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

Mathijs Schuurmans[†], Pantelis Sopasakis[‡], Panagiotis Patrinos[†] | [†]KU Leuven, Belgium, [‡] Queen's University Belfast, Northern Ireland

Background and motivation

Safe learning-based control

- Control of stochastic systems requires knowledge of underlying **probability distributions**
- In practice: distributions are **unknown**
- Distributionally robust approach [1, 2]: control assuming **worst-case distribution in ambiguity set** $\mathcal{A}$

Gather data → Ambiguity decreases → Safely reduce conservativeness

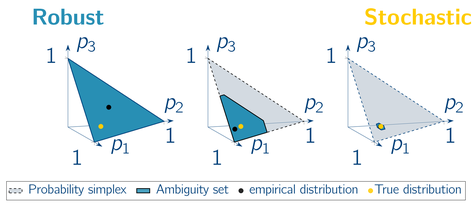

Applications

Rigorous statistical guarantees → **safety-critical** applications
e.g., autonomous driving, robotics, ... (physical interaction with humans)

Problem statement

We aim to **stabilize** a linear system

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

where random variables $w_t \in \mathcal{W} := \{1, 2, \dots, k\}$, specify the operation mode ($A(i) = A_i$, $B(i) = B_i$) at time t, and $\mathsf{P} : 2^{\mathcal{W}} \to \mathbb{R}$, with $\mathsf{P}[w = i] = \mathsf{P}[\{i\}] = p_i$ is an **unknown** probability measure.

Challenge: Mean-square (MS) stability conditions depend on true distribution p [3].

Goal: Mean-square stability in probability

For a given confidence level $1-\alpha \in (0, 1)$, compute a **linear state feedback gain** K, which renders (1) MS stable with probability at least $1 - \alpha$.

References

[1] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.

[2] P. Sopasakis, D. Herceg, A. Bemporad, and P. Patrinos, "Risk-averse model predictive control," *Automatica*, vol. 100, pp. 281–288, 2019.

[3] O. L. V. Costa, M. D. Fragoso, and R. P. Marques, *Discrete-time Markov jump linear systems*. Springer Science & Business Media, 2006.

[4] 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.

[5] 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.

Proposed approach

- **Estimate distribution** $\hat{p}$ based on N i.i.d. samples $\{w_i\}_{i=1}^N$
- **Ensure MS stability** for all $p \in \mathcal{A}_r^{\ell_1}(\hat{p}) := \{p \in \Delta_k \mid \|p-\hat{p}\|_1 \le r\}$

Subproblems

I Compute **min** r such that,

$$\mathsf{P}[p \in \mathcal{A}_r^{\ell_1}(\hat{p})] \ge 1 - \alpha \tag{2}$$

II Efficiently compute K that is MS stabilizing for all $p \in \mathcal{A}_r^{\ell_1}(\hat{p})$

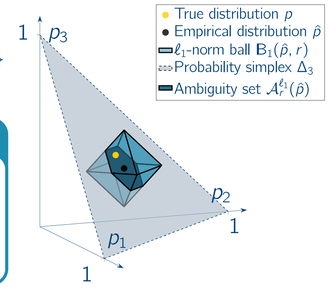

I Bounding the ambiguity

PAC-type confidence bounds for the empirical probability distribution estimate:

if $r \le \mathbf{min}\{r_{\mathsf{DKW}}, r_{\mathsf{M}}\} \Rightarrow$ (2) holds

Dvoretzky-Kiefer-Wolfowitz

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

$$\Rightarrow r_{\mathsf{DKW}}(\cdot, k, N) = \mathcal{O}\left(\frac{k}{\sqrt{N}}\right)$$

McDiarmid + min-max loss bounds [4]

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

$$\Rightarrow r_{\mathsf{M}}(\cdot, k, N) = \mathcal{O}\left(\frac{k^{3/4}}{\sqrt{N}}\right)$$

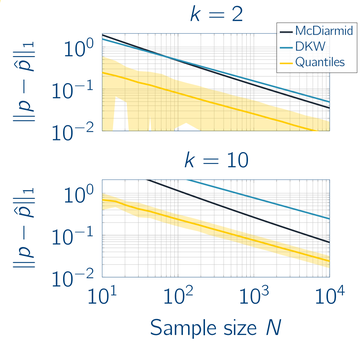

II Efficient computation of the feedback gain

Distributionally Robust Lyapunov-type stability condition:

$$\exists P \succ 0 : \max_{p \in \mathcal{A}} \sum_{i=1}^{k} p_i V(\bar{A}_i x) \le V(x) - \ell(x, Kx), \tag{3}$$

with $\bar{A}_i = A_i + B_i K$, $V(x) = x^\top P x$, and $\ell(x, u) = x^\top Q x + u^\top R u$, with $Q \succeq 0$, $R \succ 0$.

Method 1 – Vertex enumeration

ℓ_1-ambiguity set is **polytopic**: $\mathcal{A}_r^{\ell_1}(\hat{p}) = \mathbf{conv}\{p^{(i)}\}_{i=1}^{n_{\mathcal{A}}}$.

$$(3) \Leftrightarrow \max_{l \in \mathbb{N}_{[1, n_{\mathcal{A}}]}} \sum_{i=1}^{k} p_i^{(l)} V(\bar{A}_i x) \le V(x) - \ell(x, Kx)$$

→ Reduced to finite number of LMI conditions.

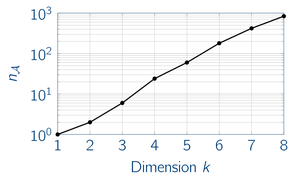

Drawback: computational cost of 1) computing vertices; and 2) solving LMI for all vertices.

Method 2 – Reformulation

Reformulation based on 3 key observations:

1 $\max_{p \in \mathcal{A}} \sum_{i=1}^{k} p_i V(\bar{A}_i x) = \sigma_{\mathcal{A}}(V(\bar{A}_i x))$

2 $\mathcal{A}_r^{\ell_1}(\hat{p}) = \Delta_k \cap \mathbf{B}_1(\hat{p}, r) = \Delta_k \cap C$ ⟶

3 $\sigma_{\Delta_k \cap C}(v) = (\sigma_{\Delta_k} \square \sigma_C)(v)$

Infimal convolution: $(f \square g)(v) = \mathbf{inf}_z\, f(v - z) + g(z)$

Close approximation leads to $2K^2$ LMIs.

Easily computable support function of elementary sets

$\sigma_{\Delta_k}(v) = \mathbf{max}\{v_1, \dots, v_k\}$
$\sigma_C(v) = r\|v\|_\infty + v^\top \hat{p}$

Experimental results

Computational cost

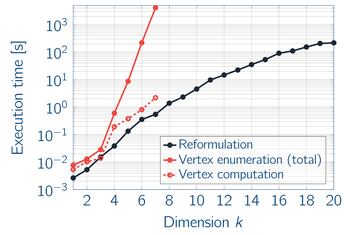

- Rapid growth of vertex count → vertex enumeration applicable for very low dimensions ($k \le 7$) of the sample space $\mathcal{W}$.
- Computation of the vertices of the ambiguity set alone is more time-consuming than solving the complete reformulated problem.

Sample complexity

Set-up

1 Given a closed-loop system

$$x_{t+1} = \Big(A(w_t) + KB(w_t)\Big)x_t \tag{4}$$

with a $\hat{p}$-MSS controller K *(stochastic approach)*

2 Define **distributional stability region** $\mathcal{S} := \{p' \mid (4)\ p'\text{-MSS}\}$

3 Compute **max** r, s.t. $\mathcal{A}_r^{\ell_1}(\hat{p}) \subseteq \mathcal{S} \Rightarrow \mathsf{P}((4)\ p\text{-MSS}) \ge (1 - \alpha)$

For Bernoulli system in [5], $\mathcal{A}_r^{\ell_1}(\hat{p}) \subseteq \mathcal{S}$ is easy to test.

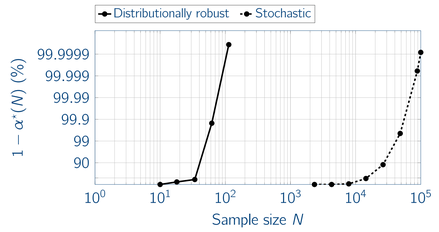

For a **given confidence level**, the distributionally robust approach provides a stabilizing controller using several orders of magnitude **less data**.

Future work

- Extend to **optimal control** setting and nonlinear dynamics
- Relax i.i.d. assumption → Markov jump systems