Stability and Learning Performance for Gaussian Process Control


Armin Lederer
armin.lederer@tum.de


Jonas Umlauft
jonas.umlauft@tum.de

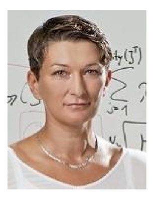
Sandra Hirche
hirche@tum.de

Motivation

Systems without analytic description require data-based modeling.

Classical system identification insufficient for complex nonlinear dynamics [6]
$\Rightarrow$ Flexibility of nonparametric models grows with available data

Closed-loop data gathering for unstable and constrained dynamical systems [1]
$\Rightarrow$ Ensure safety during closed-loop learning

How are formal guarantees provided for control on data-driven models?
How does increasing the number of training samples impact control performance?

Gaussian Process Regression

- Bayesian nonparametric modeling as "distribution over functions" [5]

$$f_{\mathcal{GP}}(\boldsymbol{x}) \sim \mathcal{GP}(m(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}'))$$

- Based on training data $\mathcal{D} = \{\boldsymbol{x}^{(i)}, y^{(i)} = f(\boldsymbol{x}^{(i)}) + \epsilon\}_{i=1}^{N}$, it provides mean and variance

$$\mu_N(\boldsymbol{x}) := \mathbb{E}\left[f_{\mathcal{GP}}(\boldsymbol{x})|\boldsymbol{x}, \mathcal{D}\right] = m(\boldsymbol{x}) + \boldsymbol{k}^\intercal(\boldsymbol{K} + \sigma_n^2\boldsymbol{I}_N)^{-1}(\boldsymbol{y} - m(\boldsymbol{x}^{(1:N)}))$$
$$\sigma_N^2(\boldsymbol{x}) := \mathbb{V}\left[f_{\mathcal{GP}}(\boldsymbol{x})|\boldsymbol{x}, \mathcal{D}\right] = k(\boldsymbol{x}, \boldsymbol{x}) - \boldsymbol{k}^\intercal(\boldsymbol{K} + \sigma_n^2\boldsymbol{I}_N)^{-1}\boldsymbol{k}$$

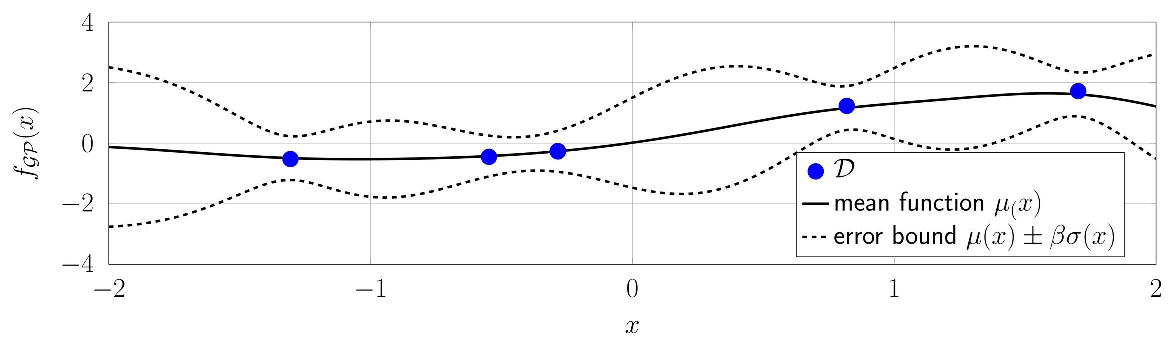

Stability Analysis using Learned Cost Functions [2]

- Infinite horizon cost with discount γ and stage cost $l(\cdot)$ as Lyapunov function

$$V(\boldsymbol{x}) = \sum_{k=0}^{\infty} \gamma^k l(\boldsymbol{f}^k(\boldsymbol{x}))$$

- Approximated cost $\tilde{V}(\boldsymbol{x})$ through Gaussian process regression with kernel

$$\tilde{k}(\boldsymbol{x}, \boldsymbol{x}') = k(\boldsymbol{x}, \boldsymbol{x}') - \gamma k(\boldsymbol{x}, \boldsymbol{f}(\boldsymbol{x}')) - \gamma k(\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}') + \gamma^2 k(\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{f}(\boldsymbol{x}'))$$

- Parallelized interval analysis of $\tilde{V}(\boldsymbol{f}(\boldsymbol{x})) - \tilde{V}(\boldsymbol{x}) < 0$ for Lyapunov decrease region
- Region of attraction as largest level set of $\tilde{V}(\boldsymbol{x})$ contained in Lyapunov decrease region

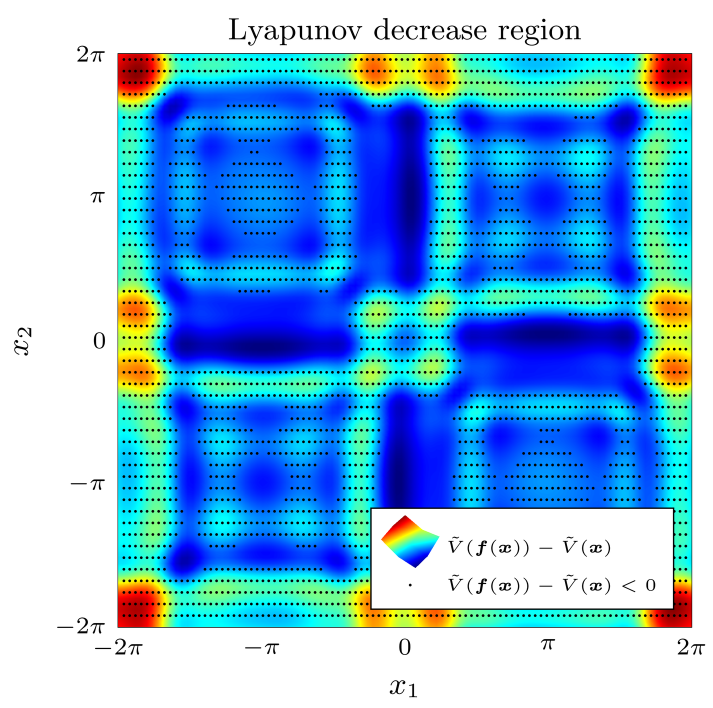

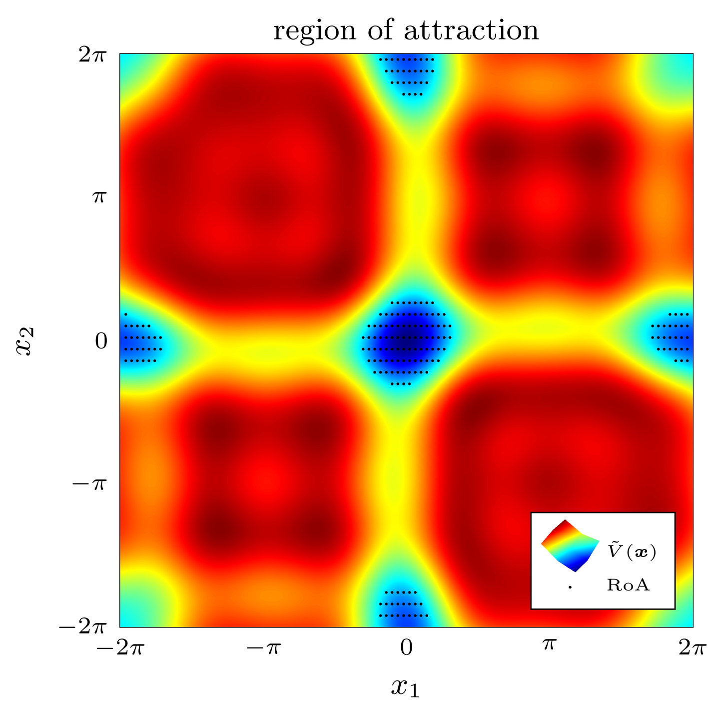

Learned infinite horizon cost function $\tilde{V}(\boldsymbol{x})$ satisfies Bellman equation at training points.

Learning Rate of Gaussian Process Regression [3]

- Informative training points in proximity of test point $\boldsymbol{x}$
- Bound GP posterior variance at $\boldsymbol{x}$ by considering m training points with distance less than ρ

$$\sigma^2(\boldsymbol{x}) \leq k(0) - \frac{k^2(\rho)}{k(0) + \sigma_n^2/m}$$

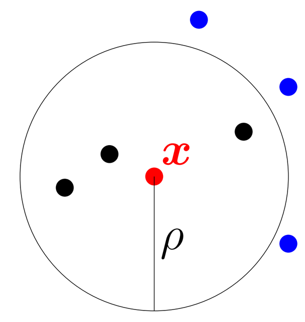

Theorem

For training data drawn from probability density $p(\cdot)$ such that there exists $\rho(N)$ and $c, \epsilon \in \mathbb{R}_+$ with

$$\lim_{N\to\infty} \rho(N) = 0$$

$$\int_{\boldsymbol{x}'\in\mathbb{R}^d:\|\boldsymbol{x}-\boldsymbol{x}'\|\leq\rho(N)} p(\boldsymbol{x}')\mathrm{d}\boldsymbol{x}' \geq cN^{-1+\epsilon},$$

the posterior variance satisfies

$$\lim_{N\to\infty} \sigma_N^2(\boldsymbol{x}) = 0 \quad \text{a.s.}$$

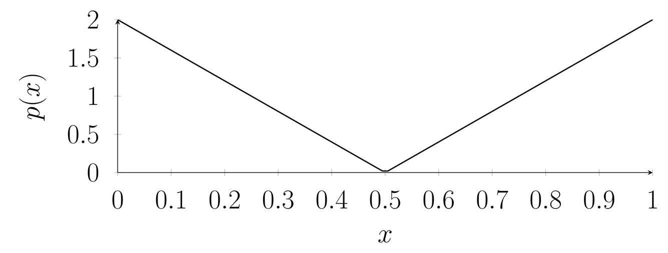

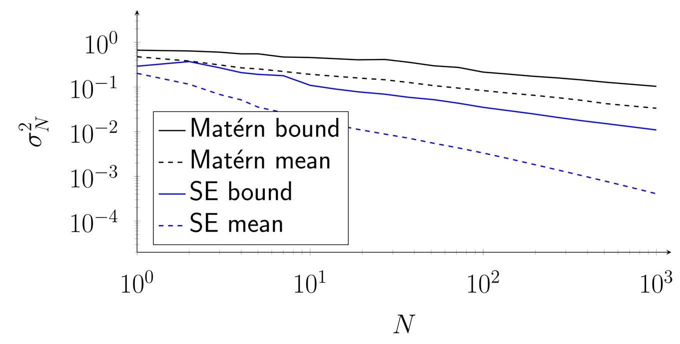

Stable Control through Uniform Error Bounds [4]

- Assumption: dynamical system $f(x, u)$ is a sample from a GP with Lipschitz constant L_f
- Lipschitz continuous posterior mean $\mu_N(\cdot)$ and standard deviation $\sigma_N(\cdot)$ with

$$\|\mu_N(x) - \mu_N(x')\| \leq L_\mu\|x - x'\| \qquad \|\sigma_N(x) - \sigma_N(x')\| \leq \omega_\sigma(\|x - x'\|)$$

Theorem

The learning error is probabilistically bounded by

$$P\left(|f(x) - \mu(x)| \leq \sqrt{\log\left(\frac{\frac{r}{\tau}+1}{\delta}\right)}\sigma_N(x) + (L_\mu + L_f)\tau + \omega_\sigma(\tau),\ \forall \boldsymbol{x} \in \mathbb{X}\right) \geq 1 - \delta$$

on the set $\mathbb{X}$ with maximal extension r for every $\delta \in (0, 1)$, $\tau \in \mathbb{R}_+$.

$\Rightarrow$ Feedback linearizing controller ensures ultimate boundedness with probability $1 - \delta$

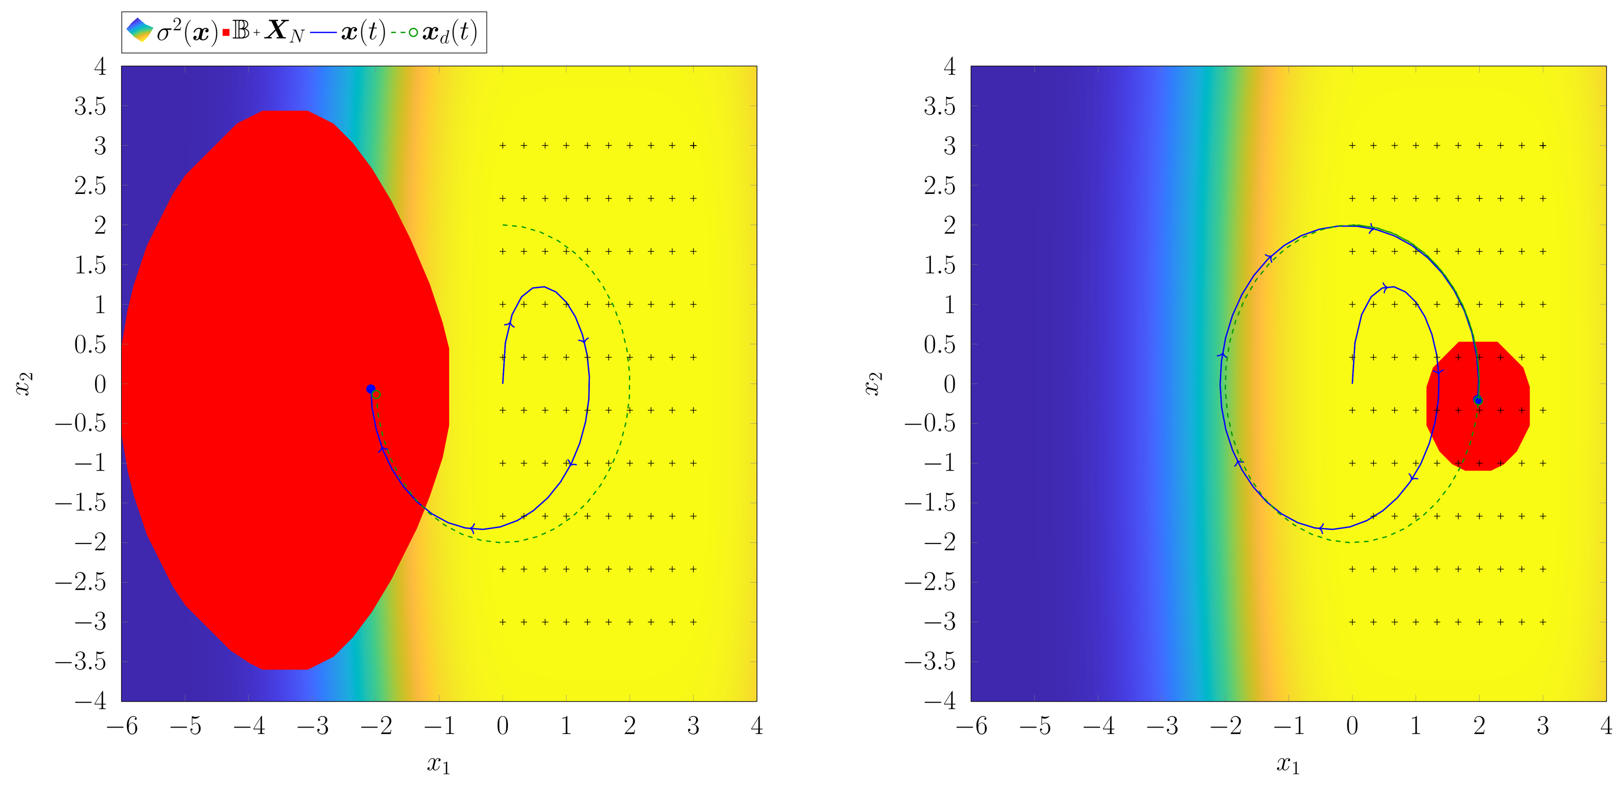

Future Work

- Performance guarantees for closed-loop learning and on-line learning with error bounds
- Optimal and safe exploration of a dynamical system in a task space
- Stable cautious MPC with reliable chance constraints for Gaussian process models
- Sampling-based analysis of models obtained through machine learning

References

[1] E. T. Campolettano, M. L. Bland, R. A. Gellner, D. W. Sproule, B. Rowson, A. M. Tyson, S. M. Duma, and S. Rowson. Ranges of injury risk associated with impact from unmanned aircraft systems. *Annals of Biomedical Engineering*, 45(12):2733–2741, 2017.

[2] A. Lederer and S. Hirche. Local asymptotic stability analysis and region of attraction estimation with Gaussian processes. Submitted to 2019 IEEE Conference on Decision and Control.

[3] A. Lederer, J. Umlauft, and S. Hirche. Posterior variance analysis of Gaussian processes with application to average learning curves. Submitted to 2019 Conference on Neural Information Processing Systems.

[4] A. Lederer, J. Umlauft, and S. Hirche. Uniform error bounds for Gaussian process regression with application to safe control. Submitted to 2019 Conference on Neural Information Processing Systems.

[5] C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, USA, Jan. 2006.

[6] H. Shen. Meet the soft, cuddly robots of the future. *Nature*, 530(7588):24–26, 2016.