Data-driven distributionally robust LQR with multiplicative noise

Peter Coppens[1] peter.coppens@kuleuven.be
Mathijs Schuurmans[1] mathijs.schuurmans@kuleuven.be
Panagiotis Patrinos[1] panos.patrinos@kuleuven.be

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

1 Introduction

We develop controllers for linear systems with time-varying parametric uncertainty. These may cover a wide range of systems. In our case the uncertainty varies stochastically. Most existing approaches then assume that the distribution of w is known [1], which is not the case in many practical applications. To this end, we adopt a data-driven distributonally robust (DR) approach [2] towards solving the infinite-horizon Stochastic Linear Quadratic Regulator (SLQR) problem. To do so 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. The result is then a controller that is mean square stabilizing with high probability.

2 Methodology

The dynamics of the linear system with stochastic parameteric uncertainty are described by:

$$\begin{aligned} &x_{k+1} = A(w_k)x_k + B(w_k)u_k \\ \text{with}\quad &A(w) = A_0 + \textstyle\sum_{i=1}^{n_w} w^i A_i, \quad B(w) = B_0 + \sum_{i=1}^{n_w} w^i B_i, \end{aligned}$$

with $x_k \in \mathbf{R}^{n_x}$ the state, $u_k \in \mathbf{R}^{n_u}$ the input and $w_k \in \mathbf{R}^{n_w}$ an indepent identically distributed copy of a random vector w distributed according to $\mathbb{P}_w$. By w^i we denothe the i'th element of w. From earlier results in literature [1] it is clear that only the mean $\mu = \mathbb{E}[w]$ and the covariance $\Sigma = \mathbb{E}[ww^\top]$ affect the optimal controller. As such we estimate these from M offline samples using $\hat{\mu} = \sum_{n=0}^{M} w_n$ and $\hat{\Sigma} = \sum_{n=0}^{M}(w_n - \hat{\mu})(w_n - \hat{\mu})^\top$. By combining results from [3, 4, 2] we then produce an ambiguity set:

$$\mathscr{A} = \left\{ \mathbb{P}_v \in \mathscr{M} \;\middle|\; \begin{array}{l} (\mathbb{E}_{\mathbb{P}_v}[v] - \hat{\mu})^\top \hat{\Sigma}^{-1} (\mathbb{E}_{\mathbb{P}_v}[v] - \hat{\mu}) \leq r_\mu^2 \\ \mathbb{E}_{\mathbb{P}_v}\left[(w-\mu)(w-\mu)^\top\right] \leq r_\Sigma \hat{\Sigma}, \end{array} \right\},$$

with $\mathbb{P}_v$ the distribution of v and $\mathscr{M}$ the set of probability measures defined on $(\mathbf{R}^{n_w}, \mathscr{B})$, with $\mathscr{B}$ the Borel σ-algebra of $\mathbf{R}^{n_w}$. The radii r_μ and r_Σ are calculated such that $\mathbb{P}[\mathbb{P}_w \in \mathscr{A}] \geq 1-\beta$, with β some tunable parameter. To do so we need to assume that w is a sub-Gaussian random vector. The distributionally robust LQR problem is then:

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

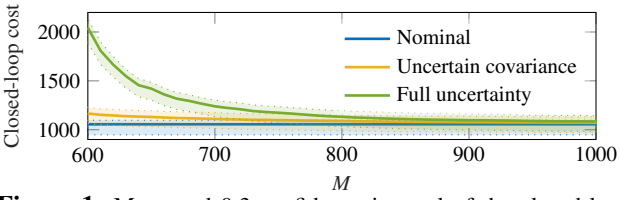

Figure 1: Mean and 0.3-confidence interval of the closed-loop cost versus sample size for controllers.

with $Q > 0$, $R > 0$ and v_k i.i.d. random vectors with distribution $\mathbb{P}_v$. When $r_\mu = 0$ it is possible to show that the solution to (1) is equal to the solution of the classical LQR problem, where we assume that the covariance of w is given by $r_\Sigma \hat{\Sigma}$. This solution is evaluated by solving a semidefinite program, which is a generalization of the one derived in [5]. For the case where $r_\mu > 0$ it is only possible to solve (1) suboptimally by using results from robust optimization [6]. We also verify that both of these controllers are mean square stabilizing with probability at least $1-\beta$.

Since the estimates $\hat{\mu}$ and $\hat{\Sigma}$ become more accurate as more data is gathered, r_μ and r_Σ decrease and therefore the controller converges to the optimal controller for the true distribution of w. This behaviour was confirmed experimentally by calculating the optimal controllers for a double integrator with stochastically varying parameters. The result is depicted in Figure 1.

3 Acknowledgements

This work was supported by the Ford KU Leuven Research Alliance; the Research Foundation Flanders (FWO) PhD grant 11E5520N and research projects 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

[1] T. Morozan. Stabilization of some stochastic discrete-time control systems. *Stoch. Anal. Appl.*, 1(1):89–116, 1983.

[2] E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.*, 58(3):595–612, 2010.

[3] D. Hsu, S. Kakade, and T. Zhang. Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electron. Commun. Probab.*, 17(52):1–13, 2012.

[4] D. Hsu, S. M. Kakadey, and T. Zhang. Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electron. Commun. Probab.*, 17(14), 2012.

[5] V. Balakrishnan and L. Vandenberghe. Semidefinite programming duality and linear time-invariant systems. *IEEE Trans. on Autom. Control*, 48(1):30–41, 2003.

[6] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. Robustness. In *Handbook of Semidefinite Programming*, pages 139–162. 2000.