Trajectory Optimisation in Learned Multimodal Dynamical Systems via Latent-ODE Collocation

Aidan Scannell[1] | Carl Henrik Ek[2] | Arthur Richards[1]

[1]University of Bristol | [1]Bristol Robotics Lab | [2]University of Cambridge

18th March 2021

bristol.ac.uk

Example

Quadcopters subject to spatially varying turbulence

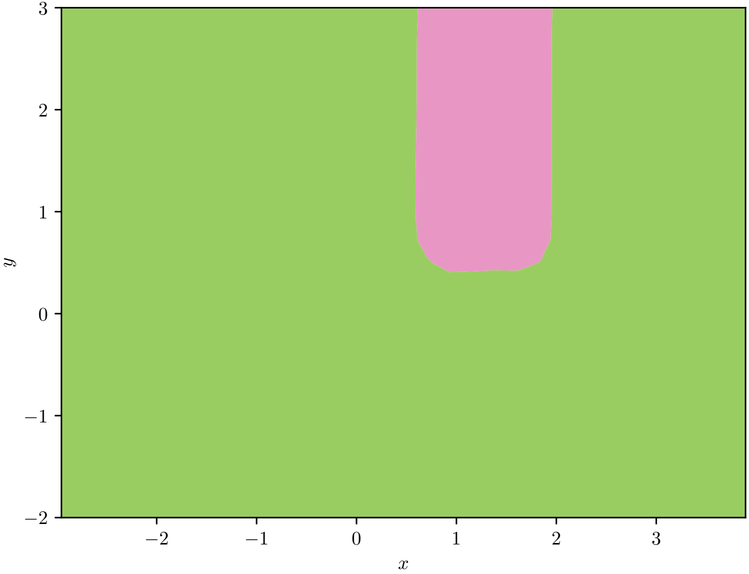
3
2
1
0
−1
−2
y
−2
−1
0
1
2
3
x

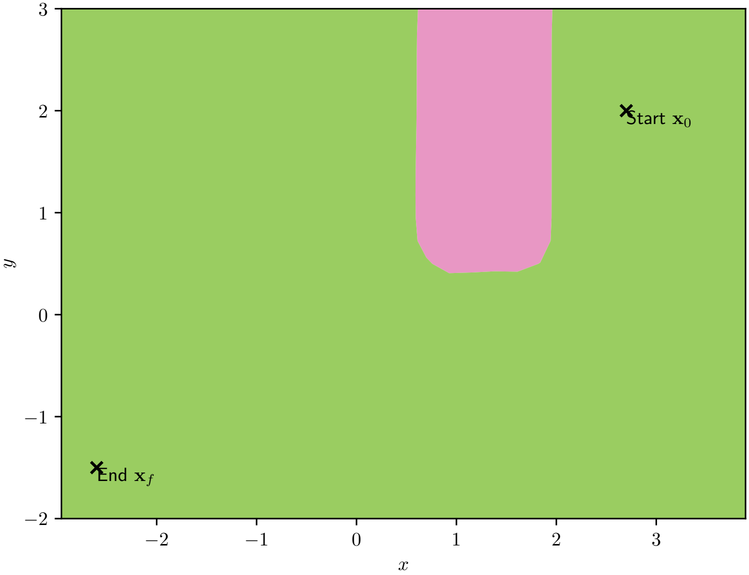

Start $\mathbf{x}_0$
End $\mathbf{x}_f$
y
x
3
2
1
0
−1
−2
−2
−1
0
1
2
3

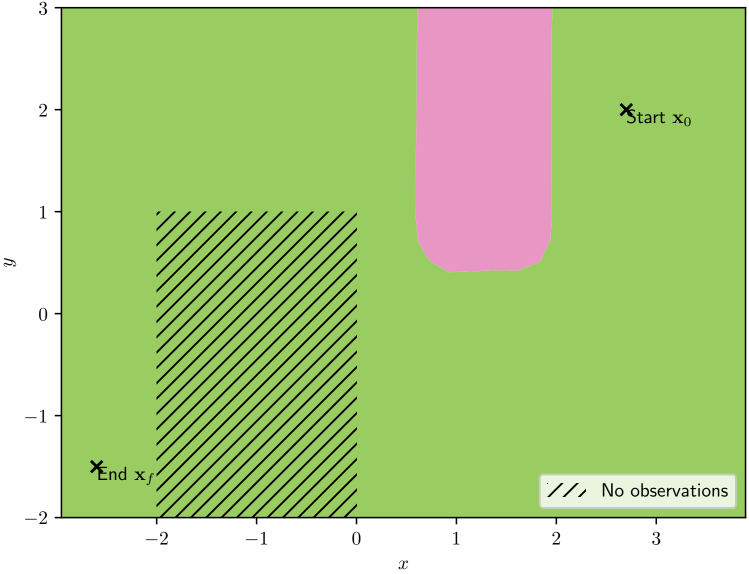

Start $\mathbf{x}_0$
End $\mathbf{x}_f$
No observations
x
y

Quadcopter Scenario

- State $\mathbf{x} \in \mathbb{R}^D$
  - $\mathbf{x} = [x, y]$
- Control $\mathbf{u} \in \mathbb{R}^F$
  - $\mathbf{u} = [\dot{x}, \dot{y}]$

$$\dot{\mathbf{x}}(t) = f^{(k)}(\mathbf{x}(t), \mathbf{u}(t)) + \epsilon^{(k)}(t) \quad \text{if} \quad \alpha(t) = k$$

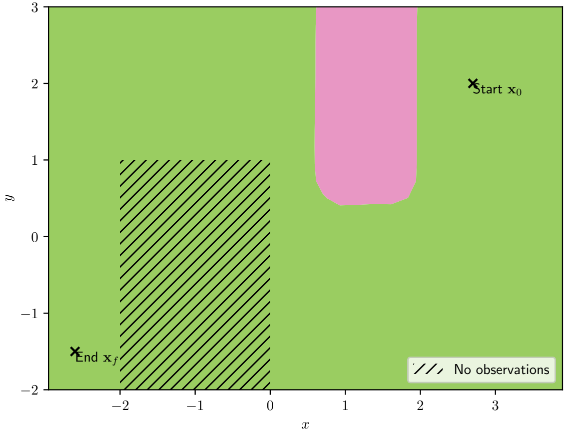

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model

Stage Two - Trajectory Optimisation

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model
- Bayesian inference

Stage Two - Trajectory Optimisation

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model
- Bayesian inference

Stage Two - Trajectory Optimisation

- Geometric objective function

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model
- Bayesian inference

Stage Two - Trajectory Optimisation

- Geometric objective function
- Implicit minimisation

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model
- Bayesian inference

Stage Two - Trajectory Optimisation

- Geometric objective function
- Implicit minimisation
  - Project trajectory optimisation onto Geodesic ODE

Method Overview

Stage One - Model Learning

- Probabilistic transition dynamics model
- Bayesian inference

Stage Two - Trajectory Optimisation

- Geometric objective function
- Implicit minimisation
  - Project trajectory optimisation onto Geodesic ODE
  - Solve with direct collocation

Stage One - Model Learning

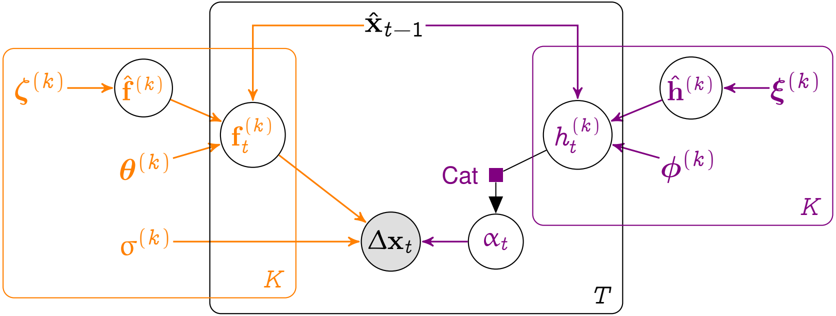

- Marginal likelihood

$$p(\Delta\mathbf{x}_{1:T}|\hat{\mathbf{x}}_{1:T}) = \prod_{t=1}^{T}\sum_{k=1}^{K}\Bigg(\underbrace{\left\langle \Pr\left(\alpha_t = k|\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{h}}\right)\right\rangle_{p(\hat{\mathbf{h}}|\boldsymbol{\xi})}}_{\text{Mixing Probability}}\underbrace{\left\langle p\left(\Delta\mathbf{x}_t|\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{f}}^{(k)}\right)\right\rangle_{p\left(\hat{\mathbf{f}}^{(k)}|\zeta^{(k)}\right)}}_{\text{Dynamics Mode } k}\Bigg)$$

GP Posterior over Desired Mode's Gating Function

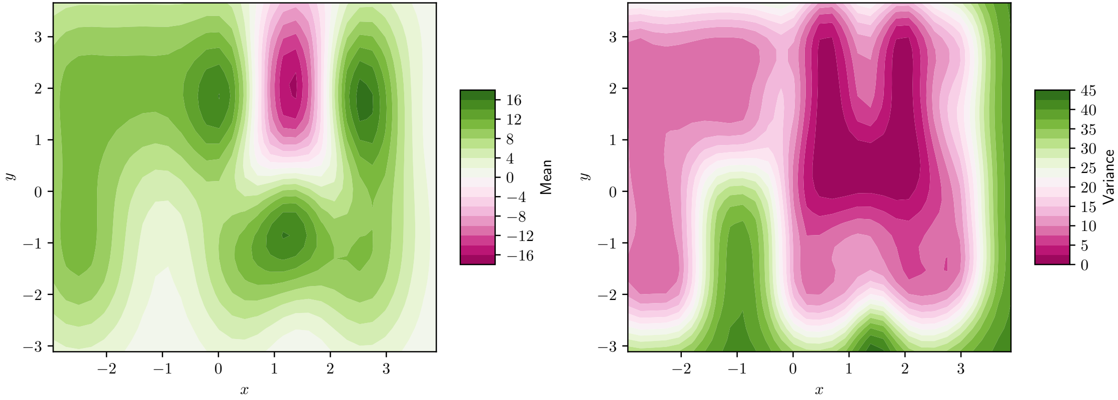

GP Posterior over Desired Mode's Gating Function

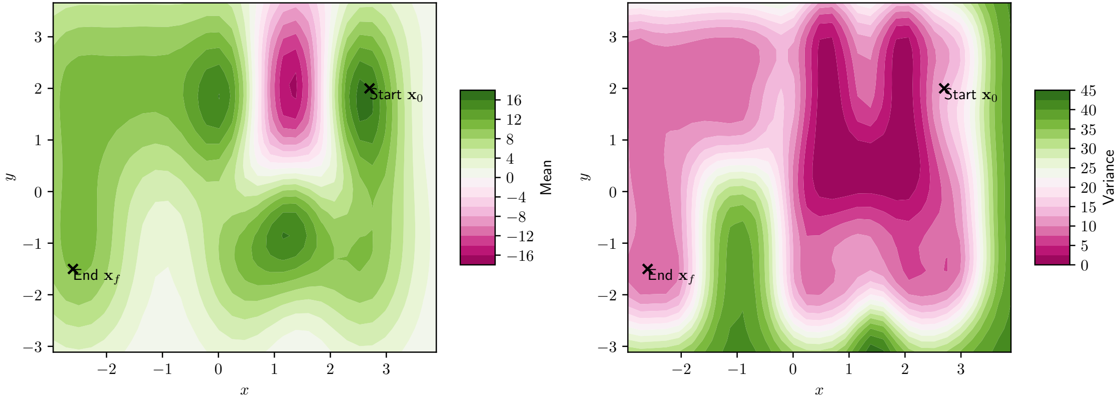

Encode Goals in Cost Function

$$\min_{\mathbf{x}(t),\mathbf{u}(t)} \int_{t_0}^{t_f} g_{\text{mode}}(\mathbf{x}(t)) + g_{\text{epistemic}}(\mathbf{x}(t))\text{d}t \quad \forall t \tag{1}$$

$$\text{s.t. transition dynamics}$$

$$\mathbf{x}(t_0) = \mathbf{x}_0 \quad \mathbf{x}(t_f) = \mathbf{x}_f$$

- g_{mode} - low in desired dynamics mode k^*
- $g_{\text{epistemic}}$ - high in regions of the dynamics with high epistemic uncertainty

Stage Two - Trajectory Optimisation

A geometric cost function?

- Desired mode's gating function $h^{(k^*)} : \mathcal{X} \to \mathbb{R}$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Desired mode's gating function $h^{(k^*)} : \mathcal{X} \rightarrow \mathbb{R}$
- State trajectory $\bar{\mathbf{x}} : [t_0, t_f] \rightarrow \mathcal{X}$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Desired mode's gating function $h^{(k^*)} : \mathcal{X} \rightarrow \mathbb{R}$
- State trajectory $\bar{\mathbf{x}} : [t_0, t_f] \rightarrow \mathcal{X}$
- Length of trajectory (in state space)

$$\text{Length}(\bar{\mathbf{x}}) = \int_{t_0}^{t_f} \|\dot{\mathbf{x}}(t)\|_2 \, \mathrm{d}t \tag{2}$$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Desired mode's gating function $h^{(k^*)} : \mathcal{X} \to \mathbb{R}$
- State trajectory $\bar{\mathbf{x}} : [t_0, t_f] \to \mathcal{X}$
- Length of trajectory (in state space)

$$\text{Length}(\bar{\mathbf{x}}) = \int_{t_0}^{t_f} \|\dot{\mathbf{x}}(t)\|_2 \, \mathrm{d}t \tag{2}$$

- Map through $h^{(k^*)}$

$$\text{Length}\left(h^{(k^*)}(\bar{\mathbf{x}})\right) = \int_{t_0}^{t_f} \left\|\dot{h}^{(k^*)}(\mathbf{x}(t))\right\|_2 \mathrm{d}t = \int_{t_0}^{t_f} \|\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)\|_2 \, \mathrm{d}t \tag{3}$$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Locally defined norm

$$\|\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)\|_2 = \sqrt{\dot{\mathbf{x}}(t)\mathbf{J}_{\mathbf{x}_t}^T\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} = \sqrt{\dot{\mathbf{x}}(t)\mathbf{G}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} \tag{4}$$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Locally defined norm

$$\|\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)\|_2 = \sqrt{\dot{\mathbf{x}}(t)\mathbf{J}_{\mathbf{x}_t}^T\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} = \sqrt{\dot{\mathbf{x}}(t)\mathbf{G}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} \tag{4}$$

- Encodes g_{mode}

$$J_{\text{geo}} = \min \text{Length}(h^{(k^*)}(\bar{\mathbf{x}})) = \min \int_{t_0}^{t_f} \|\dot{\mathbf{x}}(t)\|_{\mathbf{G}(\mathbf{x}(t))}\, \mathrm{d}t \tag{5}$$

Stage Two - Trajectory Optimisation

A geometric cost function?

- Locally defined norm

$$\|\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)\|_2 = \sqrt{\dot{\mathbf{x}}(t)\mathbf{J}_{\mathbf{x}_t}^T\mathbf{J}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} = \sqrt{\dot{\mathbf{x}}(t)\mathbf{G}_{\mathbf{x}_t}\dot{\mathbf{x}}(t)} \tag{4}$$

- Encodes g_{mode}

$$J_{\text{geo}} = \min \text{Length}(h^{(k^*)}(\bar{\mathbf{x}})) = \min \int_{t_0}^{t_f} \|\dot{\mathbf{x}}(t)\|_{\mathbf{G}(\mathbf{x}(t))} \, \mathrm{d}t \tag{5}$$

- How do we encode $g_{\text{epistemic}}$?

Stage Two - Trajectory Optimisation

Metrics for Probabilistic Geometries

[a]

$$p\left(\mathbf{J}_*|\mathbf{x}_*, \xi^{(k)}\right) = \int q\left(\hat{\mathbf{h}}^{(k)}\right) p\left(\mathbf{J}_* \mid \mathbf{x}_*, \hat{\mathbf{h}}^{(k)}, \xi^{(k)}\right) \mathrm{d}\hat{\mathbf{h}}^{(k)} = \mathcal{N}\left(\mathbf{J}_* \mid \boldsymbol{\mu}_J, \boldsymbol{\Sigma}_J\right) \tag{6}$$

$$\mathbf{G} = \mathcal{W}_D\left(p, \boldsymbol{\Sigma}_J, \mathbb{E}\left[\mathbf{J}^T\right] \mathbb{E}[\mathbf{J}]\right) \tag{7}$$

$$\mathbb{E}[\mathbf{G}] = \mathbb{E}[\mathbf{J}^T] \mathbb{E}[\mathbf{J}] + \lambda \boldsymbol{\Sigma}_J \tag{8}$$

[a] Alessandra Tosi et al. “Metrics for Probabilistic Geometries”. In: *Proceedings of the 30th Conference*. Uncertainty in Artificial Intelligence. 2014, pp. 800–808

Implicit Minimisation

Geodesic ODE

[a]

$$\ddot{\mathbf{x}}(t) = f_G(t, \dot{\mathbf{x}}, \mathbf{x})$$
$$= -\frac{1}{2}\mathbf{G}^{-1}(\mathbf{x}(t)) \left[\frac{\partial \operatorname{vec}[\mathbf{G}(\mathbf{x}(t))]}{\partial \mathbf{x}(t)}\right]^T (\dot{\mathbf{x}}(t) \otimes \dot{\mathbf{x}}(t))$$

[a] Manfredo do Carmo. *Riemannian Geometry*. Mathematics: Theory & Applications. Birkhäuser Basel, 1992

- Solve geodesic ODE subject to BCs

$$\mathbf{x}(t_0) = \mathbf{x}_0 \quad \mathbf{x}(t_f) = \mathbf{x}_f$$

Hermite-Simpson Collocation

- Approximate ODE with piecewise quadratic functions

Hermite-Simpson Collocation

- Approximate ODE with piecewise quadratic functions
- Collocation points $\{\mathbf{z}_i \in \mathcal{X}\}_{i=1}^{I}$

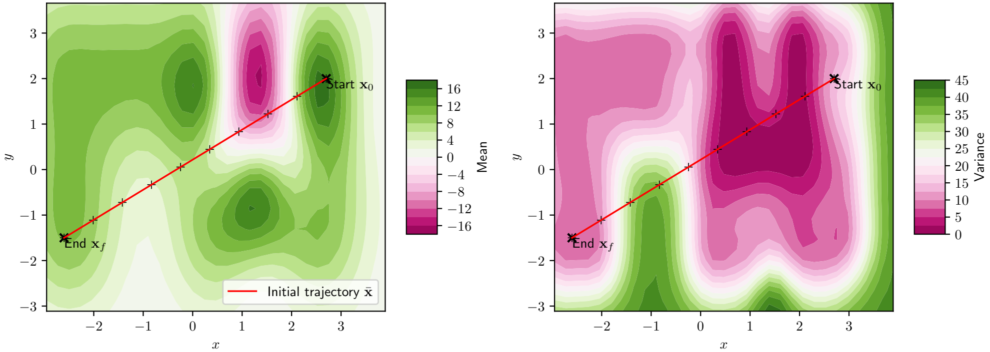

Hermite-Simpson Collocation

- Approximate ODE with piecewise quadratic functions
- Collocation defects at mid points $\Delta_i = \underbrace{\ddot{\mathbf{z}}_{i+\frac{1}{2}}}_{\text{Interpolated}} - \underbrace{f_G(t_{i+\frac{1}{2}}, \dot{\mathbf{z}}_{i+\frac{1}{2}}, \mathbf{z}_{i+\frac{1}{2}})}_{\text{ODE}}$

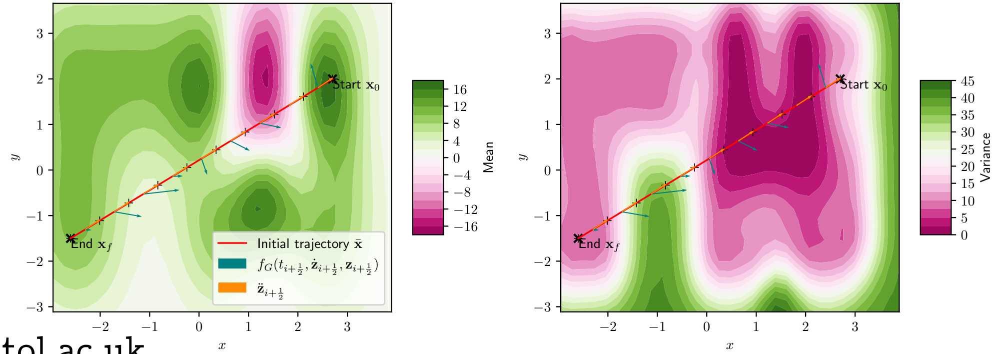

Hermite-Simpson Collocation

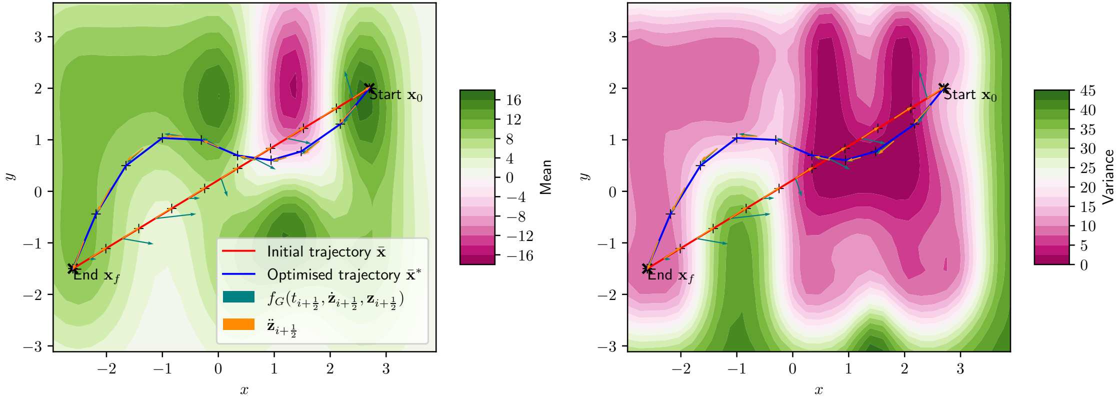

Results

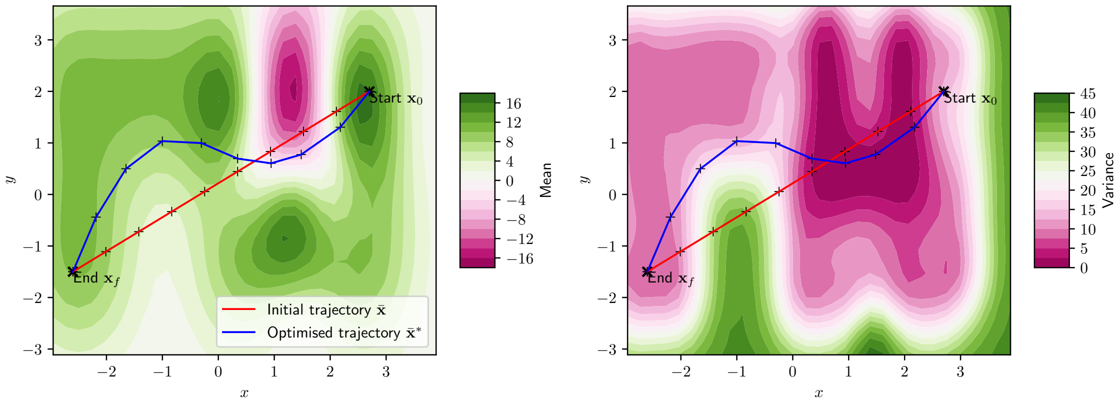

Start $\mathbf{x}_0$
End $\mathbf{x}_f$
Initial trajectory $\bar{\mathbf{x}}$
Optimised trajectory $\bar{\mathbf{x}}^*$
Mean
Variance
x
y

The Niggly Bits

- Data collection requires entering undesired dynamics modes. . .

The Niggly Bits

- Data collection requires entering undesired dynamics modes...
  - Geometry inspired active learning?

The Niggly Bits

- Data collection requires entering undesired dynamics modes. . .
  - Geometry inspired active learning?
- Is it OK to project stochastic dynamics model onto a deterministic ODE?

The Niggly Bits

- Data collection requires entering undesired dynamics modes. . .
  - Geometry inspired active learning?
- Is it OK to project stochastic dynamics model onto a deterministic ODE?
- How to set tolerance?

Thanks for Listening

Questions?

Thanks for Listening

- Manfredo do Carmo. *Riemannian Geometry*. Mathematics: Theory & Applications. Birkhäuser Basel, 1992.
- Alessandra Tosi, Søren Hauberg, Alfredo Vellido, and Neil D Lawrence. "Metrics for Probabilistic Geometries". In: *Proceedings of the 30th Conference*. Uncertainty in Artificial Intelligence. 2014, pp. 800–808.