2022 IEEE 61st Conference on Decision and Control (CDC)
December 6-9, 2022. Cancún, Mexico

Performance-Robustness Tradeoffs in Adversarially Robust Linear-Quadratic Control

Bruce D. Lee*, Thomas T.C.K. Zhang* , Hamed Hassani, and Nikolai Matni

2022 IEEE 61st Conference on Decision and Control (CDC) | 978-1-6654-6761-2/22/$31.00 ©2022 IEEE | DOI: 10.1109/CDC51059.2022.9992393

***Abstract*— While $\mathcal{H}_\infty$ methods can introduce robustness against worst-case perturbations, their nominal performance under conventional stochastic disturbances is often drastically reduced. Though this fundamental tradeoff between nominal performance and robustness is known to exist, it has not been quantitatively characterized. Toward addressing this issue, we borrow from the increasingly ubiquitous notion of adversarial training from machine learning to construct a class of controllers which are optimized for disturbances consisting of mixed stochastic and worst-case components. We find that this problem admits a stationary optimal controller that has a simple analytic form closely related to suboptimal $\mathcal{H}_\infty$ solutions. We then provide a quantitative performance-robustness tradeoff analysis, in which system-theoretic properties such as controllability and stability explicitly manifest in an interpretable manner. This provides practitioners with general guidance for determining how much robustness to incorporate based on a priori system knowledge. We empirically validate our results by comparing the performance of our controller against standard baselines, and plotting tradeoff curves.**

I. Introduction

Modern dynamical systems, from mobile robotics to power plants, require controllers that are simultaneously fast, efficient, and robust. Many control schemes attempt to achieve these desiderata by combining them into a single objective function and optimizing it, leading to a natural tradeoff. A controller optimized for speed and efficiency may perform poorly in the face of unmodeled phenomena. For instance, Linear-Quadratic Gaussian (LQG) controllers (a special case of $\mathcal{H}_2$ controllers) explicitly prioritize nominal performance by penalizing the expectation of a quadratic function of the state and input. However, such controllers can be arbitrarily fragile to small perturbations of the dynamics [1]. Replacing the LQG objective with one that considers the response of the system to worst-case dynamic uncertainty and external disturbances results in robust control methods, such as $\mathcal{H}_\infty$ and $\mathcal{L}_1$ methods [2], [3]. Such controllers are provably robust, however they tend to be overly conservative.

Toward achieving a balance between the performance of nominal and robust controllers, various approaches have been introduced, most notably mixed $\mathcal{H}_2/\mathcal{H}_\infty$ methods. However, the resulting controllers are often complicated to express and compute [4], and lack *a priori* quantitative guarantees on how much nominal performance must be given up in order to achieve a desired robustness level. Toward addressing these issues, we take inspiration from the notion of adversarial robustness from machine learning [5]–[8] and formulate a controller synthesis problem that balances performance and robustness. The goal of adversarial robustness in machine learning is to minimize the expected error under the presence of worst-case norm-bounded perturbations to the data, where the perturbations can depend on the underlying stochasticity of the problem, such as the data distribution and additive noise. We consider an analogous adversarially robust state feedback control problem where we aim to minimize an expected LQ cost subject to linear dynamics driven by process noise composed of two components: a zero-mean stochastic noise term and a norm-bounded adversarial term. We show that the solution to this problem admits a closed-form expression in terms of the solution to discrete algebraic Riccati equations (DAREs), which in turn allows for novel quantitative performance-robustness tradeoff bounds to be computed in which system parameters manifest in a natural and interpretable way.

A. Contributions

Toward analyzing the adversarially robust control problem we propose, we first show that when viewed through the lens of dynamic games [9], adversarially robust LQ control relates to a control problem introduced in [10]. We show that the optimal solution to the state feedback version of this problem is given by a central static suboptimal $\mathcal{H}_\infty$ controller, with suboptimality level γ depending on both the stochastic noise statistics and the budget given to the adversary. Furthermore, both the worst-case adversary and the corresponding optimal controller can be computed from the solution of a DARE.

We quantify the performance-robustness tradeoffs of adversarially robust LQ control, both analytically and empirically, and show a clear and interpretable dependence on underlying system theoretic properties such as controllability and stability. In particular, we show that the cost gap incurred by the adversarially robust controller in the nominal setting, relative to that achieved by the nominal controller, is upper bounded by $O(\sigma_w^2\gamma^{-4}\nu^{-1})$, where σ_w^2 is the covariance of the additive noise distribution, γ is the suboptimality level of the suboptimal $\mathcal{H}_\infty$ controller, and ν is the smallest singular value of the controllability gramian. On the other hand, the cost gap is lower bounded by $\Omega\big(\sigma_w^2\gamma^{-4}\eta^2\big)$, where η is the largest singular value of the controllability gramian for closed loop system under the nominal LQ controller with disturbances as the input.[1] These results quantitatively show that systems with uniformly good controllability have small

*Equal contribution

Department of Electrical and Systems Engineering, University of Pennsylvania. Emails: {brucele,ttz2,hassani,nmatni}@upenn.edu

[1]The controllability gramian defined by the pair $(A + BK_\star, I)$, for $K_\star$ the optimal LQR controller.

performance-robustness tradeoffs, while those that have a highly controllable mode in the nominal closed-loop system (when viewing disturbances as inputs) lead to large performance-robustness tradeoffs. We note that all proofs and further discussion may be found in [11].

B. Related Work

The mixed stochastic/worst-case problem that we consider is not the only way to strike a balance between the performance of stochastic and robust control methods. Most closely related is [10], which considers a similar problem from a deterministic perspective, in which disturbances are composed of both a bounded power component, and a bounded power spectrum component. A set description of disturbances that also interpolates between $\mathcal{H}_2$ and $\mathcal{H}_\infty$ approaches is proposed in [12]. The class of all stabilizing controllers subject to a $\mathcal{H}_\infty$ norm constraint is characterized in [13], while minimizing an $\mathcal{H}_2$ objective subject to a $\mathcal{H}_\infty$ constraint is addressed in [14], [15]. While conceptually appealing, these methods often lack a simple closed-form stationary solution. Other recent work attempts to reduce the conservatism of robust control through risk aware approaches [16], [17] or regret-minimization [18]–[20]. None of the aforementioned methods provide a characterization of the performance-robustness tradeoffs of the resulting controllers.

Analogous recent work in the machine learning community has analyzed performance-robustness tradeoffs in adversarially robust learning, including precise characterizations of the generalization errors of standard versus adversarially trained models under various theoretical models [21]–[23], and "no free lunch" theorems for obtaining adversarially robust models [24]–[26]. The successful characterization of such performance-robustness tradeoffs in machine learning motivates the control objective we consider. However, the theoretical results from this area are largely intended for the supervised learning setting and do not immediately apply to our setting. The existence of performance-robustness tradeoffs in control is shown in [27], but they are not characterized quantitatively. We end by noting that the extension of adversarial robustness results in machine learning to various control problems has recently received attention [28]–[34].

Notation: The Euclidean norm of a vector x is denoted by $\|x\|$. For a matrix A, the spectral norm is denoted $\|A\|$ and the Frobenius norm is denoted $\|A\|_F$. The spectral radius of a square matrix A is denoted $\rho(A)$. A symmetric, positive semidefinite (psd) matrix $A = A^\top$ is denoted $A \succeq 0$, and a symmetric positive definite (pd) matrix is denoted $A \succ 0$. Similarly $A \succeq B$ denotes that $A - B$ is positive semidefinite. A sequence of vectors x_t defined for $t \geq 0$ will be denoted by $\mathbf{x} = \{x_t\}_{t\geq 0}$. The ℓ^2 signal-norm of a sequence is denoted by $\|\mathbf{x}\|_{\ell^2} := (\sum_{t\geq 0} \|x_t\|^2)^{1/2}$. For an autonomous system $x_{t+1} = Ax_t$ and symmetric matrix Q, we denote the solution P to the discrete Lyapunov equation

$$A^\top P A - P + Q = 0$$

by $\texttt{dlyap}(A, Q)$. Similarly, for a controlled system $x_{t+1} = Ax_t + Bu_t$ and symmetric matrices Q, R of compatible size, we denote the solution P to the discrete algebraic Riccati equation

$$Q + A^\top P A - A^\top P B (B^\top P B + R)^{-1} B^\top P A = 0$$

by $\texttt{DARE}(A, B, Q, R)$.

II. Adversarially Robust Linear-Quadratic Control

Consider a fully observed discrete-time linear-time-invariant (LTI) system with state disturbances composed of both stochastic and adversarial components: let $x_t \in \mathbb{R}^n$ be the system state, $u_t \in \mathbb{R}^m$ the input, $w_t \in \mathbb{R}^n$ and $\delta_t \in \mathbb{R}^n$ the stochastic and adversarial components of the process disturbance, respectively. The initial condition x_0 and stochastic component of the process noise w_t are assumed to be i.i.d. zero-mean with covariance matrices $\Sigma_0,\ \Sigma_w$, respectively, and $\mathbb{E}\left[x_0 w_t^\top\right] = 0$ for all $t \geq 0$. The LTI system is then defined by the following equation:

$$x_{t+1} = Ax_t + Bu_t + w_t + \delta_t. \tag{1}$$

We assume that the adversarial perturbation sequence $\boldsymbol{\delta}$ is causal, i.e., that it can depend only on the states, inputs, and stochastic noise up to the current timestep. In particular, δ_t must be a measureable function of the randomness $x_0, w_{0:t}$. We let $Q \succeq 0$ be a weight matrix for the state cost, and $R \succ 0$ be a weight matrix for the input cost. We consider the infinite horizon objective

$$\limsup_{T\to\infty} \frac{1}{T} \mathbb{E}_{\mathbf{w},x_0} \left[x_T^\top Q x_T + \sum_{t=0}^{T-1} x_t^\top Q x_t + u_t^\top R u_t \right]$$

subject to the dynamics (1). If the adversarial perturbation is set to zero (i.e., $\boldsymbol{\delta} = 0$), then the above objective is the nominal LQR problem. If the stochasticity is set to zero (i.e., $\boldsymbol{w} = 0, x_0 = 0$), and $\boldsymbol{\delta}$ are worst-case perturbations with average power bounded by ε, the above objective is the $\mathcal{H}_\infty$ control problem. When both stochastic noise and worst-case perturbations are present, we define the resulting control task as the adversarially robust LQR problem. We denote the three corresponding objectives by NC, RC, and AC respectively:

$$\begin{aligned} \mathrm{NC}(K) := \limsup_{T\to\infty} \frac{1}{T} \mathbb{E}_{\mathbf{w},x_0}[V(T,K,Q,R,\mathbf{x})], \\ \text{s.t. } x_{t+1} = (A+BK)x_t + w_t \end{aligned} \tag{2}$$

$$\begin{aligned} \mathrm{RC}(K) := \limsup_{T\to\infty} \frac{1}{T} \max_{\substack{\boldsymbol{\delta} \text{ causal} \\ \|\boldsymbol{\delta}\|_{\ell_2}^2 \leq T\varepsilon}} V(T,K,Q,R,\mathbf{x}), \\ \text{s.t. } x_{t+1} = (A+BK)x_t + \delta_t,\ x_0 = 0 \end{aligned} \tag{3}$$

$$\begin{aligned} \mathrm{AC}(K) := \limsup_{T\to\infty} \frac{1}{T} \mathbb{E}_{\mathbf{w},x_0} \left[\max_{\substack{\boldsymbol{\delta} \text{ causal} \\ \|\boldsymbol{\delta}\|_{\ell_2}^2 \leq T\varepsilon}} V(T,K,Q,R,\mathbf{x}) \right], \\ \text{s.t. } x_{t+1} = (A+BK)x_t + w_t + \delta_t, \end{aligned} \tag{4}$$

where $V(T,K,Q,R,\mathbf{x}) := \sum_{t=0}^{T-1} x_t \big(Q + K^\top R K\big) x_t + x_T^\top Q x_T$. The adversarial budget $\|\boldsymbol{\delta}\|_{\ell^2}^2 \leq T\varepsilon$ is chosen such that the instance-wise adversarial budget satisfies $\|\delta_t\|^2 \leq \varepsilon$

on average.[2] We note that we restrict these definitions to static state feedback control policies $u_t = Kx_t$, as they are known to be optimal for NC and RC [2]. We will show, as a consequence of Lemma 2.1 and Theorem 2.1, that they are in fact also optimal for AC. In order to ensure that there exists a stabilizing controller, and that minimizing either NC or RC provides a stabilizing controller, we make the standard assumption that $(A, B, Q^{1/2})$ is stabilizable and detectable [2]. Under this assumption, there exists a stabilizing state feedback control law $u_t = Kx_t$ minimizing NC, where

$$K = -(R + B^\top PB)^{-1}B^\top PA \tag{5}$$
$$P = \texttt{DARE}(A, B, Q, R). \tag{6}$$

A solution minimizing RC may be found using Theorem 13.3.3 of [35].

The remainder of this section is devoted to finding a controller minimizing AC. Inspired by minimax dynamic games [9], we first find a controller which minimizes a soft-constrained version of the adversarial cost (4):

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\max_{\boldsymbol{\delta} \text{ causal}} V^\gamma(T, Q, R, \mathbf{x}, \boldsymbol{\delta})\right], \tag{7}$$

where $V^\gamma(T, Q, R, \mathbf{x}, \boldsymbol{\delta}) := x_T^\top Q x_T + \sum_{t=0}^{T-1} x_t(Q + K^\top RK)x_t - \gamma^2\delta_t^\top\delta_t$. The following lemma provides necessary and sufficient conditions for the existence of a stabilizing controller which minimizes the above objective.

Lemma 2.1: A controller attaining a finite value for objective (7) exists if and only if there exists a solution to the following discrete algebraic Riccati equation

$$P = \texttt{DARE}\left(A, \begin{bmatrix}B & I\end{bmatrix}, Q, \begin{bmatrix}R & 0\\ 0 & -\gamma^2 I\end{bmatrix}\right) \tag{8}$$

satisfying $0 \preceq P \prec \gamma^2 I$. When the above condition holds,

1) The controller $u_t = Kx_t$ with K given by

$$\begin{aligned} K &= -(R + B^\top MB)^{-1}B^\top MA,\\ M &= P + P(\gamma^2 I - P)^{-1}P. \end{aligned} \tag{9}$$

satisfies $\rho(A + BK) < 1$, and minimizes objective (7).

2) The optimal adversarial perturbation under the controller $u_t = Kx_t$ is given by

$$\begin{aligned} \Delta &= (\gamma^2 I - P)^{-1}P,\\ \delta_t &= \Delta((A + BK)x_t + w_t). \end{aligned} \tag{10}$$

3) The objective value (7) achieved under controller (9) and adversary (10) is $\mathrm{Tr}(M\Sigma_w)$.

The solution approach for the above problem follows that in [9] for minimax games. The finite horizon version of the problem is solved by defining a convex-concave saddle point cost-to-go, then recursing backwards in time. The causality of $\boldsymbol{\delta}$ in the stochastic signals allows the dynamic programing step to be solved in closed-form. Taking the limit as the horizon tends to infinity provides the steady state controller and adversary in Lemma 2.1. It should be noted that in contrast to most adversarially robust machine learning problems, adversarially robust LQR provides a closed-form expression for the adversarial perturbation.

We now return our attention to objective (4). We show (see Appendix B in [11]) via a strong duality and ergodicity argument that the hard-constrained problem may be solved by sequentially solving the soft-constrained problem using Lemma 2.1. Note that in contrast to the solution approach to minimize RC, dualizing the constraint in AC results in an optimal dual variable $\gamma(\varepsilon)$ that is a random variable. Therefore, it is nontrivial to exchange the order of the minimization over the dual variable with the expectation. We propose Algorithm 1 to minimize the adversarial cost (4), and Theorem 2.1 establishes its correctness.

Algorithm 1 Computing Adversarially Robust Controller: $\texttt{AdvLQR}(A, B, Q, R, \varepsilon, \gamma_{LB}, \gamma_{UB}, \texttt{tol})$

Input: State matrices A, B, cost matrices Q, R, adversary budget $\varepsilon > 0$, bounds $\gamma_{LB} < \gamma_{UB}$, tolerance $\texttt{tol}$.
// Do binary search on $\gamma \in [\gamma_{LB}, \gamma_{UB}]$ to find optimal adversary with average power $\varepsilon > 0$
While $\gamma_{UB} - \gamma_{LB} \geq \texttt{tol}$:
  $\gamma = (\gamma_{UB} + \gamma_{LB})/2$
  Compute P, M, K, Δ at level γ via (8)-(10)
  $G = \texttt{dlyap}((A + BK)^\top\Delta, \Delta\Sigma_w\Delta)$
  If $\mathrm{Tr}(G(A + BK)^\top(I + \Delta)^2(A + BK)$
    $+\Delta\Sigma_w\Delta) < \varepsilon$ (11):
  then $\gamma_{UB} = \gamma$
  else $\gamma_{LB} = \gamma$
Output: Adversarially robust LQR controller K, adversary gain Δ, optimal value of (4) $\mathrm{Tr}(M\Sigma_w) + \gamma^2\varepsilon$.

Theorem 2.1: Suppose $(A, B, Q^{1/2})$ is stabilizable and detectable. For dynamics (1), let γ_∞ denote the $\mathcal{H}_\infty$-norm of the optimal closed-loop system. Given any fixed $\varepsilon > 0$, let P, M, K, Δ satisfy equations (8)-(10) at level $\gamma_{UB} > \gamma_\infty$ and define $G := \texttt{dlyap}((A + BK)^\top\Delta, \Delta\Sigma_w\Delta)$. If γ_{UB} is sufficiently large such that the disturbance δ_t defined in equation (10) satisfies

$$\begin{aligned} &\lim_{t\to\infty}\mathbb{E}\left[\delta_t^\top\delta_t\right] =\\ &\mathrm{Tr}(G(A + BK)^\top(I + \Delta)^2(A + BK) + \Delta\Sigma_w\Delta) < \varepsilon, \end{aligned} \tag{11}$$

then, under the stated conditions, the output of Algorithm 1 $\texttt{AdvLQR}(A, B, Q, R, \varepsilon, \gamma_\infty, \gamma_{UB}, \texttt{tol})$ satisfies the following[3]:

1) The control policy $u_t = Kx_t$ minimizes $\mathrm{AC}(K)$.
2) The optimal adversarial perturbation in equation (4) under the optimal policy is $\delta_t = \Delta((A + BK)x_t + w_t)$ and satisfies $\lim_{t\to\infty}\mathbb{E}\left[\delta_t^\top\delta_t\right] \leq \varepsilon$.
3) The minimum value for the adversarial cost (4) is given by $\mathrm{Tr}(M\Sigma_w) + \gamma^2\varepsilon$.

We note that in contrast to the certainty equivalent LQR controller that is independent of the stochastic process noise

[2]This is equivalent to constraining the power semi-norm of $\boldsymbol{\delta}$ to be upper-bounded by $\sqrt{\varepsilon}$.

[3]Up to numerical precision due to the tolerance parameter $\texttt{tol}$.

statistics, the adversarially robust controller output by Algorithm 1 implicitly depends on the noise statistics through the optimal choice of γ.

Remark 2.1: The same results used to solve for the optimal controller minimizing AC may be adapted to evaluate $\mathrm{AC}(K)$ under an arbitrary stabilizing controller K. In particular, observe that the closed loop system and cost matrices $(A', B', Q', R') = (A+BK, 0, (Q+K^\top RK)^{1/2}, 0)$ satisfy the stabilizability and detectability assumptions. Then Lemma 2.1 implies that under controller $u_t = Kx_t$, equation (7) evaluates to $\mathrm{Tr}(M\Sigma_w)$, where $M = P + P(\gamma^2 I - P)^{-1}P$, and $P = \texttt{DARE}(A + BK, I, Q + K^\top RK, -\gamma^2 I)$. Similarly, Theorem 2.1 tells us that for any $\varepsilon > 0$, with properly selected γ_{UB}, γ_{LB} and $\texttt{tol}$, $\mathrm{AC}(K)$ may be determined via Algorithm 1 as $\texttt{AdvLQR}(A + BK, 0, Q + K^\top RK, 0, \varepsilon, \gamma_{LB}, \gamma_{UB}, \texttt{tol})$.

III. Performance-Robustness Tradeoff Bounds

This section presents the tradeoffs that arise in adversarial control by investigating the interplay between the objectives (2) and (4), and quantitatively bounds the resulting tradeoffs.

We consider the gap between the nominal and γ-adversarially robust controllers when applied in the nominal setting, i.e., we seek to bound the gap $\mathrm{NC}(K_\gamma) - \mathrm{NC}(K_\star)$, where K_γ is the γ-adversarially robust controller given by Lemma 2.1, and $K_\star$ is the LQR controller (5). Let $P_\star$ and P_γ solve the nominal (6) and modified (8) DAREs respectively, and let M_γ be given by equation (9).

Given an arbitrary stabilizing linear feedback controller K, results from [36] and [37] allow us to characterize the gap in the cost between K and the optimal LQR controller $K_\star$ as $\mathrm{NC}(K) - \mathrm{NC}(K_\star) = \mathrm{Tr}(\Sigma(K)(K - K_\star)^\top (R + B^\top P_\star B)(K - K_\star))$ where $\Sigma(K) := \texttt{dlyap}(A+BK, \Sigma_w)$ is the steady state covariance of the closed loop system under controller K. The following bounds on the gap $\mathrm{NC}(K_\gamma) - \mathrm{NC}(K_\star)$ then follow immediately:

$$\begin{aligned} &\sigma_{\min}(\Sigma_w)\sigma_{\min}(R + B^\top P_\star B)\left\|K_\gamma - K_\star\right\|_F^2 \\ &\le \mathrm{NC}(K_\gamma) - \mathrm{NC}(K_\star) \\ &\le \left\|\Sigma(K_\gamma)\right\| \left\|R + B^\top P_\star B\right\| \left\|K_\gamma - K_\star\right\|_F^2. \end{aligned} \tag{12}$$

We have therefore reduced the task of upper and lower bounding the cost gap between the γ-adversarially robust controller and the nominal LQR controller in the nominal setting to directly bounding the gap between the two controllers. Recalling that

$$\left\|K_\gamma - K_\star\right\|^2 \le \left\|K_\gamma - K_\star\right\|_F^2 \le \min\{m, n\}\left\|K_\gamma - K_\star\right\|^2,$$

we use the following lemma to bound the difference $\|K_\gamma - K_\star\|$ in terms of the difference between the solutions to the corresponding adversarial and nominal DAREs.

Lemma 3.1: (Adapted from Lemma 2 of [36]) Suppose that $f_1(u; x) = \frac{1}{2}u^\top Ru + \frac{1}{2}(Ax + Bu)^\top M(Ax + Bu)$ and $f_2(u; x) = \frac{1}{2}u^\top Ru + \frac{1}{2}(Ax + Bu)^\top P(Ax + Bu)$ with $M \succeq P$. Furthermore, for any x, let $u_i = K_i x = \mathrm{argmin}_u\, f_i(u, x)$. Then

$$\begin{aligned} &\frac{\left\|B^\top(M - P)(A + BK_2)\right\|}{\|R + B^\top MB)\|} \\ &\le \|K_1 - K_2\| \\ &\le \frac{\left\|B^\top(M - P)(A + BK_2)\right\|}{\sigma_{\min}(R + B^\top PB)}. \end{aligned}$$

Applying Lemma 3.1 with $P = P_\star$ and $M = M_\gamma$, computing upper and lower bounds on the gap (12) reduces to bounding the difference $M_\gamma - P_\star$. The upper bound is presented in Section III-A, and the lower bound in Section III-B.

For the rest of this paper, we assume for simplicity that $\Sigma_w = \sigma_w^2 I$. We also define γ_∞ as the minimum $\mathcal{H}_\infty$ norm for the closed loop system, i.e., the smallest value of γ for which the conditions of Lemma 2.1 hold. Similarly, we define $\tilde{\gamma}_\infty$ as the $\mathcal{H}_\infty$ norm of the closed loop system under the nominal LQR controller. Additionally, we define the ℓ-step controllability gramian as $W_\ell(A, B) := \sum_{t=0}^{\ell} A^t B B^\top (A^t)^\top$. If $\rho(A) < 1$ we define the controllability gramian as $W_\infty(A, B) := \lim_{\ell\to\infty} W_\ell(A, B)$.

A. Upper Bound

From the definition (9) of M_γ, we can write $\|P_\star - M_\gamma\| \le \|P_\star - P_\gamma\| + \frac{\|P_\gamma\|^2}{\gamma^2 - \|P_\gamma\|}$. For $\gamma > \gamma_\infty$ we have $P_\gamma \prec P_{\gamma_\infty} \prec \gamma_\infty^2 I$, which in turn implies

$$\|M_\gamma - P_\star\| \le \|P_\star - P_\gamma\| + \frac{\gamma_\infty^4}{\gamma^2 - \gamma_\infty^2}, \tag{13}$$

and thus our task reduces to bounding $\|P_\gamma - P_\star\|$, the gap between solutions to the γ-adversarial and nominal DAREs.

To bound the norm difference of DARE solutions, we show that the closed-loop dynamics under the adversary δ_t can be expressed as perturbations to the nominal system matrices. From Lemma 2.1, we have that for a noiseless γ-adversarial LQR instance, the adversary can be represented as $\delta_t = \Delta_\gamma(Ax_t + Bu_t)$, so the dynamics may be written

$$x_{t+1} = Ax_t + Bu_t + \delta_t = \hat{A}x_t + \hat{B}u_t,$$

where $\hat{A} := (I + \Delta_\gamma)A$ and $\hat{B} := (I + \Delta_\gamma)B$. It is now possible to bound the gaps $||\hat{A} - A||$, $||\hat{B} - B||$ between the system parameters (A, B) of the nominal system and the parameters $(\hat{A}, \hat{B})$ of the adversarially perturbed system in terms of γ. Bounds between the system parameters allow us to leverage existing tools for DARE perturbation analysis such as those presented in [36]. Before we present the upper bound, we introduce some notation that arises from the DARE perturbation bounds. We define the condition number on the LQR cost as

$$\kappa(Q, R) := \frac{\max\{\sigma_{\max}(Q), \sigma_{\max}(R)\}}{\min\{\sigma_{\min}(Q), \sigma_{\min}(R)\}},$$

and introduce the constants

$$\begin{aligned} \tau(A, \rho) &:= \sup\left\{\left\|A^k\right\|\rho^{-k} : k \ge 0\right\} \\ \beta &:= \max\left\{1, \frac{\gamma_\infty^2}{\gamma^2 - \gamma_\infty^2}\tau(A, \rho) + \rho\right\}, \end{aligned}$$

where $\rho > \rho(A)$. In short, $\tau(A,\rho)$ uniformly upper bounds the ratio between powers of the norm of A over ρ, and β upper bounds the effect of the adversary on the perturbed system matrix $\hat{A}$ in terms of the nominal system parameters. Since $\rho > \rho(A)$, these constants are guaranteed to exist by Gelfand's formula. Combining an upper bound on $\|P_\gamma - P_\star\|$ with equations (12) and (13) yields the following upper bound.

Theorem 3.1: Suppose $(A, B, Q^{1/2})$ is controllable and detectable. Let $\rho > \rho(A)$ and $\kappa(Q,R)$, $\tau(A,\rho)$, and β be the constants defined above. Furthermore, let ℓ be any natural number $1 \le \ell \le n$. For $\gamma > 0$ satisfying

$$\gamma^2 \ge \gamma_\infty^2 + \frac{3}{2}\ell^{3/2}\beta^{\ell-1}\sigma_{\min}(W_\ell(A,B))^{-1/2}\tau(A,\rho)^2 \cdot (\|B\|+1)\max\{\|A\|,\|B\|\}\gamma_\infty^2,$$

we have

$$\begin{aligned}&\mathrm{NC}(K_\gamma) - \mathrm{NC}(K_\star)\\ &\le O(1)\sigma_w^2\Big(\frac{\gamma_\infty^4}{\gamma^2-\gamma_\infty^2}\Big)^2 m\ell^5\beta^{4(\ell-1)}\|A+BK_\star\|^2\tau(A,\rho)^6\\ &\cdot\Big(1+\sigma_{\min}(W_\ell(A,B))^{-1/2}\Big)^2\|W_\infty(A+BK_\gamma,I)\|\\ &\cdot\frac{\|R+B^\top P_\star B\|}{\sigma_{\min}(R+B^\top P_\star B)^2}\kappa(Q,R)^2\|B\|^2(\|B\|+1)^4\|P_\star\|^2.\end{aligned}$$

As $\gamma \to \infty$, the upper bound decays to 0, since the adversarial controller converges to the nominal controller. However, the steepness of this cost gap is affected by system properties such as the minimum singular value of the ℓ-step controllability gramian, where poor controllability causes the upper bound to increase. We explicitly show the dependence of the bound on ℓ to highlight the fact that if the nominal system is controllable quickly–that is, we can choose $\ell \ll n$ such that the minimum singular value of the ℓ-step controllability gramian is large–then the upper bound improves. On the other hand, the potential exponential dependence on state dimension through β cannot be relaxed in general; consider chained-integrator systems as an example, where $\ell = \Theta(n)$ is required for the system to be controllable. We note that in contrast to the perturbation gap requirements in [36] or [38], our condition on the perturbation gap via lower bounds on γ are much less stringent. The lower bound requirement on γ arises here solely to guarantee the controllability of the adversarially perturbed system $(\hat{A}, \hat{B})$ is on the same order as that of the nominal system (A, B).

B. Lower Bound

The lower bound that we get from applying Lemma 3.1 is as follows:

$$\|K_\gamma - K_\star\| \ge \frac{\|B^\top(M_\gamma - P_\star)(A+BK_\star)\|}{\|R+BM_\gamma B\|} \ge \frac{\|B^\top(M_\gamma - P_\star)\|\,\sigma_{\min}(A+BK_\star)}{\|R+BM_\gamma B\|}.$$

Next, we add and subtract a particular DARE solution to the $M_\gamma - P_\star$ term in the above bound. Specifically, for $\gamma \ge \tilde{\gamma}_\infty$, we let $\tilde{P}_\gamma = \mathtt{DARE}(A+BK_\star, I, Q, -\gamma^2 I)$, and note that $x^\top\tilde{P}_\gamma x$ represents the cost of applying controller $K_\star$ in the noiseless adversarial setting at level γ starting from state x. The above bound becomes

$$\|K_\gamma - K_\star\| \ge \frac{\left\|B^\top(M_\gamma - \tilde{P}_\gamma + \tilde{P}_\gamma - P_\star)\right\|\sigma_{\min}(A+BK_\star)}{\|R+BM_\gamma B\|}. \quad (14)$$

From the definition (9) of M_γ and the reverse triangle inequality, we can lower bound the term $\left\|B^\top(M_\gamma - \tilde{P}_\gamma + \tilde{P}_\gamma - P_\star)\right\|$ in the above equation by

$$\left\|B^\top(P_\gamma(\gamma^2 I - P_\gamma)^{-1}P_\gamma + \tilde{P}_\gamma - P_\star)\right\| - \left\|B^\top(P_\gamma - \tilde{P}_\gamma)\right\|. \quad (15)$$

The first term may be lower bounded by expressing $\tilde{P}_\gamma - P_\star$ as the solution to a Lyapunov equation, while the second term may be bounded by leveraging the results of [37] to bound the cost gap between applying $K_\star$ and K_γ in the noiseless adversarial setting in terms of $\|K_\gamma - K_\star\|$. This is shown in the following lemma.

Lemma 3.2: Under the assumption $(A, B, Q^{1/2})$ is stabilizable and detectable, and for $\gamma \ge \tilde{\gamma}_\infty$, we have the following bounds:

$$\tilde{P}_\gamma - P_\gamma \preceq \|K_\gamma - K_\star\|^2\left\|R+B^\top M_\gamma B\right\| \cdot W_\infty(\tilde{A}+\tilde{B}K_\star, I)$$

and

$$P_\gamma(\gamma^2 I - P_\gamma)^{-1}P_\gamma + \tilde{P}_\gamma - P_\star \succeq \frac{\sigma_{\min}(P_\star)^2}{\gamma^2 - \sigma_{\min}(P_\star)}(W_\infty(A+BK_\star, I))$$

Combined with Equations (14) and (15), the above lemma gives rise to the following theorem.

Theorem 3.2: Define $\tilde{\Delta}_\gamma := I + (\gamma^2 I - \tilde{P}_\gamma)^{-1}\tilde{P}_\gamma$. Suppose $(A, B, Q^{1/2})$ is stabilizable and detectable. For $\gamma \ge \tilde{\gamma}_\infty$, and

$$\gamma^2 \ge \sigma_{\min}(P_\star) + \frac{1}{2}\sigma_{\min}(P_\star)^2\frac{\left\|B^\top W_\infty(A+BK_\star, I)\right\|}{\|R+B^\top M_\gamma B\|} \cdot \left\|B^\top W_\infty\Big(\tilde{\Delta}_\gamma(A+BK_\star), I\Big)\right\|\sigma_{\min}(A+BK_\star)^2,$$

we have

$$\mathrm{NC}(K_\gamma) - \mathrm{NC}(K_\star) \ge \frac{\sigma_w^2}{2}\left(\frac{\sigma_{\min}(P_\star)^2}{\gamma^2 - \sigma_{\min}(P_\star)}\right)^2\frac{\sigma_{\min}(R+B^\top P_\star B)}{\|R+B^\top M_\gamma B\|^2} \cdot \left\|B^\top W_\infty(A+BK_\star, I)\right\|^2\sigma_{\min}(A+BK_\star)^2.$$

The requirement that $\gamma \ge \tilde{\gamma}_\infty$ in the above theorem is due to the fact that the nominal controller must be stabilizing in the adversarial setting for the bounds to apply. The additional requirement on γ is present because the term $\left\|B^\top(P_\gamma - \tilde{P}_\gamma)\right\|$ in (15) is upper bounded in terms of $\|K_\gamma - K_\star\|^2$. If γ becomes too small (meaning $\|K_\gamma - K_\star\|$ is large), the bound becomes vacuous.

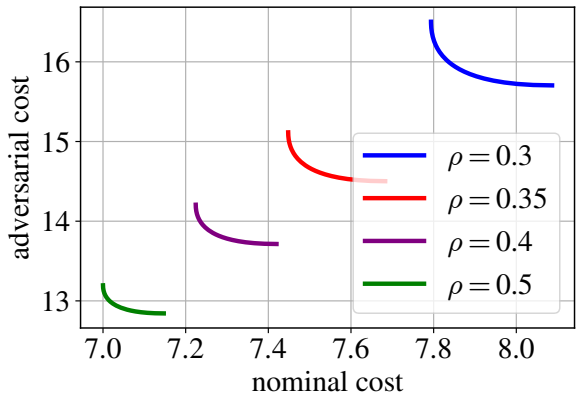

(a) Tradeoff curves

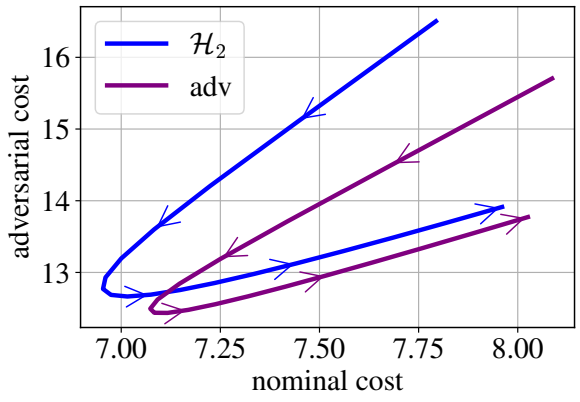

(b) Tradeoff envelope

Fig. 1. (a) For each value of the system parameter ρ, the tradeoff curves are generated by synthesizing adversarially robust controllers for adversarial levels ε ranging from $[0, 0.1]$, and then evaluating both their nominal cost, and their adversarial cost at level $\varepsilon = 0.1$. (b) We plot the nominal vs. adversarial performance (at level $\varepsilon = 0.1$) of the LQR controller and the adversarially robust controller at level $\varepsilon = 0.1$, as the ρ ranges from 0.3 to 1.2. The value of ρ increases in the directions of the arrows.

Keeping the nominal system fixed, both the upper and lower bounds decay at a rate γ^{-4}. We note that instead of the ℓ-step controllability gramian that manifests in the upper bound, we have instead the system parameter $W_\infty(A + BK_\star, I)$, which is the controllability gramian of the closed loop system under controller $K_\star$ with disturbances as inputs. That is, a large $\left\|B^\top W_\infty(A + BK_\star, I)\right\|$ implies that the nominal closed-loop system is quantifiably more controllable by the disturbance input, hence more susceptible to adversarial disturbances of fixed energy.

To apply the upper and lower bounds to the adversarial setting with fixed adversarial budget ε, one may find the optimal corresponding γ via Algorithm 1.

IV. Numerical Experiments

We now empirically study the trends suggested by our adversarially robust controller synthesis and the subsequent performance-robustness tradeoff bounds.

a) Impact of Controllability: To illustrate the dependence of the tradeoff severity on system controllability, consider the 2D integrator system defined by

$$(A, B, Q, R, \Sigma_w) = \left(\begin{bmatrix} 1 & \rho \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, I, I, I\right),$$

where we vary controllability via the parameter $\rho > 0$. When ρ is small, the system has poor controllability, and as ρ increases, controllability increases. In Figure 1a, we consider the tradeoff curves traced out by fixing ρ, then evaluating the nominal and adversarial ($\varepsilon = 0.1$) costs of adversarially robust controllers with budget ε varying between $[0, 0.1]$. We observe that as controllability decreases, the tradeoff curves shift upward and also widen. This corroborates the trend described in Theorem 3.1, where we show the bound on the nominal cost gap between adversarially robust and nominal controllers (i.e. width of the tradeoff curve) grows larger as controllability decreases. This trend is further illustrated in Figure 1b, where we plot the nominal and adversarial ($\varepsilon = 0.1$) costs attained by the $\mathcal{H}_2$ and adversarially robust ($\varepsilon = 0.1$) controllers as a function of $\rho \in [0.3, 1.2]$. We observe that for small ρ, the system has poor controllability, hence the distance between the controller costs is large. As controllability increases, this gap decreases monotonically. Note that the costs do not monotonically improve as ρ increases after some point, as the amplification of disturbances from the integrator outstrips the benefits of better controllability. Note however that in this regime, the gap between the nominal and adversarial controllers remains approximately constant, as predicted by Theorem 3.1.

b) Performance of Adversarially Robust Control: We compare the performance of the adversarially robust LQR controller against $\mathcal{H}_2$, $\mathcal{H}_\infty$, and mixed $\mathcal{H}_2/\mathcal{H}_\infty$ control for the longitudinal flight control for the linearized Boeing 747 system (see [39] for further details) defined by

$$A = \begin{bmatrix} 0.99 & 0.03 & -.02 & -.32 \\ 0.01 & 0.47 & 4.7 & .00 \\ .02 & -.06 & .40 & .00 \\ 0.01 & -.04 & .72 & .99 \end{bmatrix}, \; B = \begin{bmatrix} .01 & .99 \\ -3.44 & 1.66 \\ -.83 & .44 \\ -.47 & .25 \end{bmatrix},$$

with cost matrices $Q, R = I$ under different disturbances w_t:

(a) $w_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I)$.
(b) w_t set as worst case disturbances with power bounded by 1.
(c) $w_t \sim \mathcal{N}(\sin(0.01t), I)$.
(d) w_t given by i.i.d. $\mathcal{N}(0, I)$ Gaussian noise plus the worst case adversarial perturbation with budget $\varepsilon = 0.5$.

The initial condition is set to 0 in all cases. The optimal adversarially robust controller is generated by running Algorithm 1 with $\varepsilon = 0.5$, $Q = I, R = I$ and $\Sigma_w = I$ and the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controller is generated by fixing a $\mathcal{H}_\infty$ norm bound of $\gamma = 1000$, and approximately minimizing the $\mathcal{H}_2$ norm via the approach outlined in [40]. The $\mathcal{H}_\infty$ norm bound γ for mixed $\mathcal{H}_2/\mathcal{H}_\infty$ is chosen to achieve performance similar to the adversarially robust controller in the adversarial setting. In Figure 2a, we observe that in the zero-mean Gaussian setting, the $\mathcal{H}_2$ controller performs best as expected; however, we note that the adversarially

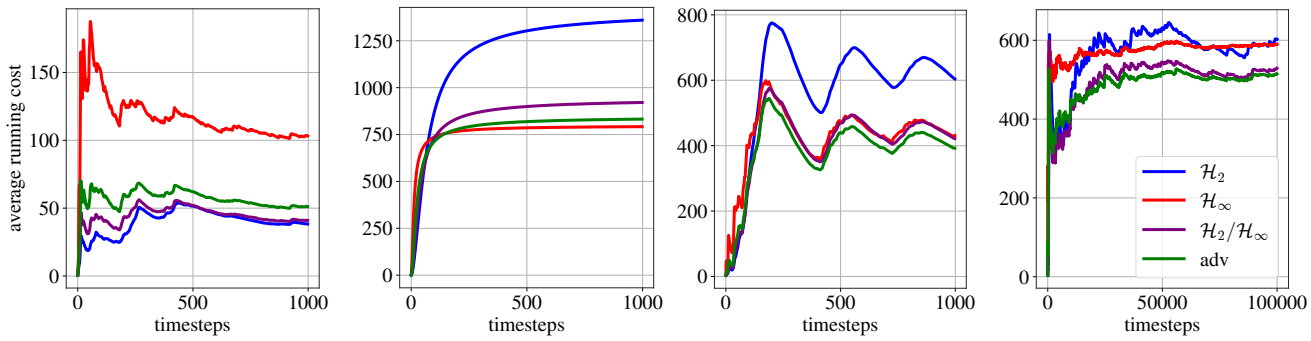

(a) Gaussian noise (b) Worst-case power bounded (c) Varying mean Gaussian (d) Adversarial setting

Fig. 2. Simulations of the performance of $\mathcal{H}_2$, $\mathcal{H}_\infty$, mixed $\mathcal{H}_2/\mathcal{H}_\infty$, and adversarially robust controllers on a linearized Boeing longitudinal flight control task. The average running cost at time t is computed as $\frac{1}{t+1}\sum_{k=0}^{t} x_k^\top Q x_k + u_k^\top R u_k$.

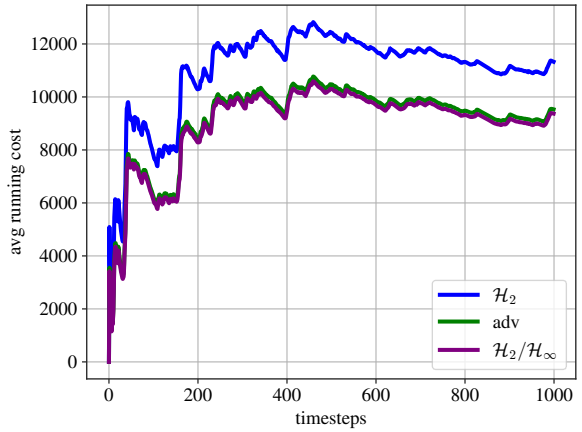

Fig. 3. Simulations of the performance of the $\mathcal{H}_2$, mixed $\mathcal{H}_2/\mathcal{H}_\infty$ ($\gamma = 10^6$), and adversarially robust ($\varepsilon = 10^{-4}$) controllers on an inverted pendulum stabilization task. The performance of $\mathcal{H}_\infty$ is not depicted due to being out of frame.

robust controller has similar performance. On worst-case power bounded disturbances in Figure 2b, $\mathcal{H}_\infty$ performs best as expected, and is closely followed by the adversarially robust controller. Interestingly, the adversarially robust controller outperforms all other controllers on the varying-mean Gaussian disturbance (Figure 2c), which is an instance of a disturbance composed of both zero-mean stochastic and deterministic components. The adversarially robust controller performs best in the adversarial setting (Figure 2d), as expected.

c) Beyond Linear Systems: For our final experiment, we demonstrate that adversarially robust control can perform favorably in non-linear systems. We compare the performance of the adversarially robust controller versus $\mathcal{H}_2$, $\mathcal{H}_\infty$, and mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controllers on an inverted pendulum system with the dynamics:

$$\ddot{\theta} = \frac{-mg\sin(\theta) - kl\dot{\theta} + u}{ml},$$

where we obtain a 2-D state space representation by setting $x := \begin{bmatrix}\theta & \dot{\theta}\end{bmatrix}^\top$. In the subsequent experiments, we set $m, k, l = 1$, $g = -9.81$, and the sampling time as $dt = 0.02$. The various controllers are computed using the linearized state matrices at the upright equilibrium point. The disturbance is generated as $w_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 2I)$, and the LQR cost matrices are set to $Q = I$, $R = I$. In Figure 3, we plot the performance of the various controllers on the inverted pendulum stabilization task. The adversarial budget is set to $\varepsilon = 10^{-4}$ for the adversarially robust controller to achieve approximately maximal performance. Note that this adversarial budget is quite small, as the true pendulum behavior is very similar to the linearized system near the upright equilibrium. Similarly the $\mathcal{H}_\infty$ norm bound γ for the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controller is chosen to maximize performance via bisection.

We observe that the adversarially robust and mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controllers perform similarly despite the adversarially robust controller being simpler to compute and implement, and notably both perform better than the $\mathcal{H}_2$ controller in this setting, despite the disturbance sequence being i.i.d. zero-mean Gaussian. This suggests that a small amount of robustness can impact the performance of a controller significantly by encompassing model errors, in this case those arising from using the linearized system to construct the controllers. At the same time, a fully robust controller is far too conservative to this end; the $\mathcal{H}_\infty$ controller is not pictured in Figure 3 due to having performance many times worse than $\mathcal{H}_2$.

V. Conclusion

We proposed an adversarially robust LQ control problem, and demonstrated that the optimal solution to this problem is given by a central static suboptimal $\mathcal{H}_\infty$ controller. An interesting aspect of this solution is that unlike pure $\mathcal{H}_2$ controllers, the adversarially robust controller depends upon the noise statistics. Experiments show that the adversarially robust controller performs similarly to mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controllers on a simple linear system, and can beat out both $\mathcal{H}_2$ and $\mathcal{H}_\infty$ simultaneously on disturbances that involve both stochastic and deterministic components, or model error arising from linearization.

We used the adversarially robust control problem as a means to study performance-robustness tradeoffs in control. In particular, we derived quantitative upper and lower bounds

on the performance gap between the nominal controller and the adversarially robust controller. The bounds show that systems with uniformly good controllability will have small performance-robustness tradeoffs, while those with a highly controllable mode in the closed-loop nominal system (viewing disturbances as inputs) will have a large performance-robustness tradeoff. These trends are corroborated by experiments on a simple linear system by tracing out tradeoff curves. Directions for future work include the extension of the problem setting considered here to the output feedback setting, and considering how other adversarial training techniques can be translated to robust controller synthesis and analysis.

ACKNOWLEDGEMENTS

Bruce D. Lee is supported by the DoD through the National Defense Science & Engineering Graduate Fellowship Program. The research of Hamed Hassani is supported by NSF Grants 1837253, 1943064, 1934876, AFOSR Grant FA9550-20-1-0111, and DCIST-CRA. Nikolai Matni is supported by NSF awards CPS-2038873, CAREER award ECCS-2045834, and a Google Research Scholar award.

REFERENCES

[1] J. Doyle, "Guaranteed margins for lqg regulators," *IEEE Transactions on Automatic Control*, vol. 23, no. 4, pp. 756–757, 1978.

[2] K. Zhou, J. Doyle, and K. Glover, *Robust and Optimal Control*, ser. Feher/Prentice Hall Digital and. Prentice Hall, 1996.

[3] M. A. Dahleh and I. J. Diaz-Bobillo, *Control of uncertain systems: a linear programming approach*. Prentice-Hall, Inc., 1994.

[4] C. Scherer, "Mixed $\mathcal{H}_2/\mathcal{H}_\infty$ control," in *Trends in Control*, A. Isidori, Ed. London: Springer London, 1995, pp. 173–216.

[5] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Machine Learning and Knowledge Discovery in Databases*. Springer Berlin Heidelberg, 2013, pp. 387–402.

[6] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[7] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.

[8] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *2017 IEEE symposium on security and privacy*. IEEE, 2017, pp. 39–57.

[9] T. Basar, $\mathcal{H}_\infty$ *Optimal Control and Related Minimax Design Problems*. Birkhäuser, 1991.

[10] J. Doyle, K. Zhou, and B. Bodenheimer, "Optimal control with mixed $\mathcal{H}_2$ and $\mathcal{H}_\infty$ performance objectives," in *1989 American Control Conference*, 1989, pp. 2065–2070.

[11] B. D. Lee, T. T. C. K. Zhang, H. Hassani, and N. Matni, "Performance-robustness tradeoffs in adversarially robust linear-quadratic control," 2022.

[12] F. Paganini, "Set descriptions of white noise and worst case induced norms," in *Proceedings of 32nd IEEE Conference on Decision and Control*, 1993, pp. 3658–3663 vol.4.

[13] K. Glover and J. C. Doyle, "State-space formulae for all stabilizing controllers that satisfy an $\mathcal{H}_\infty$-norm bound and relations to relations to risk sensitivity," *Systems & Control Letters*, vol. 11, no. 3, pp. 167–172, 1988.

[14] D. S. Bernstein and W. M. Haddad, "Lqg control with an $\mathcal{H}_\infty$ performance bound: a riccati equation approach," in *1988 American Control Conference*, 1988, pp. 796–802.

[15] M. A. Rotea and P. P. Khargonekar, "$\mathcal{H}_2$-optimal control with an $\mathcal{H}_\infty$-constraint the state feedback case," *Automatica*, vol. 27, no. 2, pp. 307–316, 1991.

[16] A. Tsiamis, D. S. Kalogerias, A. Ribeiro, and G. J. Pappas, "Linear quadratic control with risk constraints," 2021.

[17] M. P. Chapman and L. Lessard, "Toward a scalable upper bound for a cvar-lq problem," *IEEE Control Systems Letters*, vol. 6, pp. 920–925, 2022.

[18] G. Goel and B. Hassibi, "Regret-optimal control in dynamic environments," *arXiv preprint arXiv:2010.10473*, 2020.

[19] ——, "Competitive control," *arXiv preprint arXiv:2107.13657*, 2021.

[20] E. Hazan, S. Kakade, and K. Singh, "The nonstochastic control problem," in *ALT*. PMLR, 2020, pp. 408–421.

[21] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, "Theoretically principled trade-off between robustness and accuracy," *arXiv preprint arXiv:1901.08573*, 2019.

[22] A. Javanmard, M. Soltanolkotabi, and H. Hassani, "Precise tradeoffs in adversarial training for linear regression," in *Conference on Learning Theory*. PMLR, 2020, pp. 2034–2078.

[23] H. Hassani and A. Javanmard, "The curse of overparametrization in adversarial training: Precise analysis of robust generalization for random features regression," *arXiv preprint arXiv:2201.05149*, 2022.

[24] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, "Robustness may be at odds with accuracy," *arXiv preprint arXiv:1805.12152*, 2018.

[25] E. Dohmatob, "Generalized no free lunch theorem for adversarial robustness," in *International Conference on Machine Learning*. PMLR, 2019, pp. 1646–1654.

[26] D. Yin, R. Kannan, and P. Bartlett, "Rademacher complexity for adversarially robust generalization," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7085–7094.

[27] A. A. Al Makdah, V. Katewa, and F. Pasqualetti, "Accuracy prevents robustness in perception-based control," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 3940–3946.

[28] B. D. Lee, T. T. Zhang, H. Hassani, and N. Matni, "Adversarial tradeoffs in linear inverse problems and robust state estimation," *arXiv preprint arXiv:2111.08864*, 2021.

[29] T. T. Zhang, S. Tu, N. M. Boffi, J.-J. E. Slotine, and N. Matni, "Adversarially robust stability certificates can be sample-efficient," *arXiv preprint arXiv:2112.10690*, 2021.

[30] A. Havens, D. Keivan, P. Seiler, G. Dullerud, and B. Hu, "Revisiting pgd attacks for stability analysis of large-scale nonlinear systems and perception-based control," *arXiv preprint arXiv:2201.00801*, 2022.

[31] A. Pattanaik, Z. Tang, S. Liu, G. Bommannan, and G. Chowdhary, "Robust deep reinforcement learning with adversarial attacks," *arXiv preprint arXiv:1712.03632*, 2017.

[32] K. L. Tan, Y. Esfandiari, X. Y. Lee, S. Sarkar *et al.*, "Robustifying reinforcement learning agents via action space adversarial training," in *2020 American control conference (ACC)*. IEEE, 2020, pp. 3959–3964.

[33] A. Mandlekar, Y. Zhu, A. Garg, L. Fei-Fei, and S. Savarese, "Adversarially robust policy learning: Active construction of physically-plausible perturbations," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3932–3939.

[34] S. Kuutti, S. Fallah, and R. Bowden, "Arc: Adversarially robust control policies for autonomous vehicles," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 522–529.

[35] B. Hassibi, T. Kailath, and A. H. Sayed, *Indefinite-quadratic estimation and control: a unified approach to $\mathcal{H}_2$ and $\mathcal{H}_\infty$ theories*. SIAM studies in applied and numerical mathematics, 1999.

[36] H. Mania, S. Tu, and B. Recht, "Certainty equivalence is efficient for linear quadratic control," *arXiv preprint arXiv:1902.07826*, 2019.

[37] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, "Global convergence of policy gradient methods for the linear quadratic regulator," in *International Conference on Machine Learning*, vol. 80. PMLR, 10–15 Jul 2018, pp. 1467–1476.

[38] M. M. Konstantinov, P. H. Petkov, and N. D. Christov, "Perturbation analysis of the discrete riccati equation," *Kybernetika*, vol. 29, no. 1, pp. 18–29, 1993.

[39] J. Hong, N. Moehle, and S. Boyd, "Lecture notes in "introduction to matrix methods"," 2021.

[40] M. de Oliveira, J. Geromel, and J. Bernussou, "An lmi optimization approach to multiobjective controller design for discrete-time systems," in *Proceedings of the 38th IEEE Conference on Decision and Control*, vol. 4, 1999, pp. 3611–3616 vol.4.