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

Mathijs Schuurmans[1], Pantelis Sopasakis[2], Panagiotis Patrinos[1]

[1] ESAT - STADIUS, KU Leuven
Kasteelpark Arenberg 10
3001, Leuven, Belgium

[2] School of EEECS, i-AMS,
Queen's University Belfast
Ashby Building, Stranmillis Road,
Belfast BT9 5AH, Northern Ireland, UK

1 Introduction and background

In order to guarantee safe and reliable operation of control systems, it is important that the presence of uncertainty is adequately accounted for. *Distributionally robust* control is a framework that enables this by generalizing the two opposing approaches of *stochastic* and *robust* control [4]. The distributionally robust framework relaxes the robust approach by introducing a so-called *ambiguity set*, from which the underlying distribution is assumed to be chosen adversarially. The challenge is to appropriately design this ambiguity set in order to make a suitable trade-off between robustness and performance, given the required confidence and the level of uncertainty.

2 Problem statement

We consider the stochastic jump linear dynamical system with random disturbances w_t:

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

These disturbances are assumed to be drawn i.i.d. from a discrete distribution with finite support of length k. The distribution is unknown but we assume to have access to a data sample of size N. The aim is to find a linear state feedback control gain K that will render the closed-loop system mean-square stable with a given confidence level $1-\alpha \in (0,1)$. If the unknown data-generating distribution $p^\star \in \mathscr{A}$, then $u(x) = Kx$ is a mean-square stabilizing controller if there exists a $P = P^T \succ 0$, such that

$$\max_{p\in\mathscr{A}} \mathbf{E}_{\omega\sim p}\left[[A(\omega)+B(\omega)K]^T P[A(\omega)+B(\omega)K] - P\right] \prec 0.$$

Therefore, it suffices to construct an ambiguity set $\mathscr{A}$, such that $\mathbf{P}(p^\star \in \mathscr{A}) \geq 1-\alpha$.

3 Approach and Results

Similarly to the work in [2], we define the ambiguity set $\mathscr{A}_r(\hat{p})$ as the set of all distributions within a ball $\mathscr{B}_r(\hat{p})$ of radius r around the empirical distribution $\hat{p}$ of the available data, with respect to a certain statistical distance metric. In particular, we focus on the total variation distance, which induces a polytopic ambiguity set (see fig. 1). As a result, the

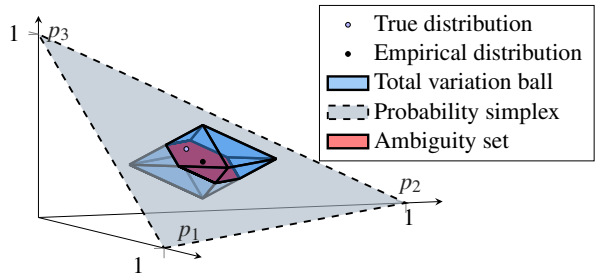


Figure 1: Construction of an ambiguity set using the total variation distance around the empirical distribution.

controller can be computed by solving a finite number of linear matrix inequalities (LMI). We additionally use conjugate duality to derive a tractable LMI formulation of the associated stability conditions, which avoids vertex enumeration and can generalize to a wider range of ambiguity sets. Combining results from statistical learning [3] with the McDiarmid inequality [1], we then derive an upper bound for the required radius as a function of the sample size N, the support length k of w and the required confidence $1-\alpha$:

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

This bound allows us to provide stability guarantees for any sample size, while being less conservative than the robust approach.

Acknowledgements This work was supported by FWO projects: G086318N; G086518N; EOS Project no 30468160 (SeLMA); Research Council KUL C14/18/068

References

[1] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

[2] P. M. Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.

[3] S. Kamath, A. Orlitsky, D. Pichapati, and A. T. Suresh. On learning distributions from their samples. In *Conference on Learning Theory*, pages 1066–1100, 2015.

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