Risk-sensitive Model Predictive Control for Reproducing Uncertain Dynamical Systems

José Ramón Medina, Klas Kronander and Aude Billard

I. Introduction

In recent years, the application of probabilistic methods in robotics has increased substantially. For instance, in the learning from demonstration paradigm, tasks models are obtained from human demonstrations as uncertain dynamical systems by means of probabilistic estimators [2]. In this particular context, the amount of uncertainty is crucial: it encodes potential task constraints or represents low confidence due to insufficient data. However, the low level control schemes used for task reproduction typically neglect uncertainty by considering only expected actions. To avoid this issue, risk-sensitive control [1] is an interesting alternative that considers not only expected values but also high-order statistics of the control problem.

In this work we study the application of risk-sensitive control for reproducing uncertain dynamical systems learned from demonstrations. To increase robustness, we explore a receding horizon implementation. We present a simple interactive simulation of our approach available online.

II. Problem setting and approach

We consider a task given by an autonomous dynamical system $\dot{\boldsymbol{\xi}}_d = \boldsymbol{f}(\boldsymbol{\xi}_d)$, where $\boldsymbol{\xi}_d$ represents a generalized state variable, e.g. joint angles or Cartesian position. A probabilistic approximation of the task

$$\mathsf{P}\left(\dot{\boldsymbol{\xi}}_d | \boldsymbol{\xi}_d\right) \tag{1}$$

is given by applying a probabilistic method, e.g. Gaussian Mixture Regression or a Gaussian Process [2]. The robot manipulator's rigid body dynamics are

$$M(\boldsymbol{\xi})\ddot{\boldsymbol{\xi}} + C(\dot{\boldsymbol{\xi}}, \boldsymbol{\xi}) + \boldsymbol{g}(\boldsymbol{\xi}) = \boldsymbol{u} \ , \tag{2}$$

where $\boldsymbol{\xi}, \boldsymbol{u}$ are the state and the control input respectively.

The problem consists of reproducing (1) with robot dynamics (2), which, from an optimality perspective corresponds to finding the controls $\boldsymbol{u}(\cdot)$ that minimize the quadratic tracking objective

$$J = \boldsymbol{e}^{\mathsf{T}}(t_f) Q_f \boldsymbol{e}(t_f) + \int_{t=0}^{t_f} \left(\boldsymbol{e}^{\mathsf{T}}(t) Q \boldsymbol{e}(t) + \boldsymbol{u}^{\mathsf{T}}(t) R \boldsymbol{u}(t)\right) dt \tag{3}$$

where $\boldsymbol{e} = \begin{bmatrix} \boldsymbol{\xi}_d \\ \dot{\boldsymbol{\xi}}_d \end{bmatrix} - \begin{bmatrix} \boldsymbol{\xi} \\ \dot{\boldsymbol{\xi}} \end{bmatrix}$, $\boldsymbol{\xi}_d(0) = \boldsymbol{\xi}(0)$ and t_f is the time horizon.

Due to the uncertainty of (1), performance (3) is a random variable. This *stochastic* optimal control problem is typically solved considering the expectation of (3), which neglects higher order cost statistics. To explicitly account for uncertainty, we perform a risk-sensitive optimization, i.e.

$$\min_{\boldsymbol{u}(\cdot)} \theta^{-1} \log \mathsf{E}\left[\exp\{\theta J\}\right] \approx \mathsf{E}\left[J\right] + \theta \mathsf{Var}\left[J\right] \ ,$$

and the influence of uncertainty in the solutions is determined by the risk-sensitivity parameter θ. Due to the nonlinearities from (1), we use a locally optimal solver [3]. In addition, we explore the application of this control scheme in a receding horizon fashion, where solutions are recomputed at each time step and the time horizon is shifted accordingly. This implementation has the advantage of being robust to external perturbations and model errors while, at the same time, it enables a time independent (closed loop) control law.

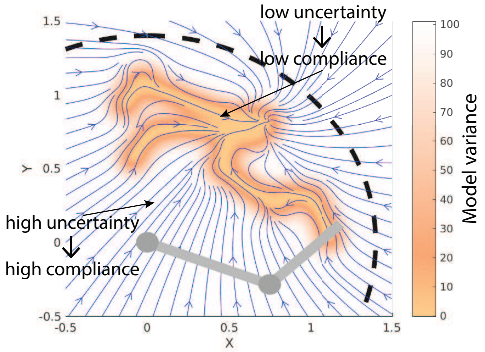

Fig. 1: Simulated 2-link robot tracking a learned dynamical system

III. Preliminary results

A simplified setting of our approach is available for simulation at https://github.com/epfl-lasa/icra-lfd-tutorial (Exercise 4). As the robot tracks the learned dynamics, under high uncertainty, the robot becomes more compliant while it becomes less compliant under low uncertainty as depicted in Fig. 1. We are currently evaluating the application of this approach in a KUKA LWR manipulator.

ACKNOWLEDGMENTS

This research is partly supported by the European Unions Horizon 2020 Research and Innovation programme (call: H2020-ICT-2014-1, RIA) under grant agreement No 643950.

References

[1] D. Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Autom. Control*, 18(2):124 – 131, apr 1973.

[2] S. M. Khansari-Zadeh and Aude Billard. Learning Stable Non-Linear Dynamical Systems with Gaussian Mixture Models. *IEEE Transaction on Robotics*, 2011.

[3] J. R. Medina and S. Hirche. Uncertainty-dependent optimal control for robot control considering high-order cost statistics. In *IEEE/RSJ IROS*, pages 3995–4002, Sept 2015.