{x}_t - x_t)^\top Q (\widetilde{x}_t - x_t) \right| \\ &\quad + \left| \sum_{t=s}^{T-1} (\widetilde{u}_t - u_t)^\top R u_t + u_t^\top R (\widetilde{u}_t - u_t) + (\widetilde{u}_t - u_t)^\top R (\widetilde{u}_t - u_t) \right| \\ &\quad + \left| (\widetilde{x}_T - x_T)^\top P x_T + x_T^\top P (\widetilde{x}_T - x_T) + (\widetilde{x}_T - x_T)^\top P (\widetilde{x}_T - x_T) \right| \\ &\leq \sum_{t=s}^{T-1} (\|Q\| + \|K^\top R K\|) \|F^{2t-2s}\| \|x_s - \widehat{x}_s\|^2 \\ &\quad + \sum_{t=s}^{T-1} 2 \|F^{t-s}\| \|x_s - \widehat{x}_s\| (\|Q\| \|x_t\| + \|RK\| \|u_t\|) \\ &\quad + 2 \|F^{T-s}\| \|P\| \|x_s - \widehat{x}_s\| \|x_T\| + \|F^{2T-2s}\| \|P\| \|x_s - \widehat{x}_s\|^2 \\ &= \sum_{t=s}^{T-1} 2 \|F^{t-s}\| \|x_s - \widehat{x}_s\| (\|Q\| \|x_t\| + \|RK\| \|u_t\|) \\ &\quad + b + 2 \|F^{T-s}\| \|P\| \|x_s - \widehat{x}_s\| \|x_T\| + O(1) \end{aligned}$$

If $\|x_t\| = O(1)$ and $\|u_t\| = O(1)$ for all t , then

$$|J_2 - \widetilde{J}_2| = O(1).$$

Otherwise, suppose $x_{i_1}, x_{i_2}, \dots, x_{i_k}$ and $u_{j_1}, u_{j_2}, \dots, u_{j_l}$ are some functions of T , then for any $1 \leq m \leq k$ and $1 \leq n \leq l$, $\|x_{i_m}\|/x_{i_m}^\top Q x_{i_m} \rightarrow 0$ and $\|u_{j_n}\|/u_{j_n}^\top R u_{j_n} \rightarrow 0$. Therefore,

$$\frac{|J_2 - \widetilde{J}_2|}{J_2} \leq 2 \|x_s - \widehat{x}_s\| \frac{\sum_{m=1}^k \|F^{i_m-s}\| \|Q\| \|x_{i_m}\| + \sum_{n=1}^l \|F^{j_n-s}\| \|RK\| \|u_{j_n}\|}{J_2} + \frac{O(1)}{J_2} \rightarrow 0.$$

Combine the two cases, we can conclude that

$$|J_2 - \widetilde{J}_2| \leq J_2 + O(1). \quad (47)$$

Therefore, from (46) and (47), we conclude that

$$\widehat{J} = \widehat{J}_1 + \widehat{J}_2 \leq J_1 + O(1) + C_{\min} \widetilde{J}_2 \leq J_1 + O(1) + C_{\min} (J_2 + O(1)) = C_{\min} J + O(1). \quad (48)$$

The proof completes by noticing that when the prediction error is zero and $\widehat{w}_t = w_t$ for all $t = 0, \dots, T-1$, the accumulated error δ will always be 0 and since the threshold σ is positive, the algorithm is always optimal and 1-consistent. As a result, Algorithm 1 is 1-consistent and $(C_{\min} + o(1))$ -robust.

E EXPERIMENTAL SETUP

In our three case studies, we consider i.i.d. prediction errors, i.e., $e_t = \widehat{w}_t - w_t$ is an i.i.d. additive prediction noise. To illustrate the effects of randomness for simulating the worst-case performance, we consider varying types of noise in the case studies. For the robot tracking case, we set $e_t = cX$ where $X \sim B(10, 0.5)$ is a binomial random variable with 10 trials and a success probability 0.5 and $c > 0$ is a scaling parameter. For the battery-buffered EV charging case, we set $e_t = Y$ where $X \sim N(0, \sigma^2)$ is a normal random variable with zero mean and σ^2 is a variance that can be varied to generate varying prediction error. For the Cart-Pole problem, we set $e_t = Zw_t$ where $w_t = 60 \times B$ with $\eta := l\left(\frac{4}{3} - \frac{m}{m+M}\right)$,

$$B := \begin{bmatrix} 0 \\ \frac{(m+M)\eta+ml}{(m+M)^2\eta} \\ 0 \\ -\frac{1}{(m+M)\eta} \end{bmatrix}$$

and $Z \sim N(0, \sigma^2)$ is a normal random variable with zero mean and σ^2 is a variance ranging between 0 to 8×10^2 . To simulate the worst-case performance of algorithms, in our experiments we run the algorithms 5 times, with a new sequence of prediction noise generated at each time and choose the one with the largest overall cost.

Finally, Table 1 and Table 2 list the detailed settings and the hyper-parameters used in the robot tracking, battery-buffered EV charging and Cart-Pole case studies.

Table 1. Hyper-parameters used in robot tracking and EV charging.

Robot Tracking	Value	EV Charging	Value
Number of Monte Carlo Tests	5	Number of Monte Carlo Tests	5
Prediction Error Type	<i>Binomial</i>	Prediction Error Type	<i>Gaussian</i>
State Dimension n	4	State Dimension n	10 (Synthetic); 52 (Realistic)
Action Dimension m	2	Action Dimension m	10 (Synthetic); 52 (Realistic)
Time Horizon Length T	Fig 2: $T = 240$ Fig 3: $T = 240$ Fig 4: $T = 200$	Time Horizon Length T	240
Initialized λ_0	0.3	Charging Efficiency	1
Scaling parameter c	Fig 4: $c \in [0, 1]$	Variance σ^2	$\sigma^2 \in [0, 10]$
CPU	Intel® i7-8850H	CPU	Intel® i7-8850H
		Energy Demand E (Synthetic Case)	5 (kWh)
		Arrival Rate (Realistic Case)	0.2

Received October 2021; revised December 2021; accepted January 2022

Table 2. Hyper-parameters used in the Cart-Pole problem.

Robot Tracking	Value
Number of Monte Carlo Tests	2000
Prediction Error Type	<i>Gaussian</i>
Action Dimension m	1
State Dimension m	4
Time Horizon Length T	200
Variance σ^2	$\sigma^2 \in [0, 800]$
Cart Mass M	$M = 10.0kg$
Pole Mass m	$m = 1.0kg$
Pole Length l	$l = 10.0m$
CPU	Intel® i7-8850H