Context-aware learning of low-dimensional stabilizing controllers in the scarce data regime

S. W. R. Werner[1] and B. Peherstorfer[1]

[1] *Courant Institute of Mathematical Sciences, New York University, New York, USA (steffen.werner@nyu.edu, pehersto@cims.nyu.edu).*

Stabilizing dynamical systems in science and engineering is challenging, especially in edge cases and limit states where typically little data are available. In this work, we propose a data-driven approach that guarantees finding stabilizing controllers from as few data samples as the dimension of the unstable dynamics, which typically is orders of magnitude lower than the state dimension of systems. The key is learning stabilizing controllers directly from data without learning (reduced) models of the systems, which would require larger numbers of data points.

The starting point for us is the concept of data informativity, which states that the direct construction of stabilizing state-feedback controllers without an intermediate model describing the dynamics is possible using as many data samples as the dimension of the observed states. We sharpened the result in [1] such that the sample complexity scales with the intrinsic (minimal) dimension of the system, rather than the dimension of the observed states, which reduces the sample complexity by several orders of magnitude. In this presentation, we build on the previous findings but go a step further by proposing a data-driven approach that guarantees finding stabilizing controllers from as few data samples as the dimension of the unstable dynamics, which typically is orders of magnitude lower than the dimension of the system. Numerical experiments with systems from chemical reactors (Figure 1) to power systems to fluid dynamics behind obstacles demonstrate that the proposed approach stabilizes systems after observing fewer than five data samples even though the dimension of observed states of the systems is up to several tens of thousands, and learning the corresponding models as well as model-free reinforcement learning with policy gradient methods requires orders of magnitude more data points.

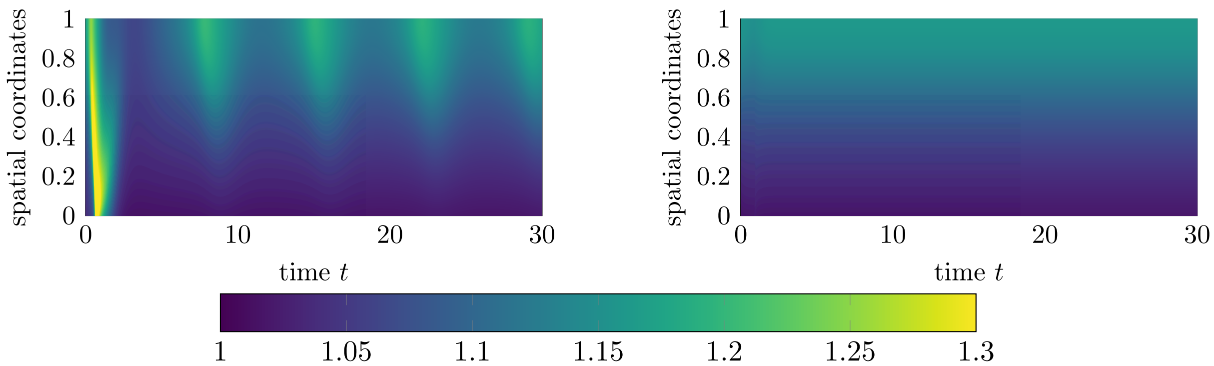

Figure 1: Temperature profiles of a tubular reactor model with unstable oscillations after perturbation of the initial state (right) and the desired steady state behavior (left). Our approach stabilizes the reactor after observing three data samples, which is a three orders of magnitude lower number of samples than required for learning a model of the dynamics.

References

[1] S. W. R. Werner and B. Peherstorfer. On the sample complexity of stabilizing linear dynamical systems from data. e-print 2203.00474, arXiv, 2022. Optimization and Control (math.OC).