

Data-driven distributionally robust LQR with multiplicative noise

Peter Coppens

PETER.COPPENS@KULEUVEN.BE

Mathijs Schuurmans

MATHIJS.SCHUURMANS@KULEUVEN.BE

Panagiotis Patrinos

PANOS.PATRINOS@KULEUVEN.BE

STADIUS, Department of Electrical Engineering (ESAT), KU Leuven, Leuven, Belgium.

Editors: A. Bayen, A. Jadbabaie, G. J. Pappas, P. Parrilo, B. Recht, C. Tomlin, M. Zeilinger

Abstract

We present a data-driven method for solving the linear quadratic regulator problem for systems with multiplicative disturbances, the distribution of which is only known through sample estimates. We adopt a distributionally robust approach to cast the controller synthesis problem as semidefinite programs. Using results from high dimensional statistics, the proposed methodology ensures that their solution provides mean-square stabilizing controllers with high probability even for low sample sizes. As the sample size increases the closed-loop cost approaches that of the optimal controller produced when the distribution is known. We demonstrate the practical applicability and performance of the method through a numerical experiment.

Keywords: data driven, distributionally robust, linear quadratic regulation, multiplicative noise, stochastic optimal control

1. Introduction

We will develop controllers for linear systems with time-varying parametric uncertainty, which may cover a wide range of system classes extensively studied in the literature. For example, we obtain Linear Parameter Varying (LPV) systems when the disturbance is observable at each time step (Wu et al., 1996; Byrnes, 1979), Linear Difference Inclusions (LDIs) when it is unknown but norm-bounded (Boyd et al., 1994) and stochastic systems with multiplicative noise when it varies stochastically (Wonham, 1967).

In many practical applications, however, the distribution of the disturbance is not known. These traditional control approaches either make a boundedness assumption on the disturbance or on its moments, which allows for a fully robust approach (El Ghaoui, 1995). Such approaches, however, disregard any statistical information that may be obtained on the distribution of the disturbances. Our aim, instead, is to design linear controllers which use sampled data to improve performance over fully robust approaches, while inheriting many of the system-theoretical guarantees of a robust control strategy. To this end, we adopt a distributionally robust (DR) approach (Dupačová, 1987; Delage and Ye, 2010) towards solving the infinite-horizon *Linear Quadratic Regulator* (LQR) problem, where we minimize the expected cost for the worst-case distribution in a so-called *ambiguity set* computed based on the available data such that it contains the true distribution with high probability. Such a DR approach addresses most of the difficulties associated with learning automatically, since the ambiguity set directly models the uncertainty in the sample-based estimates against which the controllers will be robust. Similar techniques were recently studied in Schuurmans et al. (2019)

for stochastic jump linear systems and in [Dean et al. \(2019\)](#) for deterministic systems, where the system matrices A and B are learned from data.

The work of [Gravell et al. \(2019\)](#) also deals with learning control of linear systems with multiplicative noise, albeit from a different perspective. They employ a policy gradient algorithm, which requires the initial guess for the control gain to be stabilizing. By contrast, obtaining such a controller is the main goal of our approach.

Our main contributions are summarized as follows. Leveraging recent results from high dimensional statistics, we provide practical high-probability confidence bounds for the ambiguity sets, which depend only on known quantities (Section 3). We then extend the solution of the (*nominal*) infinite horizon LQR problem with known distribution to related DR counterparts which account for the ambiguity on the disturbance distribution. Whenever the mean of the disturbance is known, we show that the DR problem is equivalent to a semidefinite program (SDP) which has the same form as the nominal one. Next, we extend the formulation to the setting in which both the mean and the covariance are only known to lie in an ellipsoidal ambiguity set (Section 4.2), for which we can only approximate the optimal controller.

Due to space limitations, only sketches of the proofs are provided here. We refer the reader to the technical report [Coppens et al. \(2019\)](#) for the full versions of the proofs.

1.1. Notation

Let \mathbb{R} denote the reals and \mathbb{N} the naturals. For symmetric matrices P, Q we write $P \succ Q$ ($P \succeq Q$) to signify that $P - Q$ is positive (semi)definite and denote by \otimes the Kronecker product. We assume that all random variables are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with Ω the sample space, \mathcal{F} its associated σ -algebra and \mathbb{P} the probability measure. Let $y : \Omega \rightarrow \mathbb{R}^n$ be a random vector defined on $(\Omega, \mathcal{F}, \mathbb{P})$. With some abuse of notation we will write $y \in \mathbb{R}^n$ to state the dimension of this random vector. Let \mathbb{P}_y denote the distribution of y , *i.e.*, $\mathbb{P}_y(A) = \mathbb{P}[y \in A]$. Then, a trajectory $\{y_i\}_{i=1}^N$ of identically and independently distributed (*i.i.d.*) copies of y is defined by the distribution it induces. That is, for any $A_0, \dots, A_N \in \mathcal{F}$ we define $\mathbb{P}_{y_0, \dots, y_N}(A_0 \times \dots \times A_N) := \mathbb{P}[y_0 \in A_0 \wedge \dots \wedge y_N \in A_N] = \prod_{i=0}^N \mathbb{P}_y(A_i)$. This definition can be extended to infinite trajectories $\{y_i\}_{i \in \mathbb{N}}$ by Kolmogorov's existence theorem [Billingsley \(1995\)](#). We will write the expectation operator as \mathbb{E} . Let \mathcal{M} denote the set of probability measures defined on $(\mathbb{R}^{n_w}, \mathcal{B})$, with \mathcal{B} the Borel σ -algebra of \mathbb{R}^{n_w} .

2. Problem statement

Consider the stochastic discrete-time system with input- and state-multiplicative noise given by:

$$x_{k+1} = A(w_k)x_k + B(w_k)u_k \tag{1}$$

$$\text{with } A(w) := A_0 + \sum_{i=1}^{n_w} w^{(i)} A_i, \quad B(w) := B_0 + \sum_{i=1}^{n_w} w^{(i)} B_i,$$

where at each time k , $x_k \in \mathbb{R}^{n_x}$ denotes the state, $u_k \in \mathbb{R}^{n_u}$ the input and $w_k \in \mathbb{R}^{n_w}$ an *i.i.d.* copy of a square integrable random variable w distributed according to \mathbb{P}_w . Let $w^{(i)}$ denote the i -th element of vector w . We introduce the following shorthands: $\mathbf{A} := [A_1^\top \dots A_{n_w}^\top]^\top$, $\mathbf{B} := [B_1^\top \dots B_{n_w}^\top]^\top$, $\mathbf{A}_0 := [A_0^\top \ \mathbf{A}^\top]^\top$, $\mathbf{B}_0 := [B_0^\top \ \mathbf{B}^\top]^\top$. We also define $\Sigma_0 := \begin{bmatrix} 1 & \mu^\top \\ \mu & \Sigma + \mu\mu^\top \end{bmatrix}$, where $\mu := \mathbb{E}[w]$, $\Sigma := \mathbb{E}[(w - \mu)(w - \mu)^\top]$.

2.1. Nominal stochastic LQR problem and solution

The primary goal is to solve the following LQR problem:

$$\begin{aligned}
 & \underset{u_0, u_1, \dots}{\text{minimize}} && \mathbb{E} \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right] \\
 & \text{subj. to} && x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad k \in \mathbb{N} \\
 & && x_0 = \bar{x},
 \end{aligned} \tag{2}$$

where we assume that $Q \succ 0$ and $R \succ 0$. The solution of (2) will yield a controller that renders the closed-loop system exponentially stable in the mean square sense, which is defined as follows.

Definition 1 ((Exponential) Mean Square Stability) *We say that an autonomous system $x_{k+1} = A(w_k)x_k$ is mean square stable (m.s.s.) iff $\mathbb{E}[x_k^\top x_k] \rightarrow 0$ as $k \rightarrow \infty$. It is exponentially mean square stable (e.m.s.s.) iff there exists a pair of positive constants $\gamma \in (0, 1)$ and c such that $\mathbb{E}[x_k^\top x_k] \leq c\gamma^k \|x_0\|^2$ for all $k \in \mathbb{N}$ and for each $x_0 \in \mathbb{R}^{n_x}$.*

This property can be verified using the classical Lyapunov operator (Morozan, 1983):

Theorem 2 (Lyapunov stability) *For the autonomous system $x_{k+1} = A(w_k)x_k$ the following statements are then equivalent: (i) it is m.s.s., (ii) it is e.m.s.s., (iii) $\exists P \succ 0$:*

$$P - \mathbf{A}_0^\top (\Sigma_0 \otimes P) \mathbf{A}_0 \succ 0. \tag{3}$$

The LQR problem (2) has been studied for many variations of (1) (Morozan, 1983; Costa and Kubrusly, 1997). The following proposition is then similar to many classical results in literature:

Proposition 3 *Consider a system with dynamics (1) and the associated LQR problem (2). Assuming that (1) is mean square stabilizable, i.e., there exists a K and $P \succ 0$ such that (3) holds for the closed-loop system $x_{k+1} = (A(w_k) + B(w_k)K)x_k$, then the following statements hold.*

I *The optimal solution of (2) is given by $K_\infty = -(R + G(P_\infty))^{-1}H(P_\infty)$, with P_∞ the solution of the following Riccati equation:*

$$P_\infty = Q + F(P_\infty) - H(P_\infty)^\top (R + G(P_\infty))^{-1} H(P_\infty), \tag{4}$$

where $F(P) := \mathbf{A}_0^\top (\Sigma_0 \otimes P) \mathbf{A}_0$, $G(P) := \mathbf{B}_0^\top (\Sigma_0 \otimes P) \mathbf{B}_0$, $H(P) := \mathbf{B}_0^\top (\Sigma_0 \otimes P) \mathbf{A}_0$.

II *The controller K_∞ stabilizes (1) in the mean square sense.*

III *The solution of the Riccati equation is found by solving the following SDP:*

$$\begin{aligned}
 & \underset{P}{\text{minimize}} && -\text{Tr } P \\
 & \text{subj. to} && \begin{bmatrix} Q - P + F(P) & H(P)^\top \\ H(P) & R + G(P) \end{bmatrix} \succeq 0, \\
 & && P \succeq 0.
 \end{aligned} \tag{5}$$

Proof The proof of statement **I** and **II** follows that of (Morozan, 1983, Theorem 1) and the proof of **III** follows from Balakrishnan and Vandenberghe (2003) (details in Coppens et al. (2019)). ■

Notice that the optimal solution to the LQR problem depends solely on the first and second moment of the random disturbance. This motivates our choice for the parametric form of the ambiguity set (6) used in the data-driven LQR problem, which we state in the next section.

2.2. Data-driven stochastic LQR problem

Consider now the case where \mathbb{P}_w is not known a priori and only a finite set of offline *i.i.d.* samples $\{\hat{w}\}_{i=0}^{M-1}$ is available. For clarity, we add a hat to imply that a random variable depends on these samples. It is apparent that under such circumstances, it is only possible to solve (2) approximately. For most applications, however, it is crucial that the approximate solution remains stabilizing, which is not trivial. For instance, the *empirical approach*, where (2) is solved using $\hat{\mu} := \frac{1}{M} \sum_{i=0}^{M-1} \hat{w}_i$ and $\hat{\Sigma} := \frac{1}{M} \sum_{i=0}^{M-1} (\hat{w}_i - \hat{\mu})(\hat{w}_i - \hat{\mu})^\top$, does not guarantee stability. We will illustrate this with an example¹, which motivates the methodology presented in this paper.

Example 1 (Motivating example) Consider the following scalar system,

$$x_{k+1} = (0.75 + w_k)x_k + u_k,$$

where $x_k \in \mathbb{R}$, $u_k \in \mathbb{R}$ and $w_k \in \mathbb{R}$ are defined as before, but now w is assumed Gaussian with $\mathbb{E}[w] = 0$ and $\mathbb{E}[w^2] = \sigma^2 = 0.5$. The empirical variance is $\hat{\sigma}^2 = \frac{1}{M} \sum_{i=0}^{M-1} \hat{w}_i^2$, using *i.i.d.* samples $\{\hat{w}_i\}_{i=0}^{M-1}$. We develop an optimal LQR controller with a stage cost given by $qx_k^2 + ru_k^2 = x_k^2 + 10^4 u_k^2$. Using results from Proposition 3, one can derive that the optimal controller is approximately given by $\hat{K}^* \approx (1 - 0.75^2 - \hat{\sigma}^2)/0.75$ for $\hat{\sigma}^2 > 0.4375$, otherwise $\hat{K}^* \approx 0$. We also assume that $\hat{\sigma}^2 < 1$ (related to the uncertainty threshold principle of Athans et al. (1977)) since otherwise the problem is infeasible as the system is not stabilizable. We can show¹ that the controller \hat{K}^* is m.s.s. iff $(0.75 + \hat{K}^*)^2 + 0.5(\hat{K}^*)^2 < 1$, which holds iff $-1.4571 < \hat{K}^* < -0.0429$. This, in turn, is the case iff $0.4697 < \hat{\sigma}^2 < 1.5303$. Only the lower bound is critical since we assumed that $\hat{\sigma}^2 < 1$. As $\hat{\sigma}^2$ is a scaled χ^2 -distributed random variable with M degrees of freedom, the probability of $\hat{\sigma}^2 < 0.4697$ occurring for $M = 500$ is 0.1693. That is, the probability that the empirical approach provides an unstable closed-loop controller in this example is almost 17%.

From this example, it is clear that underestimation of the variance of the disturbance is directly related to the probability of failure of the controller. In order to take this into account, we introduce an arbitrarily-chosen confidence level $\beta \in [0, 1]$ and a corresponding ambiguity set $\hat{\mathcal{A}} : \Omega \rightrightarrows \mathcal{M}$, which represents the uncertainty of estimators $\hat{\mu}$ and $\hat{\Sigma}$. The size of \mathcal{A} is then determined such that $\mathbb{P}(\mathbb{P}_w \in \hat{\mathcal{A}}) \geq 1 - \beta$. In particular, we parametrize \mathcal{A} as first suggested by Delage and Ye (2010):

$$\hat{\mathcal{A}} := \left\{ \mathbb{P}_v \in \mathcal{M} \left| \begin{array}{l} (\mathbb{E}[v] - \hat{\mu})^\top \hat{\Sigma}^{-1} (\mathbb{E}[v] - \hat{\mu}) \leq r_\mu(\beta) \\ \mathbb{E}[(v - \mu)(v - \mu)^\top] \leq r_\Sigma(\beta) \hat{\Sigma} \end{array} \right. \right\}, \quad (6)$$

The values of $r_\mu(\beta)$ and $r_\Sigma(\beta)$ such that $\mathbb{P}(\mathbb{P}_w \in \hat{\mathcal{A}}) \geq 1 - \beta$ are derived in Section 3. In Section 4, the following DR counterpart of (2) is solved

$$\begin{aligned} & \underset{u_0, u_1, \dots}{\text{minimize}} && \max_{\mathbb{P}_v \in \hat{\mathcal{A}}} \mathbb{E} \left[\sum_{k=1}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right] \\ & \text{subj. to} && x_{k+1} = A(v_k)x_k + B(v_k)u_k, \quad k \in \mathbb{N} \\ & && x_0 = \bar{x}, \end{aligned} \quad (7)$$

where $\{v_k\}_{k \in \mathbb{N}}$ is a trajectory of *i.i.d.* copies of v . In doing so, we can finally establish *m.s.s.* of the data-driven controller with high probability, by virtue of the following generalization of Theorem 2 to the DR case.

1. Detailed derivations in Coppens et al. (2019).

Theorem 4 (DR Lyapunov stability) Consider the matrices $\{\hat{A}_i\}_{i=0}^{n_w}$, the set $\{w_k\}_{k \in \mathbb{N}}$ consisting of i.i.d. copies of a square integrable random vector w and the autonomous system $x_{k+1} = \hat{A}(w_k)x_k$, with $\hat{A}(w) = \sum_{i=1}^M \hat{A}_i w^{(i)}$. Say that we have an ambiguity set with $\mathbb{P}(\mathbb{P}_w \in \hat{\mathcal{A}}) \geq 1 - \beta$, with \mathbb{P}_w the true distribution of w . Then if there exists a $P \succ 0$ such that:

$$P - \max_{\mathbb{P}_v \in \hat{\mathcal{A}}} \mathbb{E}[A(v)^\top P A(v)] \succ 0, \quad (8)$$

the autonomous system is e.m.s.s. with probability at least $1 - \beta$.

Proof Full proof in [Coppens et al. \(2019\)](#). ■

3. Data-driven ambiguity set estimation

We now turn to the problem of estimating the parameters $r_\Sigma(\beta)$ and $r_\mu(\beta)$ involved in the definition of the ambiguity set (6), given that we have M i.i.d. draws from the true distribution. These parameters will be estimated under the following assumption on the disturbances.

Definition 5 (Sub-Gaussianity) A random variable y is sub-Gaussian with variance proxy σ^2 if $\mathbb{E}[y] = 0$ and its moment generating function satisfies

$$\mathbb{E}[\exp(\lambda y)] \leq \exp\left(\frac{\sigma^2 \lambda^2}{2}\right) \quad \forall \lambda \in \mathbb{R}. \quad (9)$$

We denote this by $y \sim \text{subG}(\sigma^2)$. We say that a random vector $\xi \in \mathbb{R}^{n_w}$ is sub-Gaussian, or $\xi \sim \text{subG}_{n_w}(\sigma^2)$, if $z^\top \xi \sim \text{subG}(\sigma^2)$, $\forall z \in \mathbb{R}^{n_w}$ with $\|z\|_2 = 1$.

Assumption 1 We assume that (i) w is square integrable; (ii) $\{w_k\}_{k \in \mathbb{N}}$ are i.i.d. copies of w ; (iii) $\Sigma \succ 0$; and (iv) $\Sigma^{-1/2}(w_k - \mu) \sim \text{subG}_{n_w}(\sigma^2)$ for some $\sigma \geq 1$.

Note that in the specific case of Gaussian disturbances, Assumption 1(iv) holds with $\sigma^2 = 1$, so no further prior knowledge on the distribution is required. Moreover, in this case the bound on the covariance obtained in Theorem 6 can be slightly improved ([Wainwright, 2019](#)). In the case of bounded disturbances, σ^2 can be estimated in a data-driven fashion ([Delage and Ye, 2010](#)).

For the moment, we restrict our attention to obtaining concentration inequalities for moment estimators of random vectors with zero mean and unit variance — hereafter referred to as *isotropic* random vectors. We will then convert these results into ambiguity sets of the form (6) using arguments from [Delage and Ye \(2010\)](#); [So \(2011\)](#). We begin by specializing the isotropic covariance bound, derived with constants in ([Hsu et al., 2012a](#), Lemma A.1.) based on a result by [Litvak et al. \(2005\)](#) and the isotropic mean bound by ([Hsu et al., 2012b](#), Theorem 2.1).

Theorem 6 (Isotropic covariance bound) Let $\xi \sim \text{subG}_{n_w}(\sigma^2)$ be a random vector, with $\mathbb{E}[\xi] = 0$, $\mathbb{E}[\xi \xi^\top] = I_{n_w}$. Let $\{\hat{\xi}_i\}_{i=0}^{M-1}$ be M independent copies of ξ and $\hat{I} := \frac{1}{M} \sum_{i=0}^{M-1} \hat{\xi}_i \hat{\xi}_i^\top$. Then

$$\mathbb{P}[\|\hat{I} - I_{n_w}\|_2 \leq t_\Sigma(\beta)] \geq 1 - \beta,$$

with

$$t_\Sigma(\beta) := \frac{\sigma^2}{1 - 2\epsilon} \left(\sqrt{\frac{32q(\beta, \epsilon, n_w)}{M}} + \frac{2q(\beta, \epsilon, n_w)}{M} \right), \quad (10)$$

where $\epsilon \in (0, 1/2)$ is chosen freely and $q(\beta, \epsilon, n_w) := n_w \log(1 + 1/\epsilon) + \log(2/\beta)$.

Theorem 7 (Isotropic mean bound) Let $\{\hat{\xi}_i\}_{i=0}^{M-1}$ be as defined in Theorem 6, and $\hat{\zeta} := \frac{1}{M} \sum_{i=0}^{M-1} \hat{\xi}_i$. Then

$$\mathbb{P}[\|\hat{\zeta}\|_2^2 \leq t_\mu(\beta)] \geq 1 - \beta.$$

where

$$t_\mu(\beta) := \frac{\sigma^2}{M} p(\beta, n_w),$$

with $p(\beta, n_w) := \left(n_w + 2\sqrt{n_w \log(1/\beta)} + 2 \log(1/\beta) \right)$.

By combining the bounds in Theorems 6 and 7, we readily obtain the following result.

Theorem 8 (Ambiguity set) Let $w \in \mathbb{R}^{n_w}$ be a sub-Gaussian random vector, with $\mathbb{E}[w] = \mu$, $\mathbb{E}[(w - \mu)(w - \mu)^\top] = \Sigma$ and $\xi = \Sigma^{-1/2}(w - \mu) \sim \text{subG}_{n_w}(\sigma^2)$. Let $\{\hat{w}_i\}_{i=0}^{M-1}$ be independent copies of w . Let $\hat{\mu} := \frac{1}{M} \sum_{i=0}^{M-1} \hat{w}_i$ and $\hat{\Sigma} := \frac{1}{M} \sum_{i=0}^{M-1} (\hat{w}_i - \hat{\mu})(\hat{w}_i - \hat{\mu})^\top$ denote the empirical estimators for the mean and the covariance matrix, respectively. Let $\epsilon, p(\beta, n_w), q(\beta, \epsilon, n_w), t_\Sigma(\beta/2)$ and $t_\mu(\beta/2)$ be as defined in Theorems 6 and 7. Provided that

$$M > \left(\frac{\sigma^2 \sqrt{32q(\beta/2, \epsilon, n_w)} + \sqrt{32\sigma^4 q(\beta/2, \epsilon, n_w) + 8\sigma^2(1-2\epsilon)q(\beta/2, \epsilon, n_w) + 4\sigma^2(1-2\epsilon)^2 p(\beta/2, n_w)}}{2(1-2\epsilon)} \right)^2, \quad (11)$$

then with probability at least $1 - \beta$,

$$(\hat{\mu} - \mu)^\top \hat{\Sigma}^{-1} (\hat{\mu} - \mu) \leq r_\mu(\beta), \quad \Sigma \preceq r_\Sigma(\beta) \hat{\Sigma},$$

with $r_\Sigma(\beta) := \frac{1}{1 - t_\mu(\beta/2) - t_\Sigma(\beta/2)}$ and $r_\mu(\beta) := \frac{t_\mu(\beta/2)}{1 - t_\mu(\beta/2) - t_\Sigma(\beta/2)}$.

Proof Apply the procedure of (Delage and Ye, 2010, Thm. 2) to the results of Theorem 6–7. ■

4. Distributionally Robust LQR

We will tackle the solution of (2) for the ambiguity set given in (6) in two stages. Firstly, we extend the result of Proposition 3 to the case where μ is known and Σ is estimated, i.e., $r_\mu(\beta) = 0$ and $r_\Sigma(\beta) > 0$. Secondly, we present the result where both the mean and the covariance are estimated.

4.1. Uncertain covariance

The case where the mean is known is interesting since we can still formulate an exact solution to (7). This will no longer be true for the full-uncertainty case (Section 4.2).

Proposition 9 Consider that $v \in \mathbb{R}^{n_w}$ is distributed according to an element of the set

$$\hat{\mathcal{A}}_\Sigma := \left\{ \mathbb{P}_v \in \mathcal{M} \mid \mathbb{E}[(v - \mu)(v - \mu)^\top] \preceq r_\Sigma(\beta) \hat{\Sigma}, \mathbb{E}[v] = \mu \right\}, \quad (12)$$

where $\mathbb{P}(\mathbb{P}_w \in \hat{\mathcal{A}}_\Sigma) \geq 1 - \beta$. Then applying Proposition 3 with $\Sigma = r_\Sigma(\beta) \hat{\Sigma}$ results in the optimal linear controller for (7), assuming that (1) is DR mean square stabilizable, i.e., there exists a K such that the DR Lyapunov decrease (8) holds for the closed-loop system $x_{k+1} = (A(w_k) + B(w_k)K)x_k$. The optimal controller is also mean square stabilizing for (1) with probability at least $1 - \beta$.

Proof The Bellman operator associated with (7) is of the same form as the one in the proof of Proposition 3, which makes it applicable (full proof is in (Coppens et al., 2019)). ■

4.2. Full uncertainty

We finally consider the more general case using the full ambiguity set $\hat{\mathcal{A}}$ given by (6). The general *min-max problem* (7) for such sets is computationally intractable, which is why an upper bound on the quadratic cost is minimized instead by employing results from robust control (Boyd et al., 1994; Kothare et al., 1996). A common approach is to assume that the value function can be written in the quadratic form $V(x) = x^\top P x$ for some $P \succ 0$ and to solve the optimization problem

$$\begin{aligned} & \underset{V(x)}{\text{minimize}} && \mathbb{E}[V(\tilde{x})] \\ & \text{subj. to} && V(x) \geq \min_u \left\{ x^\top Q x + u^\top R u + \max_{\mathbb{P}_v \in \hat{\mathcal{A}}} \mathbb{E}[V(A(v)x + B(v)u)] \right\}, \quad \forall x, \end{aligned} \quad (13)$$

where we introduced the random initial state $\tilde{x} \in \mathbb{R}^{n_x}$. The optimal cost of (13) then upper bounds the true LQR cost of (7) for a given value of \tilde{x} as proven in Kothare et al. (1996) for a similar setup. We can then write (13) as an SDP using the following theorem.

Theorem 10 *Let $\hat{\mathcal{A}}$ be an ambiguity set of the form (6). Then we can find an approximate solution of (13) for the system (1), assuming that the initial state is given by the random vector $\tilde{x} \in \mathbb{R}^{n_x}$ with $\mathbb{E}[\tilde{x}] = 0$ and $\mathbb{E}[\tilde{x}\tilde{x}^\top] = I_{n_x}$, by solving the following SDP.*

$$\begin{aligned} & \underset{W, V, S, L}{\text{maximize}} && \text{Tr } W \\ & \text{subj. to} && \begin{bmatrix} S & r_\mu(\beta)H_1^\top & r_\mu(\beta)H_2^\top & \cdots & r_\mu(\beta)H_{n_w}^\top \\ r_\mu(\beta)H_1 & L & & & \\ r_\mu(\beta)H_2 & & L & & \\ \vdots & & & \ddots & \\ r_\mu(\beta)H_{n_w} & & & & L \end{bmatrix} \succeq 0, \end{aligned} \quad (14a) \end{aligned}$$

$$\begin{aligned} & && \begin{bmatrix} W - \sqrt{2}S & (\mathbf{A}W + \mathbf{B}V)^\top & (\hat{\mathbf{A}}W + \hat{\mathbf{B}}V)^\top & W^\top Q^{\frac{1}{2}} & V^\top R^{\frac{1}{2}} \\ \mathbf{A}W + \mathbf{B}V & \hat{\Sigma}_{\text{dr}}^{-1} \otimes W & & & \\ \hat{\mathbf{A}}W + \hat{\mathbf{B}}V & & W - \sqrt{2}L & & \\ Q^{\frac{1}{2}}W & & & I_{n_x} & \\ R^{\frac{1}{2}}V & & & & I_{n_u} \end{bmatrix} \succeq 0, \end{aligned} \quad (14b) \end{aligned}$$

where $H_i = \sum_{j=1}^{n_w} [\hat{\Sigma}^{1/2}]_{ji} (A_j W + B_j V)$, $\hat{\mathbf{A}} = A(\hat{\mu})$, $\hat{\mathbf{B}} = B(\hat{\mu})$, $\hat{\Sigma}_{\text{dr}} = r_\Sigma(\beta) \hat{\Sigma}$. Let \hat{W} and \hat{V} denote the minimizers of (14). The corresponding linear controller $u = \hat{K}x$ with $\hat{K} = \hat{V}\hat{W}^{-1}$ then achieves an upper bound of the cost (13), given by $\mathbb{E}[\tilde{x}^\top \hat{P} \tilde{x}]$, where $\hat{P} = \hat{W}^{-1}$. Moreover, \hat{K} is mean-square stabilizing for (1) with probability at least $1 - \beta$.

Proof The proof specializes (Ben-Tal et al., 2000, Theorem 6.2.1) (see Coppens et al. (2019)). \blacksquare

Remark 11 *Two approximations are made in Theorem 10. First, (Ben-Tal et al., 2000, Theorem 6.2.1), loosens $r_\mu(\beta)$ by a factor of at most $\sqrt{n_w}$. Secondly, we minimize an upper-bound of the LQR cost instead of the cost itself. We can further decrease the closed-loop cost by using a receding horizon controller for a given \tilde{x} , which too can be formulated as an SDP (Kothare et al., 1996; Coppens et al., 2019). By contrast, the assumption in Theorem 10 ensures that the solution converges to the nominal one as $r_\mu(\beta)$ and $r_\Sigma(\beta)$ go to zero (Balakrishnan and Vandenberghe, 2003).*

Remark 12 *Invertibility of $r_\Sigma(\beta) \hat{\Sigma}$ can be guaranteed by using $r_\Sigma(\beta) \hat{\Sigma} + \lambda I$ instead of $r_\Sigma(\beta) \hat{\Sigma}$ for some small λ , which increasing the size of the ambiguity set, introducing additional conservatism.*




Figure 1: Relative suboptimality versus sample size for controllers. The colored area depicts the 0.3-confidence interval around the cost. The full line depicts the mean.

5. Numerical Experiment

We experimentally quantify the sample complexity of our approach, *i.e.*, how many samples are needed before the controller becomes equivalent to the nominal one based on the true Σ and μ instead of their data-driven estimates. Consider the double integrator model with matrices:

$$A_0 = \begin{bmatrix} 1 & T_s \\ 0 & 1 - 0.4T_s \end{bmatrix}, B_0 = \begin{bmatrix} 0 \\ T_s \end{bmatrix}, A_1 = \begin{bmatrix} 0 & 0 \\ 0 & -T_s \end{bmatrix}, A_2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, B_2 = \begin{bmatrix} 0 \\ T_s \end{bmatrix},$$

where we chose $T_s = 0.02$. The dynamics are then given by (1) with w_k an independent random sequence of Gaussian random vectors with covariance $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and mean $\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

To estimate the sample complexity, we determine controllers satisfying (7) for $Q = \begin{bmatrix} 10 & 0 \\ 0 & 1 \end{bmatrix}$ and $R = 0.01$. The parameters of the ambiguity set (6) are determined using Theorem 8 with $\beta = 0.05$ and $\epsilon = 1/30$. Since w_k are Gaussian, $\xi_k = \Sigma^{-1/2}(w_k - \mu)$ are sub-Gaussian with $\sigma^2 = 1$.

The simulation setup is as follows. We compare the nominal controller, the uncertain covariance controller (Proposition 9) and the full uncertainty controller (Theorem 10). We evaluate the expected closed-loop cost for $\bar{x} = [2 \ 2]^\top$ by solving the Lyapunov equation. We start with $M = 1000$ to satisfy (11). For each value of M we produce 30 realisations of both DR controllers. Figure 1 depicts confidence intervals for relative difference between closed-loop cost of the DR controllers and the nominal controller (*i.e.*, the relative suboptimality). The figure shows that both converge with a rate $\mathcal{O}(1/M)$ to the nominal cost even though Theorem 10 only solves (7) approximately.

6. Conclusion and future work

We studied the infinite horizon LQR problem for systems with multiplicative uncertainty on both the states and the inputs. We operate in the setting where the distributions are estimated from data. We show that using results from high-dimensional statistics, high-confidence ambiguity sets can be constructed, which allow us to formulate a DR counterpart to the stochastic optimal control problem as an SDP. As a result, stability of the closed-loop system can be guaranteed with high probability.

In future work, we aim to perform an in-depth analysis of the conservatism introduced by the proposed formulations. Furthermore, we plan to study extensions towards DR Kalman filtering.

Acknowledgments

This work was supported by the Ford KU Leuven Research Alliance; the Fonds Wetenschappelijk Onderzoek PhD grant 11E5520N and research projects G0A0920N, G086518N and G086318N; Research Council KU Leuven C1 project No. C14/18/068; Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS project no 30468160 (SeLMA).

References

- Michael Athans, Richard Ku, and Stanley B. Gershwin. The uncertainty threshold principle: Some fundamental limitations of optimal decision making under dynamic uncertainty. *IEEE Transactions on Automatic Control*, 22(3):491–495, June 1977.
- Venkataramanan Balakrishnan and Lieven Vandenberghe. Semidefinite programming duality and linear time-invariant systems. *IEEE Transactions on Automatic Control*, 48(1):30–41, 2003.
- Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. Robustness. In *Handbook of Semidefinite Programming*, pages 139–162. Springer US, Boston, MA, 2000.
- Patrick Billingsley. *Probability and measure*. Wiley series in probability and mathematical statistics. Wiley, New York, 3rd edition edition, 1995.
- Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.
- Christopher I. Byrnes. On the stabilizability of linear control systems depending on parameters. In *Proceedings of the IEEE Conference on Decision and Control*, volume 1, pages 233–236. IEEE, 1979.
- Peter Coppens, Mathijs Schuurmans, and Panagiotis Patrinos. Data-driven distributionally robust LQR with multiplicative noise. *arXiv preprint arXiv:1912.09990*, 2019.
- Oswaldo L.V. Costa and Carlos S. Kubrusly. Quadratic optimal control for discrete-time infinite-dimensional stochastic bilinear systems. *IMA Journal of Mathematical Control and Information*, 14(4):385–399, 1997.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the Sample Complexity of the Linear Quadratic Regulator. *Foundations of Computational Mathematics*, 2019.
- Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.
- Jitka Dupačová. The minimax approach to stochastic programming and an illustrative application. *Stochastics*, 20(1):73–88, 1987.
- Laurent El Ghaoui. State-feedback control of systems with multiplicative noise via linear matrix inequalities. *Systems and Control Letters*, 24(3):223–228, 1995.
- Benjamin Gravell, Peyman Mohajerin Esfahani, and Tyler Summers. Learning robust control for LQR systems with multiplicative noise via policy gradient. *arXiv preprint arXiv:1905.13547*, 2019.
- Daniel Hsu, Sham M. Kakade, and Tong Zhang. Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electronic Communications in Probability*, 17(52):1–13, 2012a.

- Daniel Hsu, Sham M. Kakade, and Tong Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17(52):1–6, 2012b.
- Mayuresh V. Kothare, Venkataramanan Balakrishnan, and Manfred Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10):1361–1379, 1996.
- Alexander E. Litvak, Alalin Pajor, M. Rudelson, and N. Tomczak-Jaegermann. Smallest singular value of random matrices and geometry of random polytopes. *Advances in Mathematics*, 195(2):491–523, 2005.
- Toader Morozan. Stabilization of some stochastic discrete-time control systems. *Stochastic Analysis and Applications*, 1(1):89–116, 1983.
- Mathijs Schuurmans, Pantelis Sopasakis, and Panagiotis Patrinos. Safe learning-based control of stochastic jump linear systems: a distributionally robust approach. In *2019 IEEE Conference on Decision and Control (CDC)*, pages 6498–6503. IEEE, 2019.
- Anthony Man Cho So. Moment inequalities for sums of random matrices and their applications in optimization. *Mathematical Programming*, 130(1):125–151, 2011.
- Martin J. Wainwright. *High-Dimensional Statistics*. Cambridge University Press, 2019.
- W. Murray Wonham. Optimal Stationary Control of a Linear System with State-Dependent Noise. *SIAM Journal on Control*, 5(3):486–500, 1967.
- Fen Wu, Xin Hua Yang, Andy Packard, and Greg Becker. Induced L_2 -norm control for LPV systems with bounded parameter variation rates. *International Journal of Robust and Nonlinear Control*, 6(9-10):983–998, 1996.