Data Association with Gaussian Processes

Markus Kaiser, Clemens Otte, Thomas A. Runkler, Carl Henrik Ek
markus.kaiser@siemens.com

September 17, 2019

Siemens AG, Technical University of Munich, University of Bristol

Noise Separation

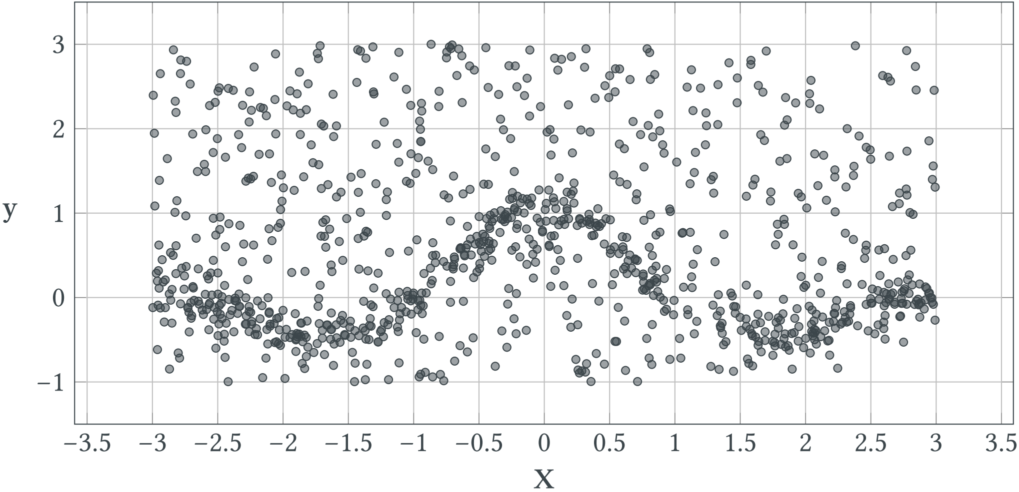

Noise Separation

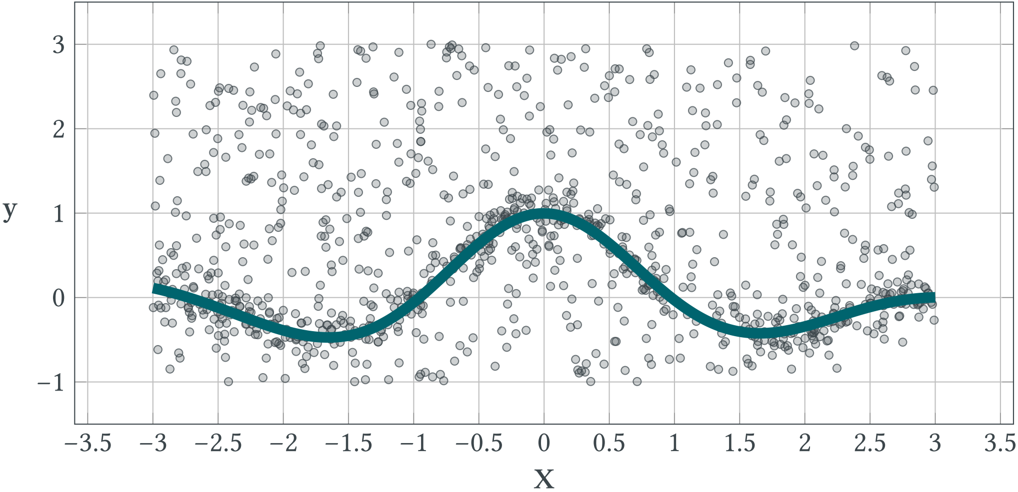

Noise Separation

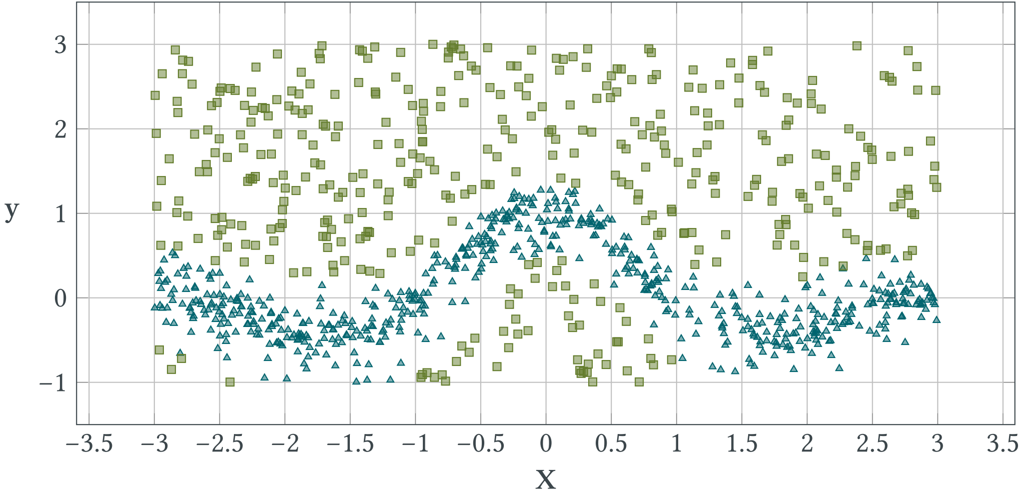

Noise Separation

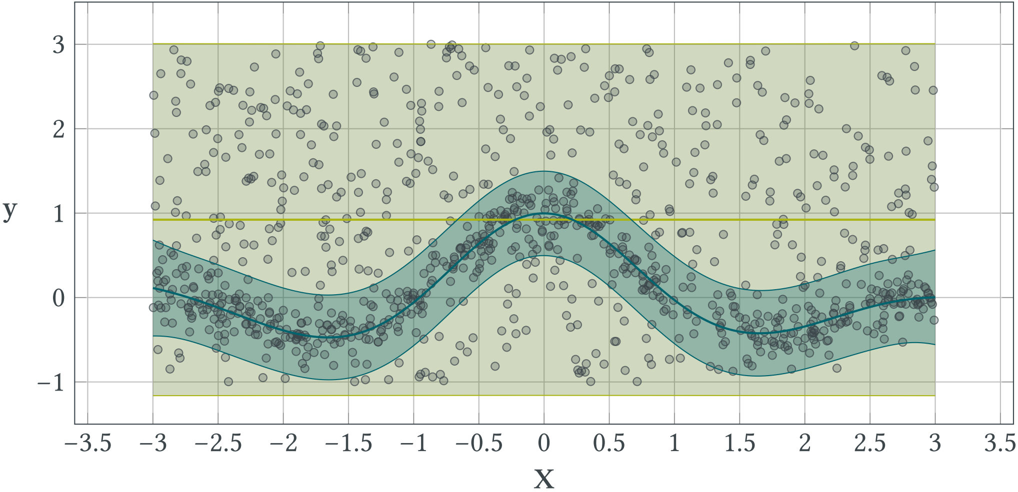

Multimodal data

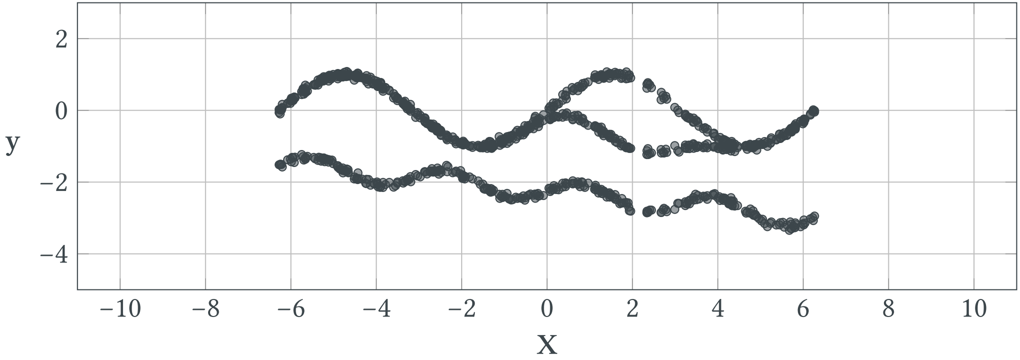

Multimodal data

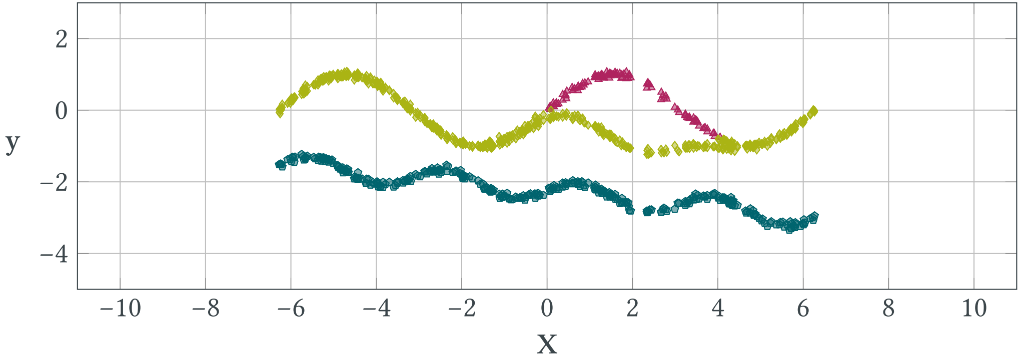

Multimodal data

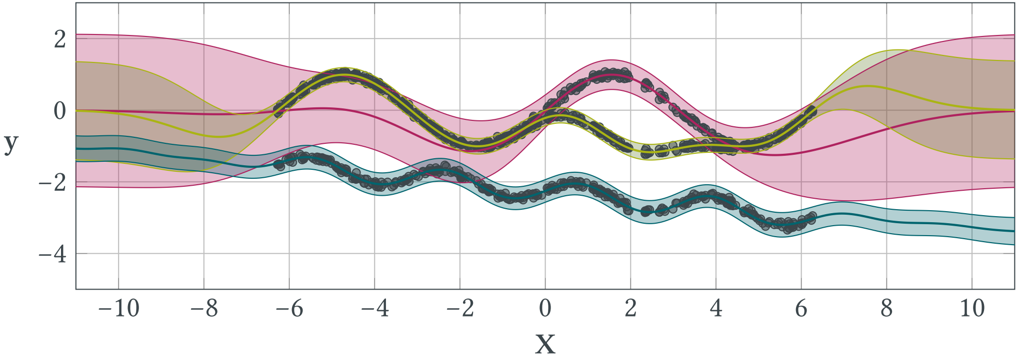

Multimodal data

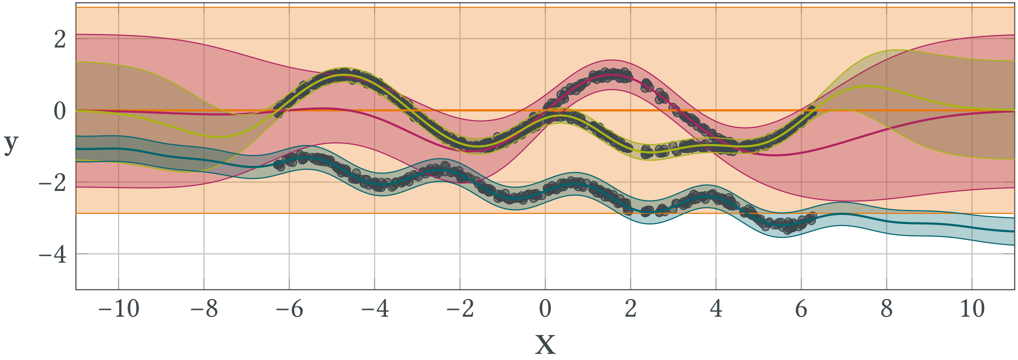

Multimodal data

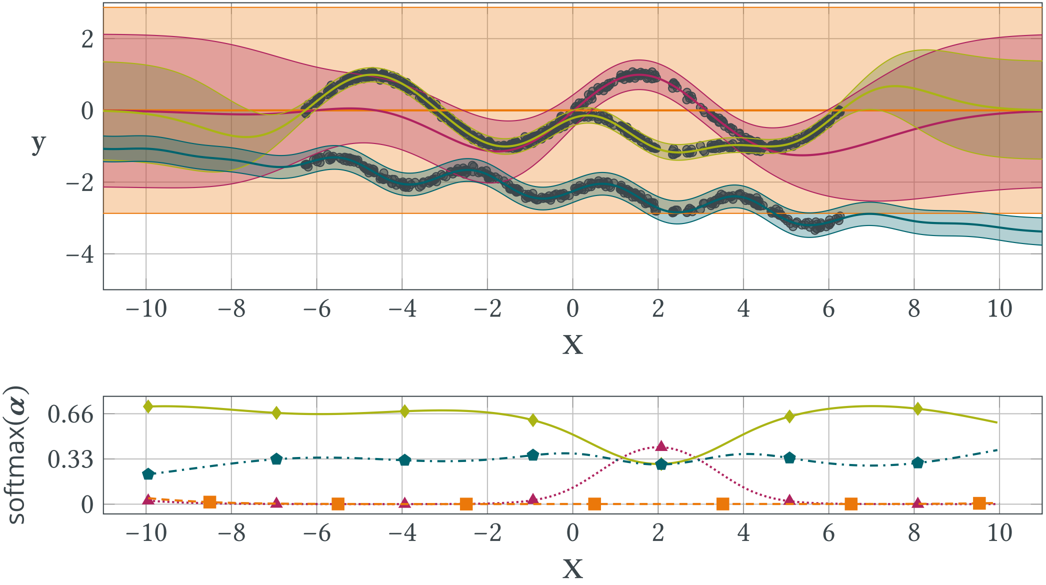

Graphical Model of DAGP

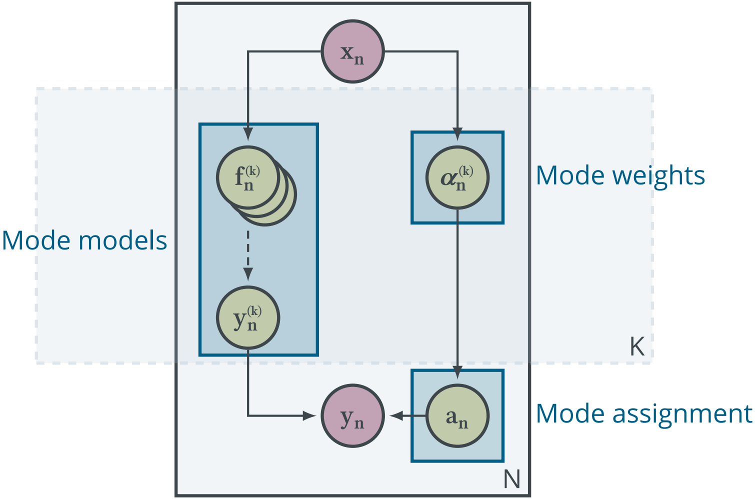

Graphical Model of DAGP

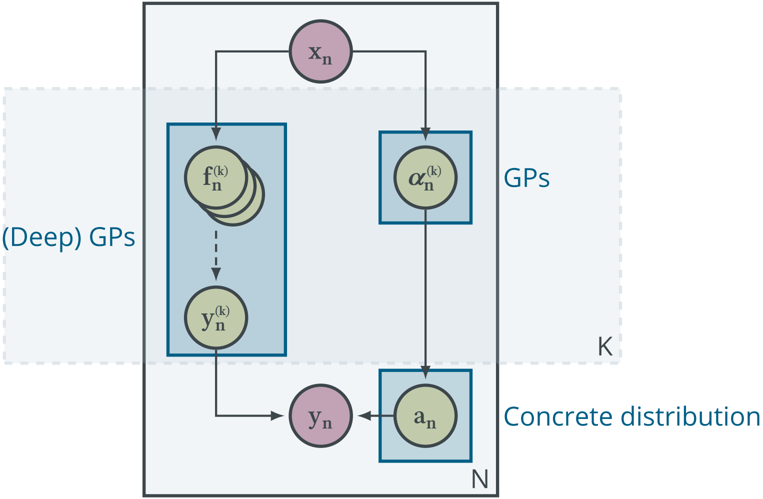

Graphical Model of DAGP

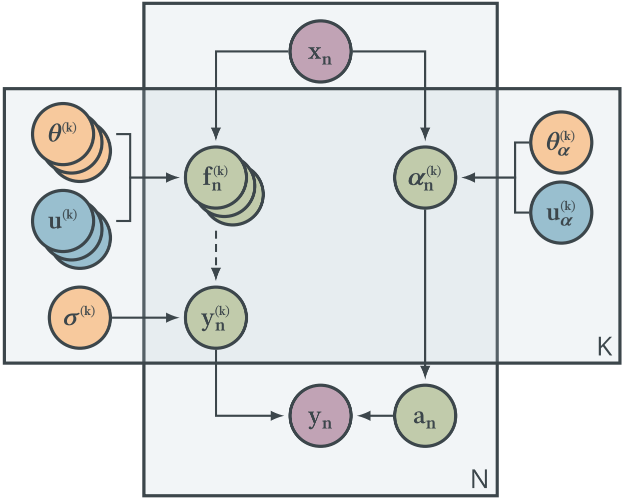

Wet-Chicken Benchmark[1]

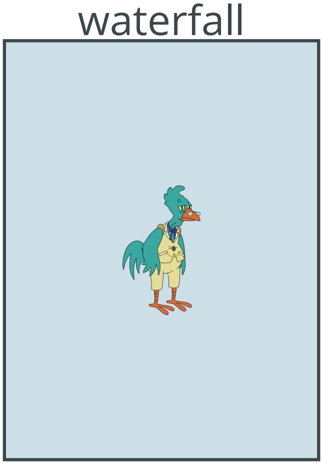

Dynamics Agent in a flowing river

Goal Get close to the waterfall

State (x, y)-position in $\mathbb{R}^2$

Action (x, y)-movement in $\mathbb{R}^2$

[1] Tresp 1994; Hans and Udluft 2009.

Wet-Chicken Benchmark[1]

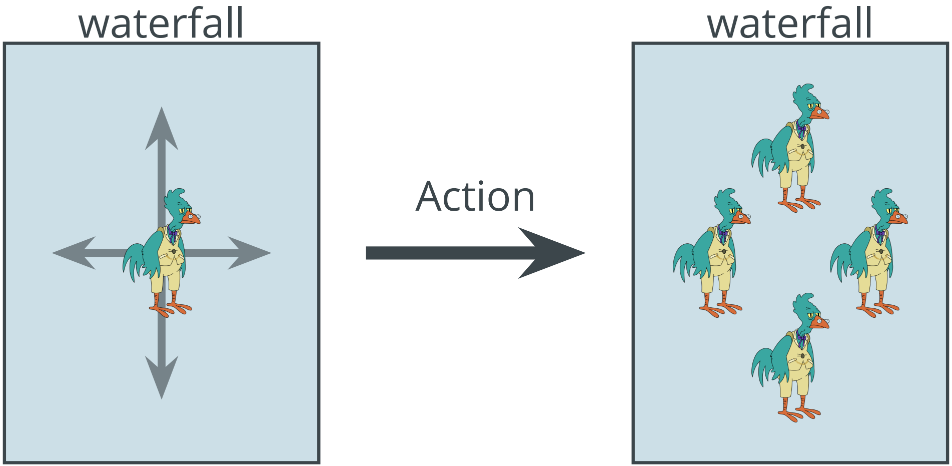

[1] Tresp 1994; Hans and Udluft 2009.

Wet-Chicken Benchmark[1]

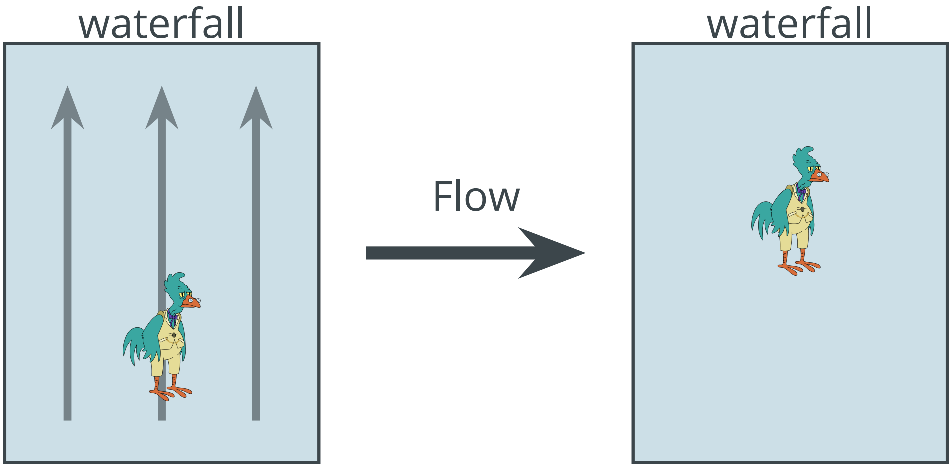

[1] Tresp 1994; Hans and Udluft 2009.

Wet-Chicken Benchmark[1]

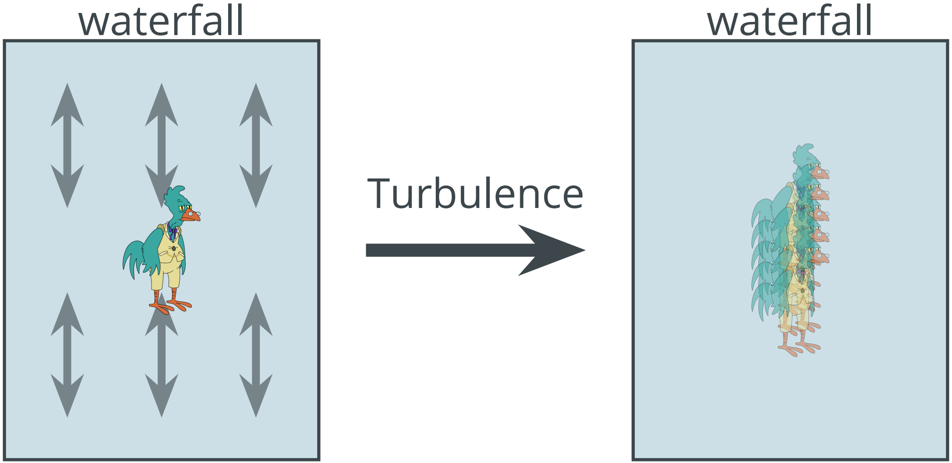

[1]Tresp 1994; Hans and Udluft 2009.

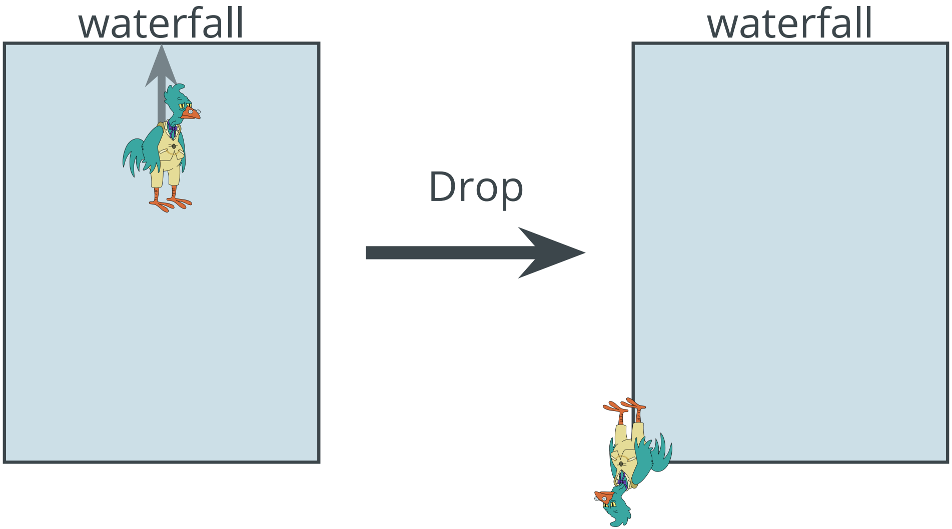

[1]Tresp 1994; Hans and Udluft 2009.

Multimodal System Dynamics

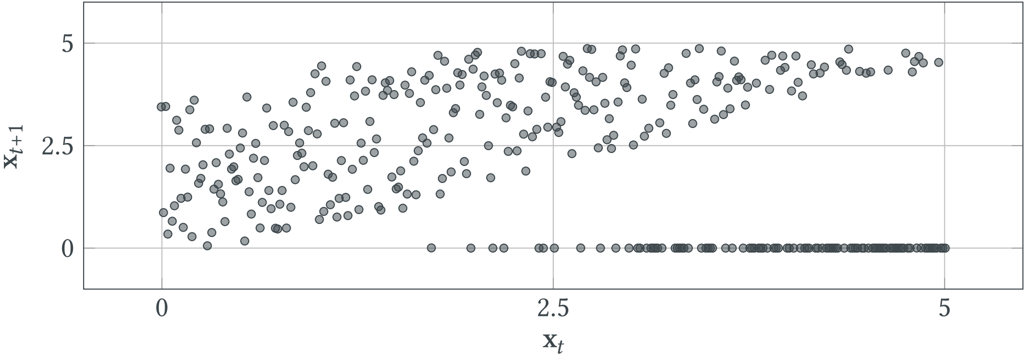

Multimodal System Dynamics

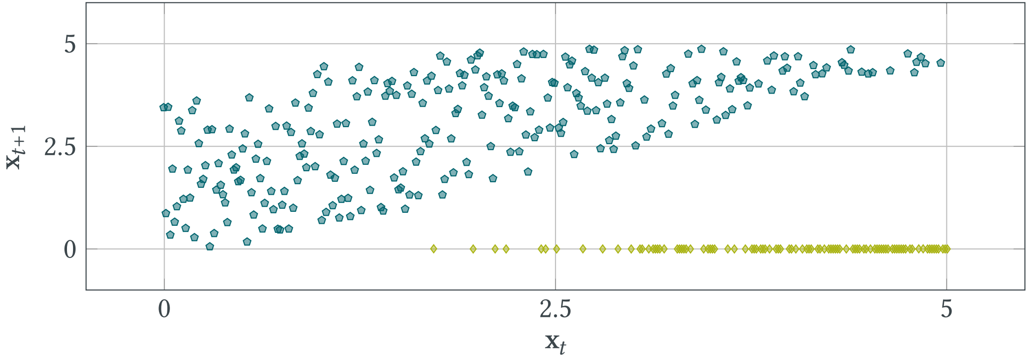

Multimodal System Dynamics

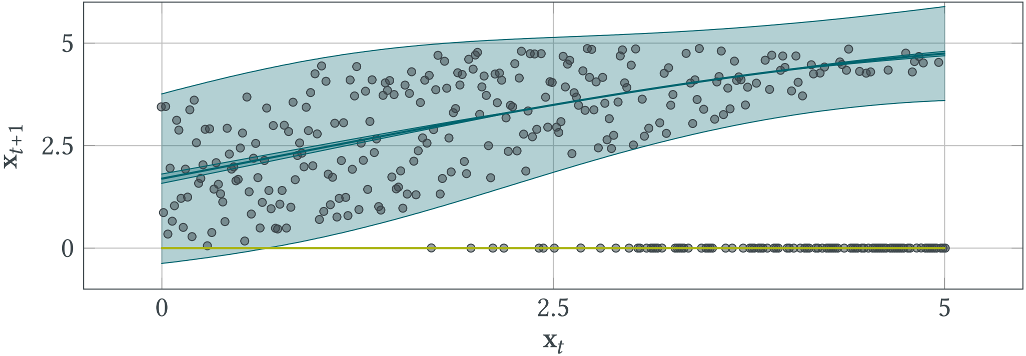

Multimodal System Dynamics

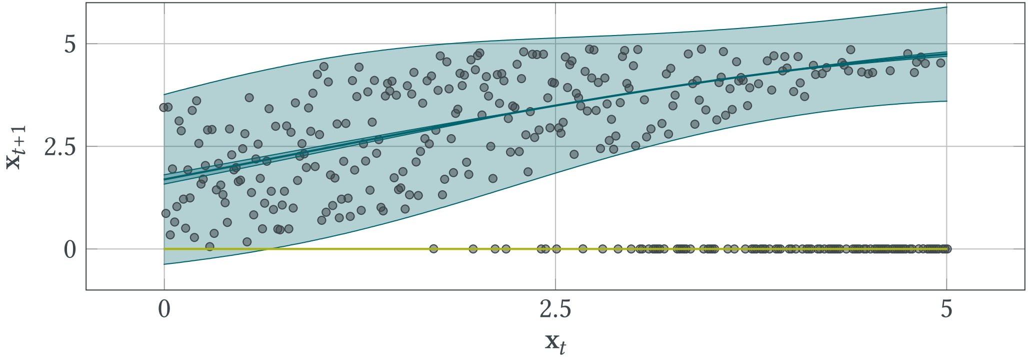

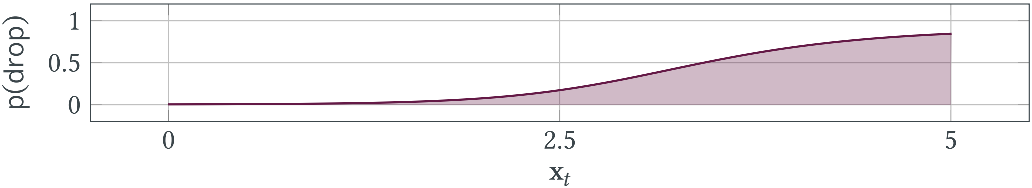

Conservative Policy Training

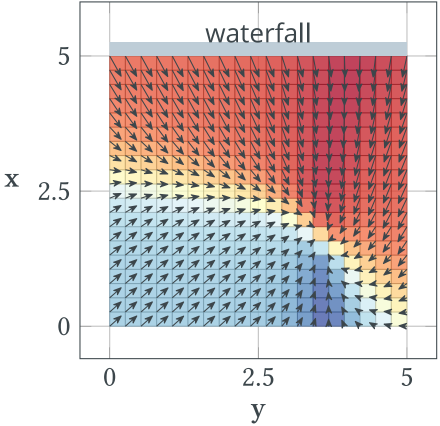

$$R(x, y) = x$$

22% drop rate

Conservative Policy Training

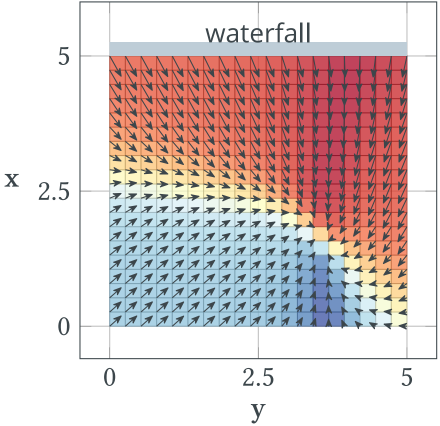

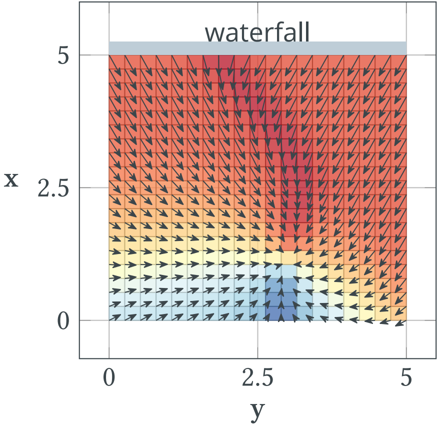

$R(x, y) = x$

22% drop rate

$R'(x, y) = R(x, y) - 5 \cdot \mathrm{p}(\mathrm{drop}\,|\,x, y)$

19% drop rate

Data Association with Gaussian Processes

Model for multimodal data

- Separate models per mode
- Predictive Associations
- Scalable inference

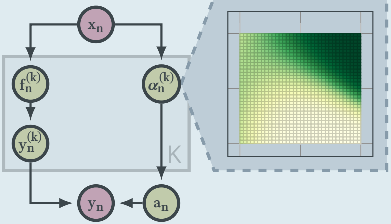

Informed decision making

- Hierarchical priors
- Interpretable sub-models
- Stochastic systems

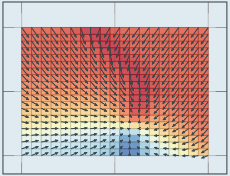

References

- Hans, Alexander and Steffen Udluft (2009). “Efficient Uncertainty Propagation for Reinforcement Learning with Limited Data”. In: *International Conference on Artificial Neural Networks*. Springer, pp. 70–79.
- Kaiser, Markus et al. (2019). “Interpretable Dynamics Models for Data-Efficient Reinforcement Learning”. In: *Computational Intelligence and Machine Learning* ESANN 2019 proceedings, p. 6.
- Tresp, Volker (1994). “The Wet Game of Chicken”. In: *Siemens AG, CT IC 4, Technical Report*.

Variational Bound

$$
\begin{aligned}
\mathscr{L}_{\mathrm{DAGP}} &= \mathbb{E}_{\mathrm{q}(\mathbf{F},\boldsymbol{\alpha},\mathbf{U})}\left[\log \frac{\mathrm{p}(\mathbf{Y},\mathbf{A},\mathbf{F},\boldsymbol{\alpha},\mathbf{U}\,|\,\mathbf{X})}{\mathrm{q}(\mathbf{F},\boldsymbol{\alpha},\mathbf{U})}\right] \\
&= \sum_{n=1}^{N} \mathbb{E}_{\mathrm{q}(\mathbf{f_n})}[\log \mathrm{p}(\mathbf{y_n}\,|\,\mathbf{f_n},\mathbf{a_n})] + \sum_{n=1}^{N} \mathbb{E}_{\mathrm{q}(\boldsymbol{\alpha}_\mathbf{n})}[\log \mathrm{p}(\mathbf{a_n}\,|\,\boldsymbol{\alpha}_\mathbf{n})] \\
&\quad - \sum_{k=1}^{K} \mathrm{KL}\left(\mathrm{q}\left(\mathbf{u}^{(\mathrm{k})}\right) \,\middle\|\, \mathrm{p}\left(\mathbf{u}^{(\mathrm{k})}\,\middle|\,\mathbf{Z}^{(\mathrm{k})}\right)\right) - \sum_{k=1}^{K} \mathrm{KL}\left(\mathrm{q}\left(\mathbf{u}_{\boldsymbol{\alpha}}^{(\mathrm{k})}\right) \,\middle\|\, \mathrm{p}\left(\mathbf{u}_{\boldsymbol{\alpha}}^{(\mathrm{k})}\,\middle|\,\mathbf{Z}_{\boldsymbol{\alpha}}^{(\mathrm{k})}\right)\right)
\end{aligned}
$$

Predictive Posterior

$$
\begin{aligned}
\mathrm{q}(\mathbf{f}_* \mid \mathbf{x}_*) &= \int \sum_{k=1}^{K} \mathrm{q}\left(a_*^{(k)} \middle| \mathbf{x}_*\right) \mathrm{q}\left(\mathbf{f}_*^{(\mathbf{k})} \middle| \mathbf{x}_*\right) \mathrm{d}\mathbf{a}_*^{(\mathbf{k})} \\
&\approx \sum_{k=1}^{K} \hat{a}_*^{(k)} \hat{\mathbf{f}}_*^{(\mathbf{k})}
\end{aligned}
$$