Model predictive safety certificates from data for learning-based control

Kim P. Wabersich and Melanie N. Zeilinger, ETH Zurich

Abstract

The control of complex systems faces a trade-off between high performance and safety guarantees, which in particular restricts the application of learning-based methods to safety-critical systems. A framework to address this issue was recently introduced in [1], which uses a safety controller, that guarantees to keep the system within a safe region of the state space.

We investigate the problem of approximate safe set and safe controller calculations based on concepts similar to robust model predictive control (MPC). For the safe control law, we propose to repeatedly solve a nominal finite horizon optimal control problem based on the latest belief about the system dynamics and the current system state. The optimal control problem is constrained such that the final system state must be contained in a terminal region, which is known to be safe. By employing an auxiliary controller that compensates belief errors in the system dynamics, we guarantee that the system tracks the belief-based optimal trajectory towards the safe terminal region using similar arguments as in tube-based MPC. Our analysis reveals that the set of states, from which we can find a finite horizon optimal control solution towards the safe terminal set implies itself a safe set. Since the offline computation of the feasible set of states is intractable except for small-scale systems, we provide an online model predictive scheme that estimates the safe set 'on-the-fly'. The scheme provides an implicit safe control law through the solution of the finite horizon optimal control problem, given the current system state.

The main advantages of this concept are as follows: It relies on mature MPC algorithms and implementations that are also efficient in case of larger scale and nonlinear systems. Further, a larger optimal control horizon leads to a larger set of feasible states and therefore determines the size of the safe set. It can thus be scaled up or scaled down, depending on computational resources.

In [2] we propose a data-driven method that allows the efficient calculation of a possibly large safe set by relying on approximations based on convex optimization problems and Lipschitz continuity of an unknown nonlinear term in the system dynamics. This method is well-suited for determining a data dependent safe terminal set for the model predictive safe set scheme.

Future work will investigate efficient computation of the tube-based controller using a learned model of the dynamics and the extension to distributed nonlinear systems.

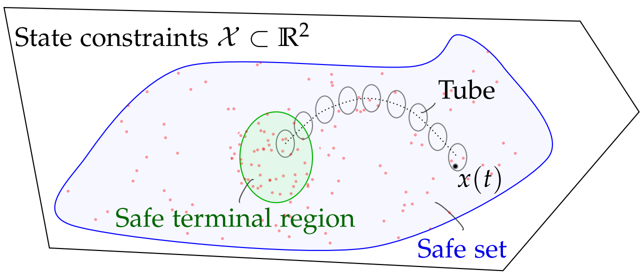

Figure 1: Model predictive safe set (blue) based on available data points (red): Belief-based safe optimal trajectory, starting nearby $x(t)$ and resulting in the safe terminal set with uncertainty tubes around it that will contain the real system state.

References

[1] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *arXiv preprint arXiv:1705.01292*, 2017.

[2] K. P. Wabersich and M. N. Zeilinger. Scalable synthesis of safety certificates from data with application to learning-based control. *arXiv preprint arXiv:1711.11417*, 2017.