R:SS Workshop on Geometry and Topology in Robotics

Gaussian Processes
on Riemannian Manifolds for Robotics

Alexander Terenin and Viacheslav Borovitskiy

Imperial College London

St Petersburg University

Gaussian Processes

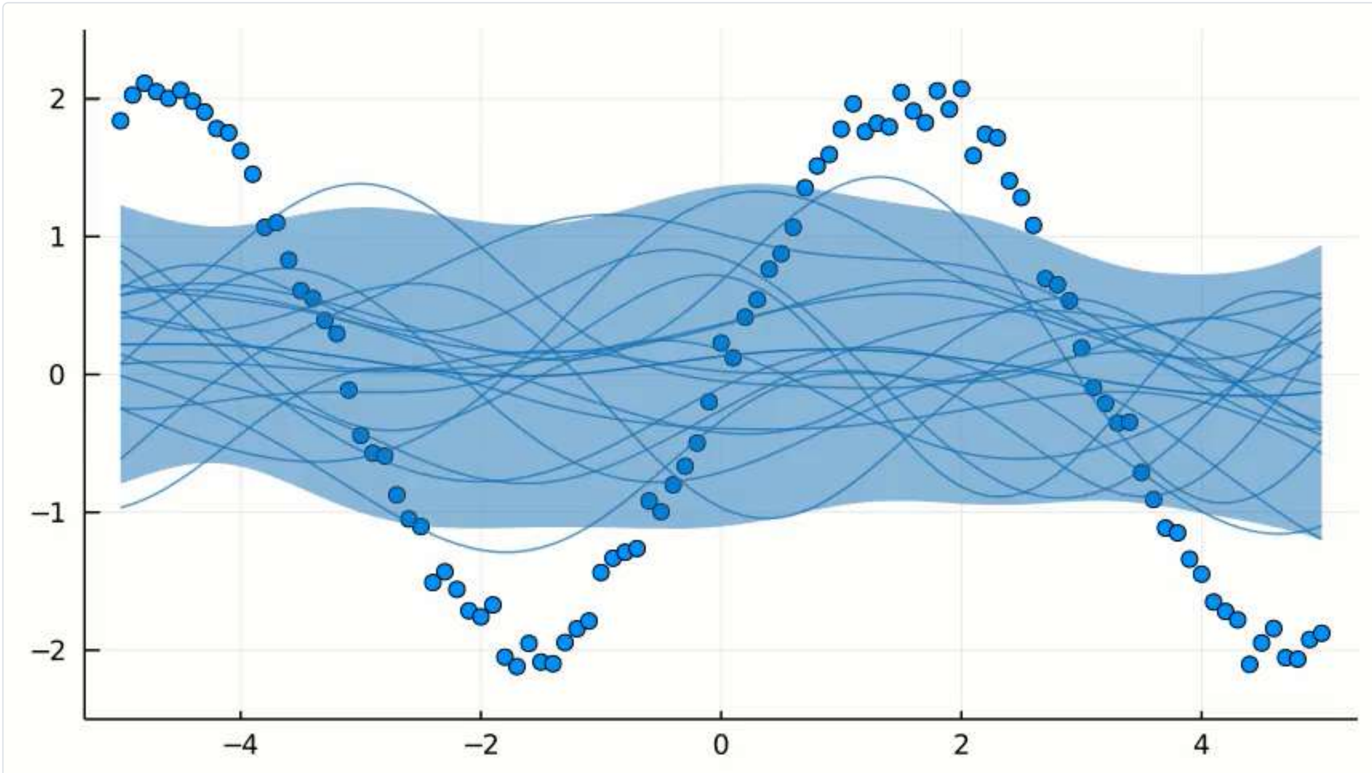

Gaussian Processes

Definition. A *Gaussian process* is random function $f : X \to \mathbb{R}$ such that for any $x_1, .., x_n$, the vector $f(x_1), .., f(x_n)$ is multivariate Gaussian.

Every GP is characterized by a *mean* $\mu(\cdot)$ and a *kernel* $k(\cdot, \cdot)$. We have

$$f(\boldsymbol{x}) \sim \mathrm{N}(\boldsymbol{\mu}_{\boldsymbol{x}}, \mathbf{K}_{\boldsymbol{x}\boldsymbol{x}})$$

where $\boldsymbol{\mu}_{\boldsymbol{x}} = \mu(\boldsymbol{x})$ and $\mathbf{K}_{\boldsymbol{x}\boldsymbol{x}'} = k(\boldsymbol{x}, \boldsymbol{x}')$.

Bayesian learning: $f \mid \boldsymbol{y}$.

Bayesian Optimization

Goal: minimize unknown function ϕ in as few evaluations as possible.

1. Build GP posterior $f \mid \boldsymbol{y}$ using data $(x_1, \phi(x_1)), .., (x_n, \phi(x_n))$.
2. Choose

$$x_{n+1} = \underset{x \in \mathcal{X}}{\arg\max}\, \alpha_{f|\boldsymbol{y}}(x)$$

   to maximize the acquisition function $\alpha_{f|\boldsymbol{y}}$, for instance expected improvement $\alpha_{f|\boldsymbol{y}} = \mathbb{E}_{f|\boldsymbol{y}} \max(0, \min_{i=1,..,n} \phi(x_n) - f(x))$.

Automatic explore-exploit tradeoff.

Bayesian Optimization

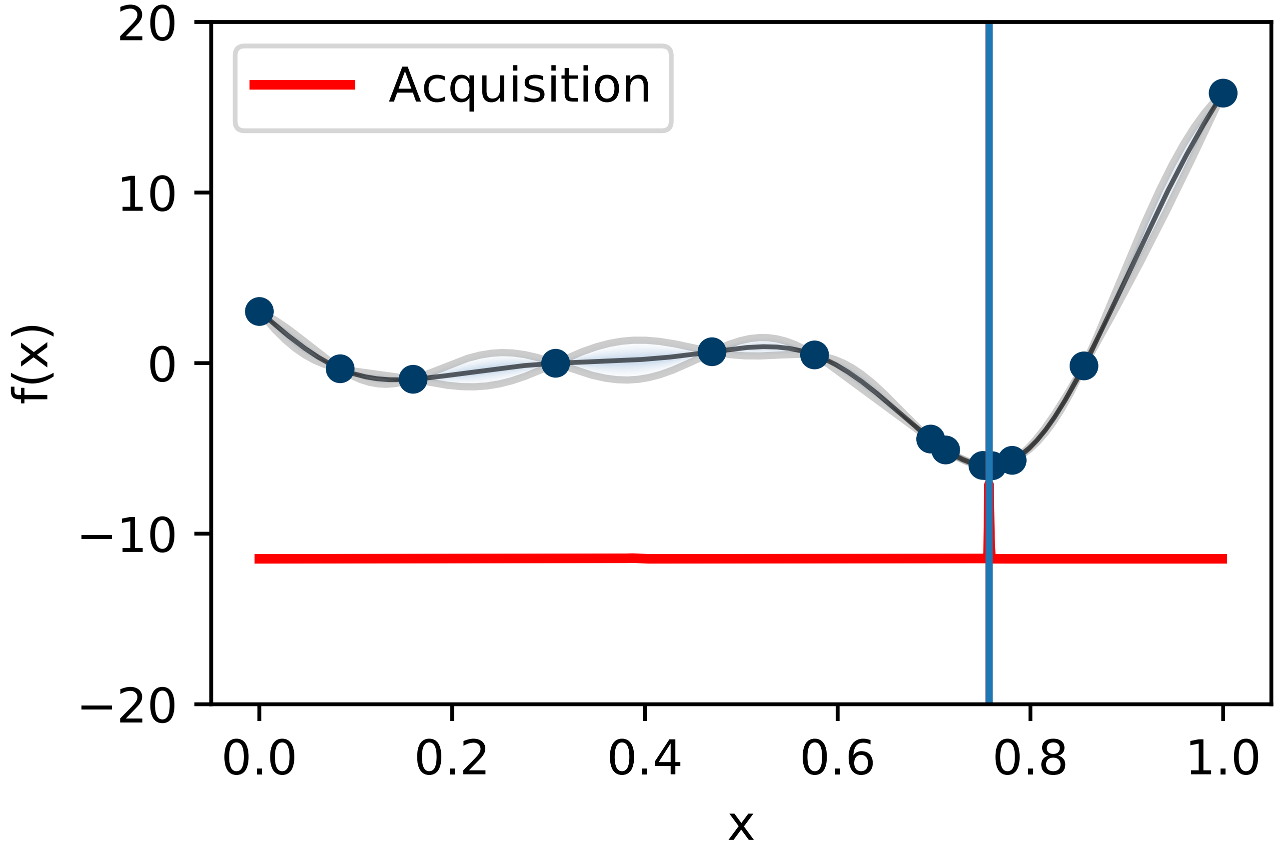

Modeling Dynamical Systems with Uncertainty

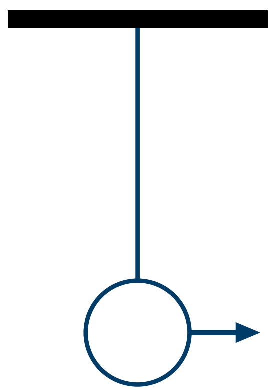

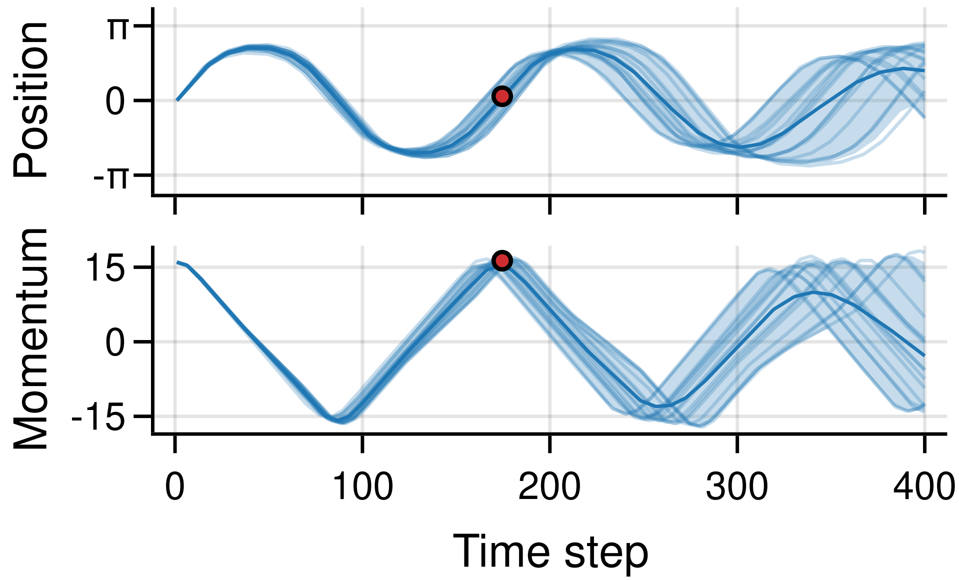

$$x_0 \qquad x_1 = x_0 + f(x_0)\Delta t \qquad x_2 = x_1 + f(x_1)\Delta t \qquad ..$$

Model-based Reinforcement Learning

Video from Deisenroth and Rasmussen (2011)

Geometry-aware Gaussian Processes

Matérn Gaussian Processes

$$k_\nu(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|x - x'\|}{\kappa} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|x - x'\|}{\kappa} \right)$$

σ^2: variance κ: length scale ν: smoothness

$\nu \to \infty$: recovers squared exponential kernel

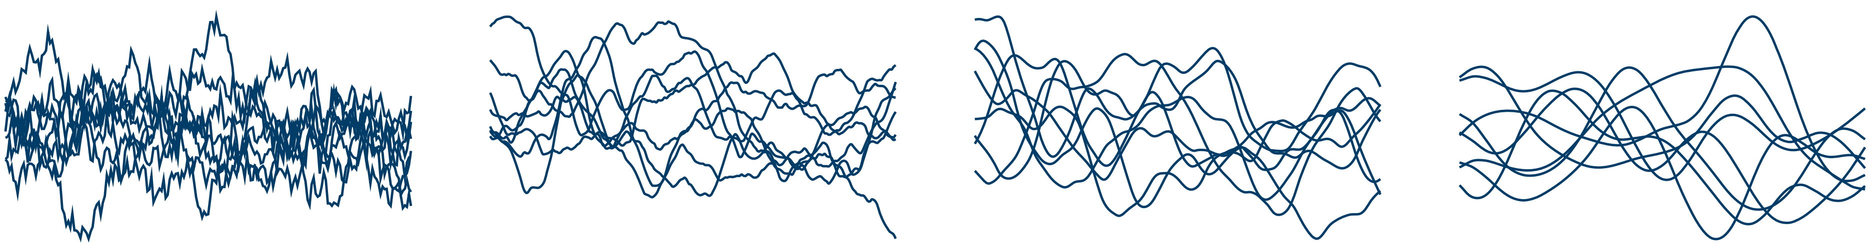

$\nu = 1/2$ $\nu = 3/2$ $\nu = 5/2$ $\nu = \infty$

Riemannian Geometry

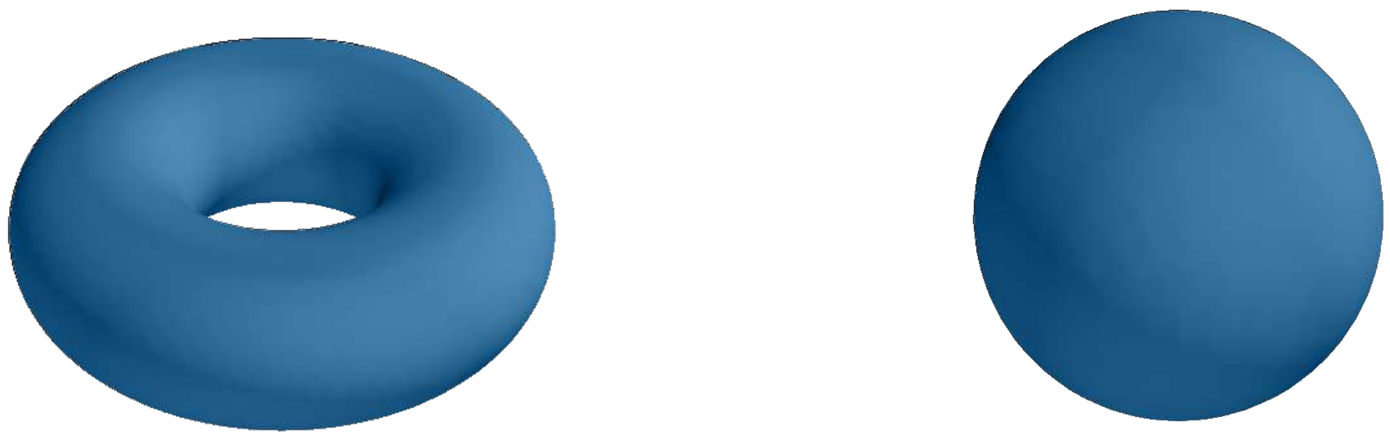

How should Matérn kernels generalize to this setting?

Geodesics

$$k_{\infty}^{(d_g)}(x, x') = \sigma^2 \exp\left(-\frac{d_g(x, x')^2}{2\kappa^2}\right)$$

Theorem. (Feragen et al.) Let M be a complete Riemannian manifold without boundary. If $k_{\infty}^{(d_g)}$ is positive semi-definite for all κ, then M is isometric to a Euclidean space.

Need a different candidate generalization

Feragen et al. (2015)

Stochastic Partial Differential Equations

$$\underbrace{\left(\frac{2\nu}{\kappa^2} - \Delta\right)^{\frac{\nu}{2}+\frac{d}{4}} f = \mathcal{W}}_{\text{Matérn}} \qquad \underbrace{e^{-\frac{\kappa^2}{4}\Delta} f = \mathcal{W}}_{\text{squared exponential}}$$

Δ: Laplacian $\quad$ $\mathcal{W}$: (rescaled) white noise

Generalizes well to the Riemannian setting

Whittle (1963)
Lindgren et al. (2011)

Riemannian Matérn Kernels: compact spaces

$$k_\nu(x, x') = \frac{\sigma^2}{C_\nu} \sum_{n=0}^{\infty} \left(\frac{2\nu}{\kappa^2} - \lambda_n \right)^{\nu - \frac{d}{2}} f_n(x) f_n(x')$$

λ_n, f_n: Laplace–Beltrami eigenpairs

Analytic expressions for circle, sphere, ..

Riemannian Matérn Kernels

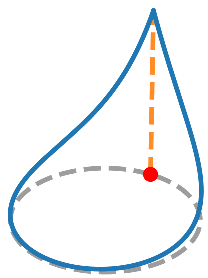

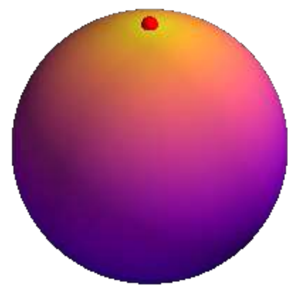

$$k_{1/2}(\bullet, \cdot)$$

Example: regression on the surface of a dragon

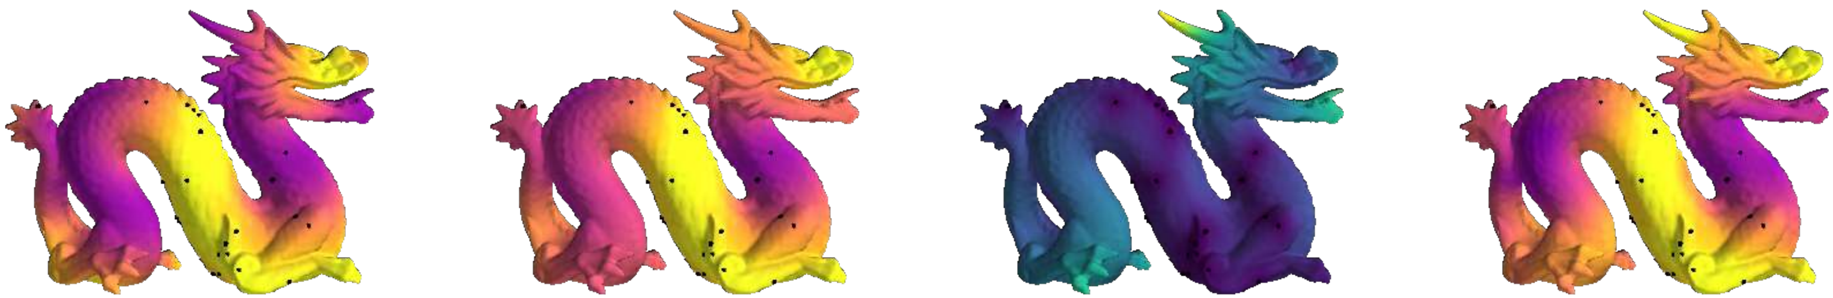

(a) Ground truth (b) Posterior mean (c) Standard deviation (d) Posterior sample

Applications

Pendulum Dynamics: GP on a cylinder

(a) Ground truth

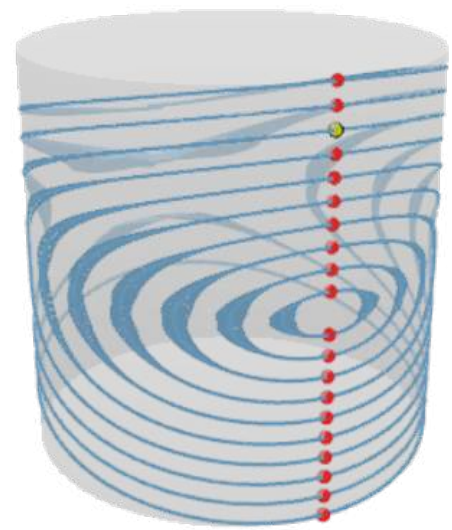

(b) Gaussian process

Bayesian Optimization in Robotics

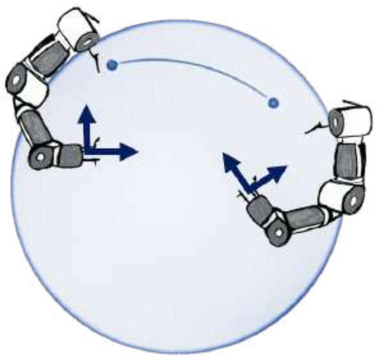

Orientation: sphere

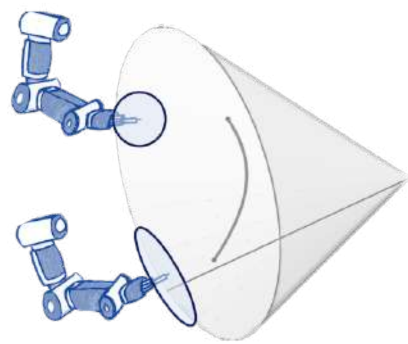

Manipulability:
SPD manifold

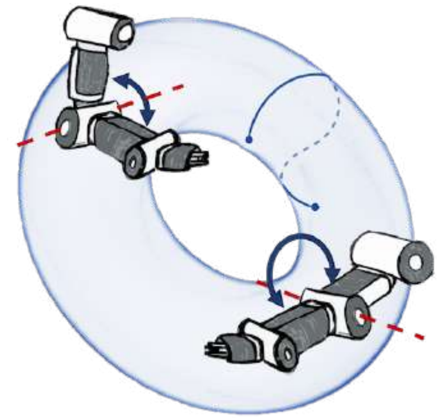

Joint postures: torus

Jaquier et al. (2020)

Matérn Gaussian Processes on Riemannian Manifolds

NEURAL INFORMATION PROCESSING SYSTEMS

   

Viacheslav* Borovitskiy

Alexander* Terenin

Peter* Mostowsky

Marc Peter Deisenroth

St Petersburg University

Imperial College London

UCL

*equal contribution

Matérn Gaussian Processes on Graphs

AISTATS






Viacheslav* Borovitskiy

Iskander* Azangulov

Alexander* Terenin

Peter* Mostowsky

Marc Peter Deisenroth

Nicolas Durrande

St Petersburg University

Imperial College London

UCL

Secondmind

*equal contribution

Thank you!

V. Borovitskiy, I. Azangulov, A. Terenin, P. Mostowsky, M. P. Deisenroth. Matérn Gaussian Processes on Graphs. *Artificial Intelligence and Statistics*, 2021.

V. Borovitskiy, A. Terenin, P. Mostowsky, M. P. Deisenroth. Matérn Gaussian Processes on Riemannian Manifolds. *Advances in Neural Information Processing Systems*, 2020.

St Petersburg University

Imperial College London