Lecture 17

Learning Control Contraction Metrics

In the previous lecture, we explored how learning can be used to find a Lyapunov function and corresponding controller via mixed-integer programming. Approaches have also been developed to find control contraction metrics from data using neural networks.

Let's review the control contraction metric conditions:

$$B_\perp^T(-\dot{W} + AW + WA^T + 2\lambda W)B_\perp < 0 \tag{17.1}$$

where $B_\perp B = 0$ and

$$\frac{\partial W}{\partial x_i} b_i - \frac{\partial b_i}{\partial x} W - W \frac{\partial b_i}{\partial x}^T = 0, \quad \forall i \tag{17.2}$$

where b_i is the i^{th} column of B. Can we parameterize W as a neural network?

17.1. Learning CCMs from Data

To learn a CCM from data, we specify a composite loss function over the domain $B = \{x_1, x_2, ...x_{N_b}\}$. The first loss function L_1 promotes satisfaction of the CCM condition on the training data and is given as

$$L_1(B) = -\alpha_1 \frac{1}{N_b} \sum_{i=1}^{N_b} \exp\left[\max(\bar{\lambda}(C_1(x_i)), \tau)\right]. \tag{17.3}$$

The second loss ensures that the CCM is well conditioned and is given as

$$L_2(B) = \alpha_2 \frac{1}{N_b} \sum_{i=1}^{N_b} \sqrt{\frac{\bar{\lambda}(M(x_i))}{\underline{\lambda}(M(x_i))}}. \tag{17.4}$$

The third loss encourages large contraction rates and is given as

$$L_3(B) = -\alpha_3 \min_{1 \le i \le N_b} \lambda(x_i). \tag{17.5}$$

Here α_1 and α_2 are tunable parameters and τ is a threshold parameter. $\bar{\lambda}$ and $\underline{\lambda}$ signify the maximum and minimum eigen values.

How can we verify over the whole region of inerest?

17.2. Extreme Value Theory

One way to provide stochastic guarantees is the Fisher-Tippet-Gnedenko Theorem. This theorem guarantees that sample batch maximums of random variable Z converge to the Generalized Extreme Value Distribution.

The maximum can then be estimated (e.g., using Maximum Likelihood Estimation (MLE)) with confidence ψ.

For our scenarios, assume $|f(x,u) - \hat{f}(x,u)| \leq w(x,u)$.

$$Z \sim \max_k \Big(|f(x,u) - \hat{f}(x,u)| - w(x,u) \Big)_k \tag{17.6}$$

Use MLE to fit a sample of batch maximums of Z to the Generalized Extreme Value (GEV) distribution:

$$H(y) = \exp\left\{ -\left[1 + \xi \left(\frac{y-\mu}{\sigma} \right) \right]^{-\frac{1}{xi}} \right\}, \quad -\infty < \mu, \quad \xi < \infty, \quad \sigma > 0 \tag{17.7}$$

where μ is the location parameter, ξ is the shape parameter, and $\sigma > 0$ is the scale parameter. If $\xi < 0$, the maximum of Z is finite and can be estimated with confidence ψ. This maximum value is given as ϕ and can provide a robust bound as follows:

$$|f(x,u) - \hat{f}(x,u)| \leq w(x,u) + \phi. \tag{17.8}$$

Bibliography

[1] Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *Conference on Robot Learning*, pages 1519–1539. PMLR, 2021.

[2] Craig Knuth, Glen Chou, Necmiye Ozay, and Dmitry Berenson. Planning with learned dynamics: Probabilistic guarantees on safety and reachability via lipschitz constants. *IEEE Robotics and Automation Letters*, 6(3):5129–5136, 2021.