Cautious Nonlinear Covariance Steering using Variational Gaussian Process Predictive Models

Alex Tsolovikos and Efstathios Bakolas

Department of Aerospace Engineering and Engineering Mechanics
The University of Texas at Austin

Modeling, Estimation, and Control Conference 2021
Austin, TX, October 2021

Problem Formulation

- **Unknown** stochastic nonlinear dynamics

$$\mathbf{z}_{t+1} = \mathbf{g}(\mathbf{z}_t, \mathbf{u}_t) + \boldsymbol{\epsilon}_t$$

- State measurements from arbitrary control inputs
- Compute a feedback control policy that will steer the stochastic state of the unknown nonlinear system
- Goal: Steer the state **mean** and **covariance** from a given initial distribution to a target terminal one in **finite** time

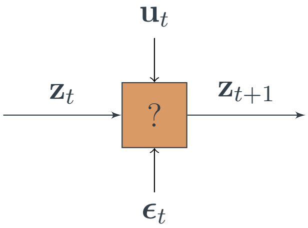

Training trajectories:
$\{\mathbf{z}_t, \mathbf{u}_t, \mathbf{z}_{t+1}\}_{t=1}^{N}$

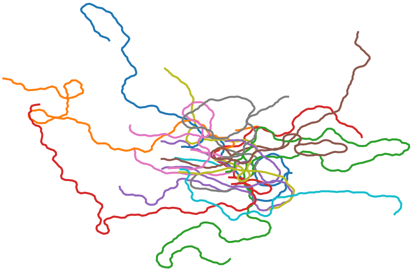

Learning a Model: Gaussian Process Regression

- Non-parametric regression models
- Distributions over unknown functions

$$f(\cdot) : \mathbb{R}^n \to \mathbb{R}$$

- Advantages:
  - Flexible
  - Provide uncertainty estimates (model uncertainties + process noise)
  - Degrade gracefully – they know what they don't know
- Assume $f(\cdot)$ belongs to a family of functions with a Gaussian prior:

$$f(\mathbf{x}) \sim \mathcal{N}\left(f(\mathbf{x} \mid m(\mathbf{x}), k(\mathbf{x}, \mathbf{x})\right)$$

- Mean function $m(\mathbf{x})$:
  - usually zero, constant, or linear
- Kernel function $k(\mathbf{x}, \mathbf{x}')$:
  - measures closeness between two points
  - specifies smoothness and continuity properties of $f(\cdot)$
- N observations at locations $\mathbf{x}_i$

$$y_i = f(\boldsymbol{x}_i) + \epsilon_i, \quad i = 1, \ldots, N$$

with observation likelihood

$$p(y_i \mid f(\mathbf{x}_i)) = \mathcal{N}\left(y_i \mid f(\mathbf{x}_i), \sigma_\epsilon^2\right)$$

Learning a Model: Gaussian Process Regression

- Marginal observation likelihood:

$$p(\mathbf{y};\mathbf{X}) = \mathcal{N}\left(\mathbf{y} \mid m(\mathbf{X}), k(\mathbf{X},\mathbf{X}) + \sigma_\epsilon^2 I\right)$$

- Optimize mean/kernel hyperparameters:

$$\mathbf{\Theta}_{\text{opt}} = \arg\min_{\mathbf{\Theta}} \left(-\log p(\mathbf{y};\mathbf{X})\right)$$

- Prediction:

$$p(y_*;\mathbf{x}_*,\mathbf{y},\mathbf{X}) = \int p(y_*,\mathbf{y};\mathbf{x}_*,\mathbf{X})d\mathbf{y} = \mathcal{N}\left(y_* \mid \mu_*, \sigma_*^2\right)$$

$$\mu_* = m(\mathbf{x}_*) + k(\mathbf{x}_*,\mathbf{X})\left[k(\mathbf{X},\mathbf{X}) + \sigma_\epsilon^2 I\right]^{-1}(\mathbf{y} - m(\mathbf{X}))$$
$$\sigma_*^2 = k(\mathbf{x}_*,\mathbf{x}_*) - k(\mathbf{x}_*,\mathbf{X})\left[k(\mathbf{X},\mathbf{X}) + \sigma_\epsilon^2 I\right]^{-1} k(\mathbf{X},\mathbf{x}_*)$$

- Inference: invert $N \times N$ matrix – Does not scale to more than a few thousand data points

Scaling GPs to Big Data: Sparse Variational GP*

- Sparse approximation of GPs
- Introduce inducing locations

$$\mathbf{Z} = \begin{bmatrix} \mathbf{z}_1 & \cdots & \mathbf{z}_M \end{bmatrix}^\top, \ \mathbf{u} = f(\mathbf{Z})$$

  where $M \ll N$

- Introduce variational posterior

$$q(\mathbf{f}, \mathbf{u}) = p(\mathbf{f} \mid \mathbf{u}; \mathbf{X}, \mathbf{Z}) q(\mathbf{u}),$$
$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u} \mid \mathbf{m}, \mathbf{S})$$

  where $\mathbf{m}$, $\mathbf{S}$, along with $\mathbf{Z}$, are variational parameters

- Optimize variational parameters and hyperparameters: maximize lower bound $\mathcal{L}$

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{f},\mathbf{u})} \left[\log \frac{p(\mathbf{y}, \mathbf{f}, \mathbf{u})}{q(\mathbf{f}, \mathbf{u})} \right] \leq \log p(\mathbf{y} \mid \mathbf{X})$$

- Lower bound $\mathcal{L}$:
  - Factorizes as a sum over training data
  - Minimize $-\mathcal{L}$ using stochastic gradient descent
  - Scales to big data
- Inference in $\mathcal{O}(M^3)$

*James Hensman, Nicolò Fusi, and Neil D. Lawrence. "Gaussian Processes for Big Data". In: *29th Conference on Uncertainty in Artificial Intelligence*. UAI'13. Bellevue, WA, 2013, pp. 282–290.

System Identification with SVGPs

- Consider stochastic nonlinear dynamics:

$$\mathbf{z}_{t+1} = \mathbf{g}(\mathbf{z}_t, \mathbf{u}_t) + \epsilon_t$$

- Dynamics are unknown, but we have state observations

$$\mathbf{y}_i = \mathbf{g}(\mathbf{z}_i, \mathbf{u}_i) + \epsilon_i$$

at known locations

$$\mathbf{x}_i = [\mathbf{z}_i; \mathbf{u}_i]$$

- Data: $\mathcal{D} = \{\mathbf{y}_i, \mathbf{z}_i, \mathbf{u}_i\}_{i=1}^N$
- Learned dynamics: a multitask SVGP trained on $\mathcal{D}$, such that

$$\boxed{\mathbf{z}_{t+1} = G(\mathbf{z}_t, \mathbf{u}_t) + \mathbf{w}_t}$$

where

$$G(\mathbf{z}_t, \mathbf{u}_t) = \boldsymbol{\mu}_f([\mathbf{z}_t; \mathbf{u}_t])$$

and

$$\mathbf{w}_t \sim \mathcal{N}\left(\mathbf{w}_t \mid 0, \Sigma_f([\mathbf{z}_t; \mathbf{u}_t], [\mathbf{z}_t; \mathbf{u}_t]) + \sigma_\epsilon^2\right)$$

Nonlinear Covariance Steering

- Find an optimal control policy

$$\mathbf{u}_t = \pi_t(\mathbf{z}_t), \quad t = 0, \ldots, T-1$$

that will take the initial state distribution

$$\mathbb{E}[\mathbf{z}_0] = \boldsymbol{\mu}_0, \quad \mathrm{Cov}[\mathbf{z}_0] = \Sigma_0$$

to a terminal distribution

$$\mathbb{E}[\mathbf{z}_T] = \boldsymbol{\mu}_\mathrm{f}, \quad \Sigma_\mathrm{f} - \mathrm{Cov}[\mathbf{z}_T] \succcurlyeq 0$$

in a **finite** number of time steps

- Final state: reached with a guaranteed upper bound on the uncertainty

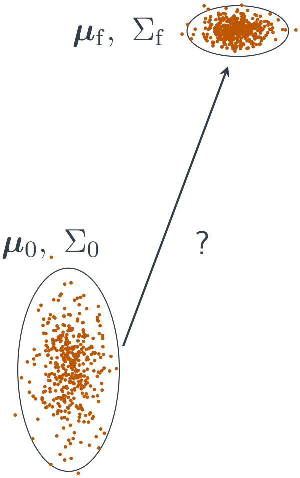

Nonlinear Covariance Steering with SVGPs

- For a **linear dynamical system**

$$\mathbf{z}_{t+1} = A_t \mathbf{z}_t + B_t \mathbf{u}_t$$

the optimal feedback policy

$$\pi_t(\{\mathbf{z}_i\}_{i=0}^t) = \boldsymbol{v}_t + \sum_{i=0}^{t} K_{t,i} \mathbf{z}_i$$

is the solution of a semi-definite program (SDP)

- For the **nonlinear SVGP model** we apply a greedy control algorithm†
- For $t = 0, \ldots, T$:
  1. **Linearize** dynamics around latest state mean $\boldsymbol{\mu}_t$ and corresponding input $\mathbf{u}_t^*$
  2. Solve the **linearized covariance steering** (SDP) from t to T
  3. Estimate the next mean $\boldsymbol{\mu}_{t+1}$ and covariance Σ_{t+1} using the **Unscented Transform**

†E. Bakolas and A. Tsolovikos. "Greedy finite-horizon covariance steering for discrete-time stochastic nonlinear systems based on the unscented transform". In: *2020 American Control Conference (ACC)*. IEEE. 2020, pp. 3595–3600.

Linearization of SVGP Dynamics

- Linearize around $\mathbf{z}_*$, $\mathbf{u}_*$:

$$\mathbf{z}_{t+1} \approx A_* \mathbf{z}_t + B_* \mathbf{u}_t + \mathbf{d}_*$$

where

$$A_* = \frac{\partial}{\partial \mathbf{z}} G(\mathbf{z}_*, \mathbf{u}_*) = \frac{\partial}{\partial [\mathbf{z}; \mathbf{u}]} \boldsymbol{\mu}_f([\mathbf{z}; \mathbf{u}]) \begin{bmatrix} I_{n_z} \\ 0 \end{bmatrix} \Bigg|_{\substack{\mathbf{z}=\mathbf{z}_* \\ \mathbf{u}=\mathbf{u}_*}}$$

$$B_* = \frac{\partial}{\partial \mathbf{u}} G(\mathbf{z}_*, \mathbf{u}_*) = \frac{\partial}{\partial [\mathbf{z}; \mathbf{u}]} \boldsymbol{\mu}_f([\mathbf{z}; \mathbf{u}]) \begin{bmatrix} 0 \\ I_{n_u} \end{bmatrix} \Bigg|_{\substack{\mathbf{z}=\mathbf{z}_* \\ \mathbf{u}=\mathbf{u}_*}}$$

$$\mathbf{d}_* = -A_* \mathbf{z}_* - B_* \mathbf{u}_* + G(\mathbf{z}_*, \mathbf{u}_*)$$

- Use automatic differentiation

Experiments: 4D Nonlinear System

- Consider unicycle car dynamics:

$$\begin{aligned} s_{x,t+1} &= s_{x,t} + v_t \tau \cos\theta_t + \epsilon_t^{s_x} \\ s_{y,t+1} &= s_{y,t} + v_t \tau \sin\theta_t + \epsilon_t^{s_y} \\ \theta_{t+1} &= \theta_t + u_t^{\theta} v_t \tau + \epsilon_t^{\theta} \\ v_{t+1} &= v_t + u_t^{v} \tau + \epsilon_t^{v} \end{aligned}$$

- Assume dynamics are unknown but a black-box simulator is available
- Collect data and run stochastic gradient descent → SVGP Dynamics
- Number of training data: 16000
- Number of inducing locations: 256

- Mean & kernel functions:

$$m(\mathbf{x}) = const.$$
$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_{L^{-1}}^2\right)$$

- Initial distribution:

$$\boldsymbol{\mu}_0 = [0, 0, 0, 1]^\top$$
$$\Sigma_0 = \mathrm{diag}([0.1, 0.2, 0.1, 0.1])^2$$

- Target distribution:

$$\boldsymbol{\mu}_f = [1, 2, 0, 1]^\top$$
$$\Sigma_f = \mathrm{diag}([0.1, 0.05, 0.05, 0.05]^2$$

- Number of time steps: $T = 30$

Experiments: 4D Nonlinear System

- SVGP Model:

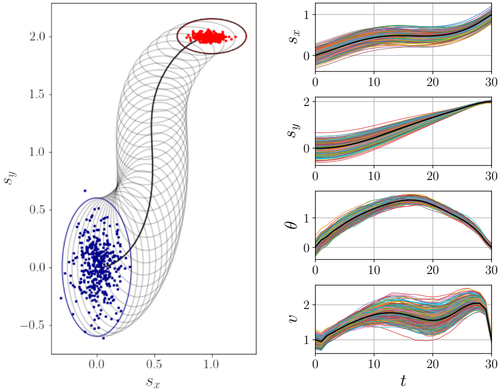

- Exact Model:

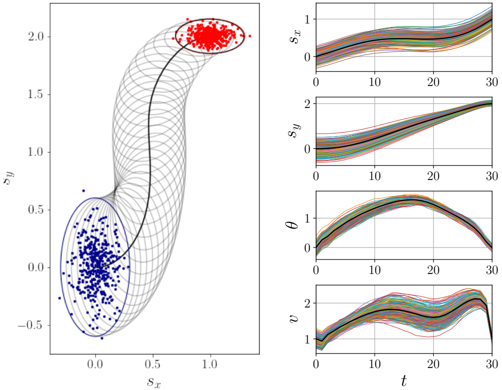

Conclusions and Future Work

- Scalable Gaussian Process predictive models were used for nonlinear covariance steering of an unknown stochastic nonlinear system
- GP models capture both process noise and model uncertainties, leading to "cautious" control policies
- Next steps:
  - Systems with incomplete state measurements
  - Active learning of variational GP models

The code for this work is available at:
`https://github.com/alextsolovikos/greedyGPCS`

The University of Texas at Austin
Aerospace Engineering
and Engineering Mechanics
Cockrell School of Engineering

This work has been supported in part by:

NSF

ONASSIS FOUNDATION