

 Open access • Proceedings Article • DOI:10.1109/CDC40024.2019.9029946

Safe Learning-Based Control of Stochastic Jump Linear Systems: a Distributionally Robust Approach — [Source link](#)

Mathijs Schuurmans, Pantelis Sopasakis, Panagiotis Patrinos

Institutions: Katholieke Universiteit Leuven, Queen's University Belfast

Published on: 01 Dec 2019 - Conference on Decision and Control

Topics: Probability distribution, Sample space, Stability (probability), Empirical distribution function and Sample size determination

Related papers:

- [Safe Learning-Based Control of Stochastic Jump Linear Systems: a Distributionally Robust Approach.](#)
- [Discrete-Time Markov Jump Linear Systems](#)
- [Dynamic Programming and Optimal Control](#)
- [Distributionally Constrained Stochastic Gradient Estimation Using Noisy Function Evaluations](#)
- [Regret Bounds for the Adaptive Control of Linear Quadratic Systems](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/safe-learning-based-control-of-stochastic-jump-linear-3ct7mh04me>

Safe Learning-Based Control of Stochastic Jump Linear Systems: a Distributionally Robust Approach

Mathijs Schuurmans[†], Pantelis Sopasakis[‡] and Panagiotis Patrinos[†]

Abstract— We consider the problem of designing control laws for stochastic jump linear systems where the disturbances are drawn randomly from a finite sample space according to an unknown distribution, which is estimated from a finite sample of i.i.d. observations. We adopt a distributionally robust approach to compute a mean-square stabilizing feedback gain with a given probability. The larger the sample size, the less conservative the controller, yet our methodology gives stability guarantees with high probability, for any number of samples. Using tools from statistical learning theory, we estimate confidence regions for the unknown probability distributions (ambiguity sets) which have the shape of total variation balls centered around the empirical distribution. We use these confidence regions in the design of appropriate distributionally robust controllers and show that the associated stability conditions can be cast as a tractable linear matrix inequality (LMI) by using conjugate duality. The resulting design procedure scales gracefully with the size of the probability space and the system dimensions. Through a numerical example, we illustrate the superior sample complexity of the proposed methodology over the stochastic approach.

I. INTRODUCTION

A. Background and motivation

The ever-decreasing costs of measuring, communicating and storing data have led to a variety of opportunities to apply learning-based and data-driven approaches in control [1], [2]. These opportunities are of particular interest for systems with inherent stochastic uncertainty, as data-driven methodologies may be used to reduce conservativeness in controller design, while retaining safety guarantees.

A natural way of addressing this trade-off is by adopting a *distributionally robust* approach [3], [4], which is gaining popularity in many fields including machine learning [5], [6] and control [7], [8]. It provides a framework which inherently accounts for uncertainty on probability estimates by generalizing two opposing approaches of *stochastic* and *robust* control [9]. Performance and safety guarantees of the former [10] require full knowledge of the underlying

probability distribution of involved random variables, which in practice is only available by approximation. The robust approach, on the other hand, aims at providing guarantees in the worst possible realization of the uncertain variables. This disregard for available statistical knowledge typically leads to overly conservative solutions or infeasibility. By contrast, the distributionally robust framework imposes robustness only with respect to a given set of probability distributions, often called *ambiguity sets*. The challenge is to appropriately design this ambiguity set in order to make a suitable trade-off between safety and performance.

In the past few years, the stochastic optimization community has proposed several methods for building ambiguity sets from data, and solving corresponding optimization problems [11], [12]. One popular approach is to estimate the unknown distribution (e.g., by the *empirical estimate*) and to construct the ambiguity set as the set of distributions within some statistical distance, such as the *Wasserstein distance* [8], [13]–[15] or *ϕ -divergences* [16], [17] from this estimate. In this paper, we follow this line of reasoning and restrict the considered class of ambiguity sets to be the ℓ_1 -norm ball around the empirical probability estimate. Many of the obtained results can however be extended to more general classes of convex ambiguity sets, given the appropriate modifications.

B. Main contributions

Firstly, we propose a data-driven, distributionally robust design methodology for synthesizing static feedback control gains for stochastic jump linear systems, which, for *any* finite sample size grants mean-square stability to the closed-loop system at a given confidence level.

Secondly, we propose a reformulation of the resulting stability conditions and approximate it by a tractable linear matrix inequality (LMI), which avoids enumerating the extreme points of the polytopic ℓ_1 -ambiguity set. We demonstrate the computational gains of this formulation and show that, in practice, the induced conservativeness is very limited.

C. Notation

Let I_n be the $n \times n$ identity matrix and let the sets of symmetric positive definite and positive semi-definite $n \times n$ matrices be denoted as \mathbb{S}_{++}^n and \mathbb{S}_+^n , respectively. We denote by \otimes the Kronecker product. For $x, y \in \mathbb{R}$, we define $\mathbf{1}_y(x)$ to be equal to 1 if $x = y$, and 0 otherwise. We denote the expectation operator by $\mathbb{E}[\cdot]$ and the probability simplex by $\Delta_k := \{p \in \mathbb{R}^k \mid p_i \geq 0, \sum_{i=1}^k p_i = 1\}$. We define $\mathbb{N}_{[a,b]} := \{i \in \mathbb{N} \mid a \leq i \leq b\}$. The spectral radius of a

[†]M. Schuurmans and P. Patrinos are with the Department of Electrical Engineering (ESAT-STADIUS), KU Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium. Email: {mathijs.schuurmans, panos.patrinos}@esat.kuleuven.be

[‡]P. Sopasakis is with Queen’s University Belfast, School of Electronics, Electrical Engineering and Computer Science, Centre for Intelligent Autonomous Manufacturing Systems, BT9 5AH, Northern Ireland, UK. Email: p.sopasakis@qub.ac.uk

The work of the first and third authors was supported by: FWO projects: No. G086318N; No. G086518N; Fonds de la Recherche Scientifique – FNRS, the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project No. 30468160 (SeLMA), Research Council KU Leuven C1 project No. C14/18/068 and the Ford–KU Leuven Research Alliance project No. KUL0023.

matrix A is denoted $\rho(A)$. We denote the dimensions of a vector x by n_x , and finally, we denote the ℓ_p -norm ball of radius r around $x \in \mathbb{R}^{n_x}$ as $\mathbb{B}_p(x, r)$.

II. PROBLEM STATEMENT

A. Stabilizing control of stochastic jump linear systems

This paper concerns the control of discrete-time stochastic jump linear dynamical systems with random disturbances w_t :

$$x_{t+1} = A(w_t)x_t + B(w_t)u_t. \quad (1)$$

The disturbances w_t take values on the finite sample space $\mathcal{W} := \mathbb{N}_{[1,k]}$ equipped with the discrete σ -algebra $2^{\mathcal{W}}$. For all $i \in \mathbb{N}_{[1,k]}$, we introduce the notation $A_i := A(i)$ and $B_i := B(i)$. Furthermore, let $P : 2^{\mathcal{W}} \rightarrow \mathbb{R}$, with $P[w = i] = P[\{i\}] = p_i$ be a probability measure, such that $(\mathcal{W}, 2^{\mathcal{W}}, P)$ defines a probability space. Note that system (1) is a specific type of Markov jump linear system (MJLS), where all the rows in transition probability matrix are identical. Furthermore, consider the analogously defined autonomous system

$$x_{t+1} = A(w_t)x_t, \quad (2)$$

for which the following fundamental notion of stability is defined.

Definition II.1 (Mean-square stability [18, Def. 3.8]). *We say that the autonomous system (2) is mean-square stable (MSS) if*

- 1) $\|\mathbb{E}[x_t]\| \rightarrow 0$ as $t \rightarrow \infty$ and
- 2) $\|\mathbb{E}[x_t x_t^\top]\| \rightarrow 0$ as $t \rightarrow \infty$,

for each $x_0 \in \mathbb{R}^{n_x}$.

This property can be verified by means of the following well-known conditions.

Theorem II.2 (Conditions for MSS). *Defining the operator $T : \mathbb{R}^k \rightarrow \mathbb{R}^{k n_x^2 \times k n_x^2}$ as*

$$T(p) := (p^\top \otimes \mathbf{1}_k \otimes I_{n_x^2}) \cdot \text{blkdiag}(\{A_i^\top \otimes A_i\}_{i \in \mathbb{N}_{[1,k]}}), \quad (3)$$

the following statements are equivalent:

- (S1) System (2) is MSS.
- (S2) $\rho(T(p)) < 1$.
- (S3) $\exists P \in \mathbb{S}_{++}^k : \sum_{i=1}^k p_i A_i^\top P A_i - P \prec 0$.

Proof. These results follow directly from [18, thm. 3.9 and Cor. 3.26]. \square

Ideally, our objective is to compute a linear state feedback gain K , such that the closed loop system

$$x_{t+1} = \bar{A}(w_t)x_t = (A(w_t) + B(w_t)K)x_t \quad (4)$$

is MSS. Unfortunately, however, application of **Theorem II.2** requires the knowledge of p , which is not available in practice. Instead, we assume to have access to a finite sample $\{w_i\}_{i=1}^N$ of N independent, identically distributed (i.i.d.) disturbance values. We will show that it is possible to leverage non-asymptotic statistical information to design linear feedback laws which lead to a mean-square stable closed loop with *high probability*.

B. Mean-square stability in probability

The proposed distributionally robust approach to certifying MSS in probability entails the use of the available data to determine a set of probability distributions $\mathcal{A} \subseteq \Delta_k$ so that with high confidence, $p \in \mathcal{A}$ — such a set is called an *ambiguity set* [9]. The requirement that the closed-loop system (4) is MSS for all $\pi \in \mathcal{A}$, leads to the *distributionally robust* variant of the Lyapunov-type stability condition (S3):

$$\exists P \in \mathbb{S}_{++}^{n_x} : \max_{\pi \in \mathcal{A}} \sum_{i=1}^k \pi_i \bar{A}_i^\top P \bar{A}_i - P \prec 0. \quad (5)$$

Due to the dual representation of coherent risk measures [19, Thm. 6.4], the resulting property is equivalent to risk-square stability with respect to the risk measure induced by \mathcal{A} [9].

Thus, given an ambiguity set \mathcal{A} which includes the true distribution p at a given confidence level, one can be equally confident that a controller for which the closed-loop system satisfies (5), is mean-square stabilizing.

The existence of such a controller depends on the system at hand. Therefore, it is useful to define the following required property of the open-loop system (1), which can be tested by feasibility of the problems described in **Section IV**.

Definition II.3 (Linear distributionally robust stabilizability). *We say that system (1) is linearly distributionally robustly stabilizable with respect to an ambiguity set \mathcal{A} (\mathcal{A} -LDRS) if there exists a linear state feedback law $u(x) = Kx$ such that the corresponding closed-loop system (4) is p -MSS for all $p \in \mathcal{A}$.*

Remark II.4. Based on **Definition II.3**, we may additionally define *linear robust stabilizability* (LRS) of (1) as Δ_k -LDRS, and *linear stochastic stabilizability* with respect to the distribution $\hat{p} \in \Delta_k$ (\hat{p} -LSS) as $\{\hat{p}\}$ -LDRS. Since for any $\mathcal{A}_1 \subseteq \mathcal{A}_2$, \mathcal{A}_2 -LDRS \Rightarrow \mathcal{A}_1 -LDRS, LRS and LSS can be viewed as the extreme cases of LDRS.

Remark II.5. Provided that system (1) is LRS, the proposed approach can certify MSS with arbitrary confidence, regardless of the sample size. In contrast to the robust approach, however, by collecting a (small) data sample, MSS can still be certified when the system is only \mathcal{A} -LDRS for some ambiguity set \mathcal{A} . The required sample size is prescribed by the bounds described below. We illustrate this in **Section V-C**.

III. LEARNING-BASED AMBIGUITY ESTIMATION

Given N independent samples w_1, \dots, w_N from the distribution of the disturbance, we define the empirical measure $\hat{p} = (\hat{p}_i)_{i=1}^k$, where

$$\hat{p}_i := \frac{1}{N} \sum_{j=1}^N \mathbf{1}_i(w_j), \quad (6)$$

for all $i \in \mathbb{N}_{[1,k]}$. We now derive upper bounds on the radius r of the ℓ_1 -ambiguity set $\mathcal{A}_r^{\ell_1}(\hat{p}) := \{\mu \in \Delta_k \mid \|\mu - \hat{p}\|_1 \leq r\}$ (see **Fig. 1**), such that for given $\alpha \in (0, 1)$

$$P(p \in \mathcal{A}_r^{\ell_1}(\hat{p})) \geq 1 - \alpha. \quad (7)$$




Fig. 1. Illustration of the ℓ_1 -ambiguity set $\mathcal{A}^{\ell_1}(\hat{p})$ for $k = 3$.

Given such an ambiguity set, it is then possible to use the aforementioned stability condition (5) to design controllers which are MS stabilizing with confidence $1 - \alpha$. This is discussed further in Section IV.

A. Dvoretzky-Kiefer-Wolfowitz

A statistical upper bound on the ℓ_1 -norm of $(\hat{p} - p)$ can be easily obtained by means of the *Dvoretzky-Kiefer-Wolfowitz* (DKW) inequality [20], which probabilistically bounds the error on the empirical estimate of the cumulative probability distribution. We observe that this bound can readily be translated to the error on the probability distribution p .

Theorem III.1 (DKW inequality for estimates of p). *For $p, \hat{p} \in \Delta_k$ and N as defined above, the ℓ_1 -distance between p and \hat{p} is bounded in probability by*

$$\mathbb{P}(\|\hat{p} - p\|_1 > 2k\epsilon) \leq 2e^{-2N\epsilon^2}. \quad (8)$$

Proof. The proof is in the Section. \square

From this confidence bound, we can directly derive the required radius, leading to the following result.

Corollary III.2. *For a probability space of dimension k , sample size N , and any given confidence level $1 - \alpha$, (7) holds with*

$$r = r_{\text{DKW}}(\alpha, k, N) := 2k\sqrt{\frac{\ln 2/\alpha}{2N}}. \quad (9)$$

Proof. This result follows directly from Theorem III.1. Define $r := 2k\epsilon$, and $\alpha := 2e^{-2N\epsilon^2}$ to obtain (9). \square

B. McDiarmid bounds

Alternatively, we may obtain a bound on the radius of the ℓ_1 -ambiguity set based on the following well-known measure concentration result, often referred to as *McDiarmid's inequality*.

Lemma III.3 (McDiarmid's inequality [21, Thm. 6.2]). *If a function $f : \mathcal{W}^N \rightarrow \mathbb{R}$ has the bounded differences property, i.e., there exist some constants $c_1, \dots, c_N \geq 0$ such that,*

$$\sup_{\substack{w_1, \dots, w_N \\ w'_i \in \mathcal{W}}} |f(w_1, \dots, w_N) - f(w_1, \dots, w'_i, \dots, w_N)| \leq c_i, \quad (10)$$

and $\{w_i\}_{i=1}^N$ are independent random variables, then

$$\mathbb{P}[f(w_1, \dots, w_N) - \mathbb{E}[f(w_1, \dots, w_N)] > \epsilon] \leq e^{-\epsilon^2/(2v)}, \quad (11)$$

where

$$v = \frac{1}{4} \sum_{i=1}^N c_i^2.$$

Theorem III.4. *For a probability space of dimension k , sample size N , and any given confidence level $1 - \alpha$, (7) holds with*

$$r = r_{\text{M}}(\alpha, k, N) := \sqrt{-\frac{2 \ln(\alpha)}{N}} + \sqrt{\frac{2(k-1)}{\pi N}} + \frac{4k^{1/2}(k-1)^{1/4}}{N^{3/4}}. \quad (12)$$

Proof. First, we define a function $\psi : \mathcal{W}^N \rightarrow [0, 2]$ and show that it satisfies the bounded differences condition (10):

$$\psi(w_1, \dots, w_N) := \|\hat{p} - p\|_1 = \sum_{j=1}^k |\hat{p}_j(w_1, \dots, w_N) - p_j|.$$

Due to the discrete support of w_i , modifying one of the values $w_i = l \in \mathcal{W}$ to $w'_i = m \in \mathcal{W}$ corresponds to increasing \hat{p}_l , and decreasing \hat{p}_m by an amount $1/N$. For ease of notation, we omit the function arguments and define $\psi' := \psi(w_1, \dots, w'_i, \dots, w_N)$ and $\hat{p}' := \hat{p}(w_1, \dots, w'_i, \dots, w_N)$, so that

$$\begin{aligned} \psi' &= \|\hat{p}' - p\|_1 \\ &= |\hat{p}'_l - p_l| + |\hat{p}'_m - p_m| + \sum_{\substack{j=1, \\ j \neq l, j \neq m}}^k |\hat{p}_j - p_j| \\ &= |\hat{p}_l - p_l - \frac{1}{N}| + |\hat{p}_m - p_m + \frac{1}{N}| + \sum_{\substack{j=1, \\ j \neq l, j \neq m}}^k |\hat{p}_j - p_j|. \end{aligned}$$

Thus, (10) holds with $c_i = 2/N$, and consequently $v = 1/N$. By Lemma III.3, then,

$$\mathbb{P}[\|\hat{p} - p\|_1 > \epsilon + \mathbb{E}[\|\hat{p} - p\|_1]] \leq e^{-\frac{N\epsilon^2}{2}}. \quad (13)$$

Moreover, from [22, Lemma 7], we obtain a tight upper bound for the expected ℓ_1 -norm of the estimation error:

$$\mathbb{E}[\|\hat{p} - p\|_1] \leq \sqrt{\frac{2(k-1)}{\pi N}} + \frac{4k^{1/2}(k-1)^{1/4}}{N^{3/4}}. \quad (14)$$

Using this result, (13) can easily be brought into the required form: Let $\alpha = e^{-\frac{N\epsilon^2}{2}} \Rightarrow \epsilon = \sqrt{2 \ln(1/\alpha)/N}$, and substitute (14) into (13) to obtain the bound (12). \square

The behavior of both bounds in terms of the sample size is similar; both decrease with the sample size N as $\mathcal{O}(1/\sqrt{N})$. In terms of k , however, by virtue of (14), $r_{\text{M}} \sim \mathcal{O}(\sqrt{k})$. This is an improvement to r_{DKW} , which increases linearly with k . See Section V-A for a numerical comparison between these bounds and an empirical estimation of their tightness.

IV. DESIGN OF DISTRIBUTIONALLY ROBUST CONTROLLERS

We revisit the Lyapunov-type stability condition (5), and restate it in a slightly more general form that is more convenient when applied for constructing stabilizing terminal conditions for a receding horizon strategy.

We denote the closed-loop dynamics corresponding to $w_t = i$ by $f_i(x, Kx) = A_i + B_i Kx$, define $\ell(x, u) := x^\top Qx + u^\top Ru$ with $Q \in \mathbb{S}_{++}^{n_x}$ and $R \in \mathbb{S}_{++}^{n_u}$, and denote the quadratic candidate Lyapunov function as $V(x) := x^\top Px$. Due to the homogeneity of (5), we may replace the strict inequality by a non-strict inequality and introduce the negative definite quadratic form $-\ell(x, Kx)$ in the right-hand side, to obtain an equivalent condition in these terms:

$$\exists P \in \mathbb{S}_{++}^{n_x} : \max_{\pi \in \mathcal{A}} \sum_{i=1}^k \pi_i V(f_i(x, Kx)) \leq V(x) - \ell(x, Kx), \quad (15)$$

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

Given that \mathcal{A} is a polytope, it has a finite set of extreme points, that is, $\mathcal{A} = \text{conv}\{a^l\}_{l=1}^{n_{\mathcal{A}}}$. Since the maximum over a polyhedral set is attained at an extreme point [23, Thm. 32.2], (15) is equivalent to requiring that $\sum_{i=1}^k a_i^l V(f_i(x, Kx)) \leq V(x) - \ell(x, Kx)$ for all $l \in \mathbb{N}_{[1, n_{\mathcal{A}}]}$. However, the enumeration of the vertices of \mathcal{A} is typically computationally intensive and $n_{\mathcal{A}}$ grows rapidly with k (see Section V-B for timings using $\mathcal{A}_r^{\ell_1}(\hat{p})$).

We therefore present a methodology for the determination of a gain K and a matrix P that satisfies (15) for the ℓ_1 -ambiguity set $\mathcal{A}_r^{\ell_1}(\hat{p})$, without enumerating its vertices. This methodology is based on the following lemma.

Lemma IV.1. *Let $v(x) = (v_1(x), \dots, v_k(x))$ with*

$$v_i(x) = V(f_i(x, Kx)) = x^\top (A_i + B_i K)^\top P (A_i + B_i K)x,$$

for $x \in \mathbb{R}^{n_x}$, and let $\hat{p} \in \Delta_k$ denote the empirical estimate (6). Given that the ambiguity set is equal to the ℓ_1 -ambiguity set of radius r around \hat{p} , i.e., $\mathcal{A} = \mathcal{A}_r^{\ell_1}(\hat{p})$, the distributionally robust Lyapunov-type stability condition (15) is equivalent to the condition that there exist k functions $z_i : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$, such that

$$v_i(x) - z_i(x) \pm r z_j(x) + z(x)^\top \hat{p} \leq V(x) - \ell(x, Kx) \quad (16)$$

for all $i, j \in \mathbb{N}_{[1, k]}$ and $x \in \mathbb{R}^{n_x}$.

Proof. The left-hand side of the inequality in (15) is equivalent to the definition of the support function of \mathcal{A} , evaluated at $v(x)$. Computing $\sigma_{\mathcal{A}}(v(x))$ directly is seemingly not an easy task. However, \mathcal{A} can be written as the intersection of two sets with easily computable support functions,

$$\mathcal{A} = \Delta_k \cap C,$$

where $C := \mathbb{B}_1(\hat{p}, r)$. In fact,

$$\sigma_{\Delta_k}(v) = \max\{v_1, \dots, v_k\} \quad (17a)$$

$$\sigma_C(v) = r \|v\|_\infty + v^\top \hat{p} \quad (17b)$$

By the Attouch-Brézis theorem [24, Thm. 15.3],

$$\sigma_{\Delta_k \cap C}(v) = (\sigma_{\Delta_k} \square \sigma_C)(v)$$

where \square denotes infimal convolution given by

$$(\sigma_{\Delta_k} \square \sigma_C)(v) = \inf_z \sigma_{\Delta_k}(v - z) + \sigma_C(z).$$

Therefore we can equivalently express (15) as

$$\inf_z \sigma_{\Delta_k}(v(x) - z) + \sigma_C(z) \leq V(x) - \ell(x, Kx), \quad (18)$$

for all $x \in \mathbb{R}^{n_x}$. Eq. (18) is true if and only if there exists a $z(x) = (z_1(x), \dots, z_k(x))$ such that

$$\sigma_{\Delta_k}(v(x) - z(x)) + \sigma_C(z(x)) \leq V(x) - \ell(x, Kx), \quad (19)$$

for all $x \in \mathbb{R}^{n_x}$. Using (17), we express (19) as

$$\max_{i \in \mathbb{N}_{[1, k]}} \{v_i(x) - z_i(x)\} + r \|z(x)\|_\infty + z(x)^\top \hat{p} \leq V(x) - \ell(x, Kx) \quad (20)$$

In turn, this is true if and only if

$$v_i(x) - z_i(x) \pm r z_j(x) + z(x)^\top \hat{p} \leq V(x) - \ell(x, Kx)$$

for all $i, j \in \mathbb{N}_{[1, k]}$ and $x \in \mathbb{R}^{n_x}$, which is exactly condition (16). \square

We shall proceed by assuming that the components of $z(x)$ are quadratic functions of x of the form $z_i(x) = x^\top H_i x$, where $H_i \in \mathbb{R}^{n_x \times n_x}$ are symmetric matrices, which allows to cast (16) as a set of $2k^2$ matrix inequalities

$$(A_i + B_i K)^\top P (A_i + B_i K) - H_i \pm r H_j + \sum_{l=1}^k p_l H_l - P + Q + K^\top R K \preceq 0. \quad (21)$$

for $i, j \in \mathbb{N}_{[1, k]}$, which can be described by an LMI as shown in the following proposition.

Proposition IV.2. *The matrix inequality (21) is equivalent to the LMIs*

$$\begin{bmatrix} -W - \hat{H}_i \pm r \hat{H}_j + \sum_{l=1}^k p_l \hat{H}_l & W A_i^\top + Z^\top B_i^\top & W Q^{\frac{1}{2}} & Z^\top R^{\frac{1}{2}} \\ * & -I_{n_x} & 0 & 0 \\ * & * & -I_{n_u} & 0 \\ * & * & * & -W \end{bmatrix} \preceq 0$$

for all $i, j \in \mathbb{N}_{[1, k]}$, where $P := W^{-1}$, for $W \in \mathbb{S}_{++}^{n_x}$, $\hat{H} = W H_i W$ and $Z = K W$.

Proof. We pre- and post- multiply (21) by W to obtain,

$$(A_i W + B_i Z)^\top W^{-1} (A_i W + B_i Z) - W + \hat{H}_i \pm r \hat{H}_j + \sum_{l=1}^k p_l \hat{H}_l + (W Q^{\frac{1}{2}})^\top (W Q^{\frac{1}{2}}) - (Z R^{\frac{1}{2}})^\top (Z R^{\frac{1}{2}}) \preceq 0. \quad (22)$$

Now define

$$\Theta_i := \begin{bmatrix} A_i W + B_i Z \\ Q^{\frac{1}{2}} W \\ R^{\frac{1}{2}} Z \end{bmatrix}, \quad D := \begin{bmatrix} W & \\ & I_{n_x + n_u} \end{bmatrix}$$

to obtain $\Theta_i^\top D^{-1} \Theta_i + (-\hat{H}_i \pm r \hat{H}_j + \sum_{l=1}^k p_l \hat{H}_l) \preceq 0$, which, by the Schur complement lemma [25, Sec. 2.1] is equivalent to the LMI

$$\begin{bmatrix} -W - \hat{H}_i \pm r \hat{H}_j + \sum_{l=1}^k p_l \hat{H}_l & \Theta_i^\top \\ \Theta_i & D \end{bmatrix} \preceq 0, \\ W \succ 0,$$

which expands to the given LMI. \square

The assumption that the components of $z(x)$ are quadratic can be justified by noting that a mapping z that minimizes the lefthand side in (18) can be taken to be a piecewise affine function of v [26]. In fact, due to homogeneity of the support functions, it can be easily seen that z can be taken to be piecewise linear. Therefore, $z(x)$ is in fact piecewise quadratic and homogeneous of degree two. However, the task of computing the exact expression of z is equivalent to solving a parametric linear program, hence as complex as enumerating the vertices of \mathcal{A} . Therefore, a sensible approximation is to impose that $z(x)$ is simply quadratic. Moreover, in Section V, we demonstrate that in practice, the induced conservativeness is limited, whereas the computational advantage of the reformulation in Proposition IV.2 compared to vertex enumeration allows us to solve problems of a significantly larger scale.

Lastly, note that the derivation leading to Lemma IV.1 is not limited to ℓ_1 -based — or even polytopic — ambiguity sets, as it can easily be repeated for other ambiguity sets which consist of an intersection of convex sets with easily computable support functions.

V. NUMERICAL EXPERIMENTS

A. Data-driven ambiguity bounds

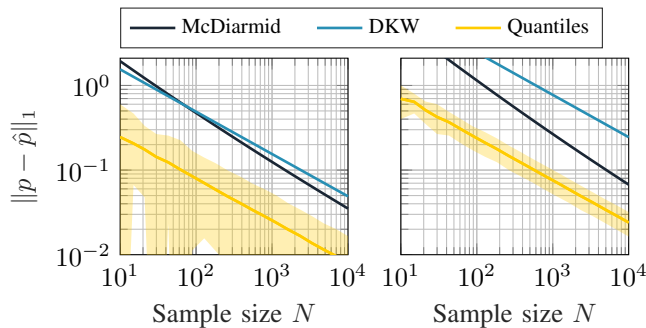


Fig. 2. Comparison of the derived bounds r_M and r_{DKW} at confidence level $1 - \alpha = 0.9$ for (Left) $k = 2$ and (Right) $k = 10$. Additionally, the shaded area is bounded by the empirical 0.1 and 0.9 quantiles of $\|p - \hat{p}\|_1$, based on 10,000 randomly generated data sets.

We compare the behavior of the DKW-based radius (III-A) and the radius based on McDiarmid’s inequality (III-B) with respect to increasing sample sizes. Figure 2 shows a comparison for two values of k . Since r_M scales better with k ($\mathcal{O}(\sqrt{k})$) compared to r_{DKW} ($\mathcal{O}(k)$), r_{DKW} is generally lower than r_M , especially for large values of k . However, for very low values of k and N , Figure 2 demonstrates that r_{DKW} is tighter, albeit only by a small margin. In practice, we

may of course exploit the closed-form expressions to obtain a tighter bound which is simply $r = \min\{r_M, r_{DKW}\}$. Figure 3 illustrates the corresponding ℓ_1 -ambiguity sets for $k = 3$.




Fig. 3. Probability estimate \hat{p} and ambiguity set $\mathcal{A}_r^{\ell_1}(\hat{p})$ at confidence level $1 - \alpha = 0.9$ for (Left) $N = 10$, (Middle) $N = 100$ and (Right) $N = 5000$.

B. Methods for controller design

1) *Timings*: In Section IV, we derived an approximation of the Lyapunov-type stability condition (15) which allows us to compute stabilizing controllers with high confidence without solving as many LMIs as the number of vertices of the polytopic ambiguity set $\mathcal{A}_r^{\ell_1}(\hat{p})$. In Figure 4 we present a comparison of these approaches in terms of computational complexity for a system with $n_x = n_u = 2$. For $k > 7$, the vertex enumeration approach fails due to excessive memory requirements caused by the rapid increase in the number of vertices n_A . On the same machine, using the proposed reformulation, problems of at least $k = 30$ could still be solved without running out of memory. Moreover, we observe that simply computing the vertices of $\mathcal{A}_r^{\ell_1}(\hat{p})$ already proves to be more time-consuming a problem than solving the complete LMI of the reformulation (21).




Fig. 4. Time to solve the LMI formulations of (15). We compare vertex enumeration and the reformulation in Proposition IV.2. The solid lines represent the total time of the methods. For the vertex enumeration approach, the dashed line separately shows the time to compute the vertices of the ambiguity set. (Vertex computations are performed using the MPT [27] toolbox, LMIs were solved using MOSEK [28], on an Intel Core i7-7700K CPU at 4.20GHz.)

2) *Approximation quality*: We observe that in practice, the conservativeness introduced by the reformulation is often negligible. During experimentation, we have not been able to find a system for which no feasible feedback gain could be found through the reformulation while there could through

vertex enumeration. This is further illustrated by the following example. Consider the system with dynamics

$$A_1 = \begin{bmatrix} 0.9 & 1 \\ 0 & 0.99 \end{bmatrix}, A_2 = \begin{bmatrix} 1.5 & 2.5 \\ 0 & 2.5 \end{bmatrix}, B_1 = B_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

For $\hat{p}_1 = \hat{p}_2 = 0.5$, $r = 0.1$, $Q = 10^{-4}I_2$ and $R = 10^{-4}$, we estimate the sets \mathcal{F} and $\hat{\mathcal{F}}$ of feasible control gains for the exact approach (using vertex enumeration) and the reformulated LMI of Proposition IV.2, respectively. That is, $\mathcal{F} := \{K \in \mathbb{R}^2 \mid (15) \text{ holds}\}$, and $\hat{\mathcal{F}} := \{K \in \mathbb{R}^2 \mid \exists H_i, i \in \mathbb{N}_{[1,k]} : (21) \text{ holds}\}$. We construct a regular grid of potential feedback gains $K = [K_1 \ K_2]$ and verify whether a P (or equivalently, W) exists such that the involved LMI is satisfied. This point is then marked with the corresponding color in Figure 5. Since feasibility of (21) implies feasibility of (15), it is $\hat{\mathcal{F}} \subseteq \mathcal{F}$. We find that the experimental estimates of $\hat{\mathcal{F}}$ and \mathcal{F} nearly fully overlap. In fact, in this set of 10,000 samples of K , only 4 instances out of 2825 that are in \mathcal{F} , are not in $\hat{\mathcal{F}}$.




Fig. 5. Estimates of the feasible set \mathcal{F} and $\hat{\mathcal{F}}$, defined in Section V-B. All sampled points $K \in \mathcal{F} \setminus \hat{\mathcal{F}}$ are encircled. Infeasible points ($K \notin \mathcal{F}$) are left blank.

C. Comparison with stochastic and robust approaches

The following example demonstrates (i) the superior sample complexity of the distributionally robust approach over the stochastic approach, based on the bounds obtained in Section III; and (ii) the improved applicability in comparison with the robust approach.

This example is based on estimating the *distributional stability region* of an autonomous dynamical system (4). We define this as the set \mathcal{S} of all probability vectors p for which the system is MSS. Using the operator T , defined in (3), we can denote this set as

$$\mathcal{S} := \{p \in \Delta_k \mid \rho(T(p)) < 1\}. \quad (23)$$

While it is easy to test whether the system is p -MSS for some given p , it does not seem to be easy to determine \mathcal{S} . Indeed, since the spectral radius of a matrix is generally not convex, aside from very specific cases, this set is difficult to analyze for most systems.

However, for the following simple system

$$A_1 = A, B_1 = B, A_2 = 0, B_2 = 0, \quad (24)$$

which is of particular interest in networked systems, it is shown in [29] that \mathcal{S} for the closed-loop system with $u(x) = Kx$ can be written explicitly as

$$\mathcal{S} = \left\{ p \in \Delta_2 \mid p_1 < \frac{1}{\rho(A+BK)^2} \right\}. \quad (25)$$

This set simply defines a half-open line section in \mathbb{R}^2 and is thus convex. Using the convexity of this set, we may devise a simple procedure to estimate a lower bound on the confidence that a given linear controller is MSS for the true distribution, given only that it is stabilizing for \hat{p} , which is estimated based on N i.i.d. data points. In fact, we compute $r^* = \max\{r \in [0, 2] \mid \mathcal{A}_r^{\ell_1}(\hat{p}) \subseteq \mathcal{S}\}$. Since the inclusion $\mathcal{A}_r^{\ell_1}(\hat{p}) \subseteq \mathcal{S}$ can be verified easily using (25), r^* is readily computed numerically by means of a simple bisection scheme. The bounds derived in Section III now associate each r^* with a lower bound $(1 - \alpha^*(N))$ on the probability that a closed-loop system is p -MSS. We have that $\alpha^*(N) = \min\{\alpha_M^*(N), \alpha_{\text{DKW}}^*(N)\}$, which, by rearranging the terms in (9) and (12), and setting $k = 2$, can be shown to be

$$\alpha_M^*(N) = e^{-\frac{N}{2}(\sqrt{2/\pi N} + 2\sqrt{2}/N^{3/4} - r^*)^2},$$

$$\alpha_{\text{DKW}}^*(N) = 2e^{-\frac{N(r^*)^2}{8}}.$$

Consider now the open-loop stochastic jump linear system of the form (24), with

$$A = \begin{bmatrix} 1.05 & 1.8 \\ 0 & 1.1 \end{bmatrix}, B = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

and with unknown distribution $p \in \Delta_2$. Given N i.i.d. observations of the disturbance w , we obtain an empirical probability estimate \hat{p} , and a feedback gain K according to the stochastic approach, i.e., the closed-loop system satisfies (15) for $\mathcal{A} = \{\hat{p}\}$. We compute r^* , such that $\mathcal{A}_r^{\ell_1}(\hat{p})$ is a tight under-approximation of \mathcal{S} . We repeat this process for increasing values of N and plot the corresponding confidence $1 - \alpha^*(N)$ that the system is p -MSS in Figure 6.

Similarly, to evaluate the distributionally robust approach, we compute the largest r , such that (15) is feasible for $\mathcal{A}_r^{\ell_1}(\hat{p})$ and obtain a feedback gain K from solving the corresponding LMI problem. We again repeat this for increasing values of N and plot the lower bound on $P(p \in \mathcal{A}_r^{\ell_1}(\hat{p}))$ in Figure 6.

Note that the given system is not LRS, i.e., no linear controller exists that can stabilize the system in mean-square sense for all $p \in \Delta_2$. Therefore, the robust approach is not applicable. However, using the distributionally robust approach in a data-driven manner, it suffices to acquire 62 data points in order to find a controller for this same system, which is mean-square stabilizing with over 99.8% confidence. By contrast, obtaining similar guarantees from the stochastic approach, requires nearly 50,000 data points.

VI. CONCLUSION

We studied the problem of data-driven synthesis of a static linear state feedback gain for stochastic jump linear systems that grants MSS *with high probability*. To this end, we adopted a distributionally robust approach, focusing specifically on ℓ_1 -ambiguity sets. We derived bounds that guarantee the inclusion of the true distribution in this set at the given confidence level and impose MS stability for all distributions within this ambiguity set. To efficiently solve this problem, we derived an LMI formulation which approximates the corresponding Lyapunov-type stability condition, but grows polynomially with the support of the dimension of the




Fig. 6. Lower bound on the probability of obtaining a MSS controller with respect to the true distribution. For example, in order to be at least 99.8% confident that the closed-loop system is mean-square stable, the stochastic approach requires a sample of about 50,000 data points, whereas the distributionally robust approach requires merely 62 data points. Note that no common controller exists which stabilizes both modes of this system (the system is not LRS).

probability space. Our findings were illustrated and verified through several numerical experiments.

In future work, we aim to generalize this methodology to Markovian disturbances and nonlinear systems. Additionally, we aim to study the use of this methodology to design learning-based terminal conditions for model predictive control.

REFERENCES

- [1] U. Rosolia and F. Borrelli, “Learning model predictive control for iterative tasks. a data-driven control framework,” *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2018.
- [2] A. Domahidi, M. N. Zeilinger, M. Morari, and C. N. Jones, “Learning a feasible and stabilizing explicit model predictive control law by robust optimization,” in *2011 50th IEEE conference on decision and control and European control conference*, pp. 513–519, IEEE, 2011.
- [3] J. Dupačová, “The minimax approach to stochastic programming and an illustrative application,” *Stochastics: An International Journal of Probability and Stochastic Processes*, vol. 20, no. 1, pp. 73–88, 1987.
- [4] A. Shapiro and A. Kleywegt, “Minimax analysis of stochastic problems,” *Optimization Methods and Software*, vol. 17, no. 3, pp. 523–542, 2002.
- [5] J. Hou, K. Huang, and Z. Zhang, “Distributionally robust optimization leads to better generalization: on sgd and beyond,” 2018.
- [6] E. Smirnova, E. Dohmatob, and J. Mary, “Distributionally robust reinforcement learning,” *arXiv preprint arXiv:1902.08708*, 2019.
- [7] B. P. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, “Distributionally robust control of constrained stochastic systems,” *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2016.
- [8] I. Yang, “Wasserstein distributionally robust stochastic control: A data-driven approach,” *arXiv preprint arXiv:1812.09808*, 2018.
- [9] P. Sotasakis, D. Herceg, A. Bemporad, and P. Patrinos, “Risk-averse model predictive control,” *Automatica*, vol. 100, pp. 281–288, 2019.
- [10] P. Patrinos, P. Sotasakis, H. Sarimveis, and A. Bemporad, “Stochastic model predictive control for constrained discrete-time Markovian switching systems,” *Automatica*, vol. 50, no. 10, pp. 2504–2514, 2014.
- [11] E. Delage and Y. Ye, “Distributionally robust optimization under moment uncertainty with application to data-driven problems,” *Operations research*, vol. 58, no. 3, pp. 595–612, 2010.
- [12] D. Bertsimas, V. Gupta, and N. Kallus, “Data-driven robust optimization,” *Mathematical Programming*, vol. 167, no. 2, pp. 235–292, 2018.
- [13] P. M. Esfahani and D. Kuhn, “Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations,” *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.
- [14] C. Zhao and Y. Guan, “Data-driven risk-averse stochastic optimization with Wasserstein metric,” *Operations Research Letters*, vol. 46, no. 2, pp. 262–267, 2018.
- [15] R. Gao and A. J. Kleywegt, “Distributionally robust stochastic optimization with Wasserstein distance,” *arXiv preprint arXiv:1604.02199*, 2016.

- [16] R. Jiang and Y. Guan, “Data-driven chance constrained stochastic program,” *Mathematical Programming*, vol. 158, no. 1-2, pp. 291–327, 2016.
- [17] A. Ben-Tal, D. Den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen, “Robust solutions of optimization problems affected by uncertain probabilities,” *Management Science*, vol. 59, no. 2, pp. 341–357, 2013.
- [18] O. L. V. Costa, M. D. Fragoso, and R. P. Marques, *Discrete-time Markov jump linear systems*. Springer Science & Business Media, 2006.
- [19] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- [20] P. Massart, “The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality,” *The Annals of Probability*, pp. 1269–1283, 1990.
- [21] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [22] S. Kamath, A. Orlitsky, D. Pichapati, and A. T. Suresh, “On learning distributions from their samples,” in *Conference on Learning Theory*, pp. 1066–1100, 2015.
- [23] R. Rockafellar, *Convex analysis*. Princeton university press, 1997.
- [24] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer International Publishing, 2017.
- [25] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. SIAM, 1994.
- [26] P. Patrinos and H. Sarimveis, “Convex parametric piecewise quadratic optimization: Theory and algorithms,” *Automatica*, vol. 47, no. 8, pp. 1770–1777, 2011.
- [27] M. Kvasnica, P. Grieder, and M. Baotić, “Multi-Parametric Toolbox (MPT),” 2004.
- [28] MOSEK ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 8.1.*, 2017.
- [29] K. Gatsis and G. J. Pappas, “Sample complexity of networked control systems over unknown channels,” in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 6067–6072, IEEE, 2018.

APPENDIX

♠ Proof of Theorem III.1.

Let F denote the cumulative distribution function (cdf) of w and define \hat{F} to be the empirical cumulative distribution given N samples $\{w_j\}_{j=1}^n$, that is

$$\hat{F}_i := \frac{1}{N} \sum_{j=1}^N 1_{w_j \leq i}.$$

The DKW inequality [20] states that

$$\mathbb{P} \left(\max_{i \in \mathcal{W}} |\hat{F}_i - F_i| > \epsilon \right) \leq 2e^{-2N\epsilon^2}, \quad \epsilon > 0, N \in \mathbb{N}. \quad (26)$$

the cdf F and the probability mass function (pmf) p of a discrete distribution are related as

$$\begin{cases} p_1 &= F_1 \\ p_i &= F_i - F_{i-1}, \text{ for } i \in \mathbb{N}_{[2,k]}. \end{cases}$$

The same relation holds between the empirical counterparts \hat{p} and \hat{F} . Therefore, a bound of the form $|\hat{F}_i - F_i| \leq \epsilon$ implies that, for $i > 1$

$$\begin{aligned} |\hat{p}_i - p_i| &= |\hat{F}_i - \hat{F}_{i-1} - (F_i - F_{i-1})| \\ &\leq |\hat{F}_i - F_i| + |\hat{F}_{i-1} - F_{i-1}| \leq 2\epsilon \end{aligned}$$

For $i = 1$, this inequality trivially hold as well. Thus, writing (26) in terms of the probability mass function, we obtain

$$\mathbb{P} \left(\max_{i \in \mathcal{W}} |\hat{p}_i - p_i| > 2\epsilon \right) \leq 2e^{-2N\epsilon^2},$$

which directly implies (8). \square