Uncertainties Need a Purpose

Markus Kaiser
mrksr.de

10 October 2019

Siemens AG, Technical University of Munich

Marginal Uncertainties and Samples

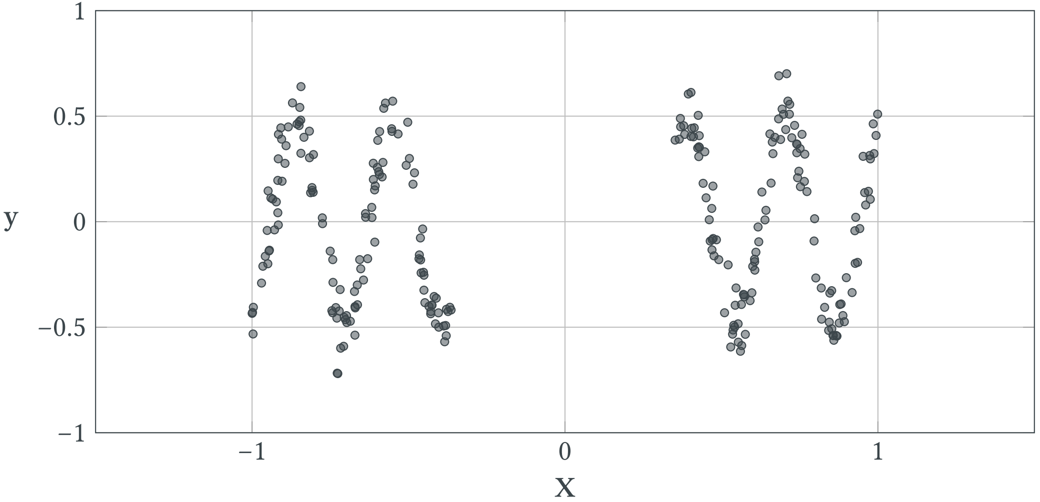

Marginal Uncertainties and Samples

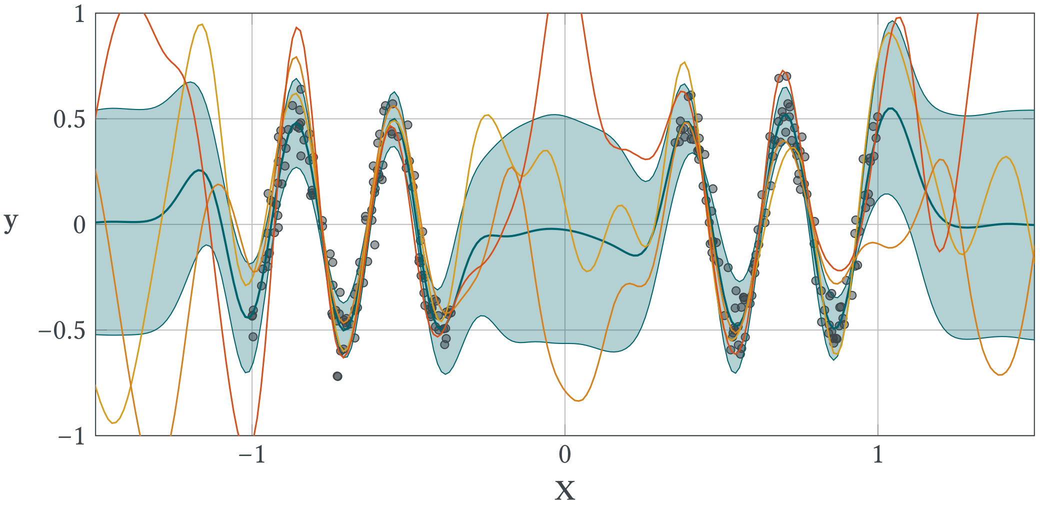

Marginal Uncertainties and Samples

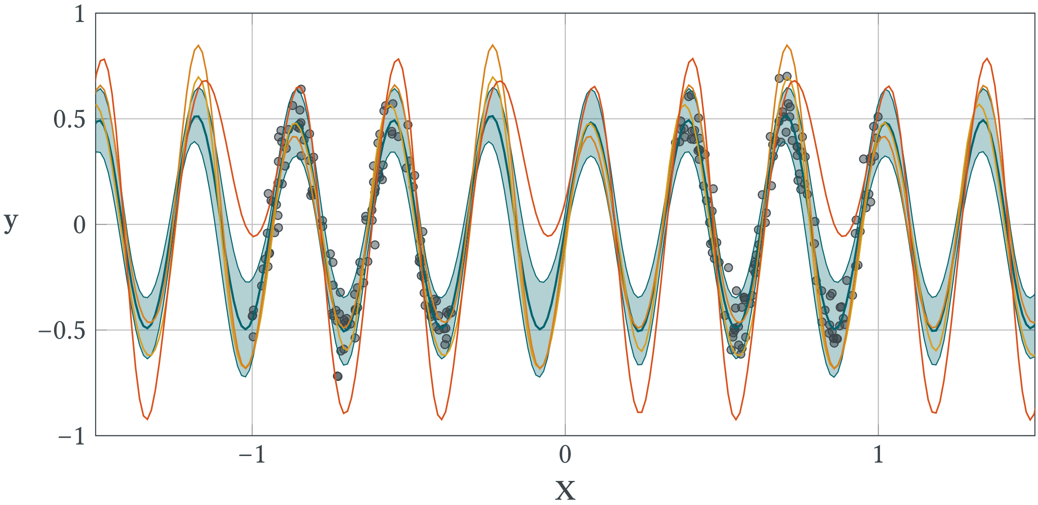

Marginal Uncertainties and Samples

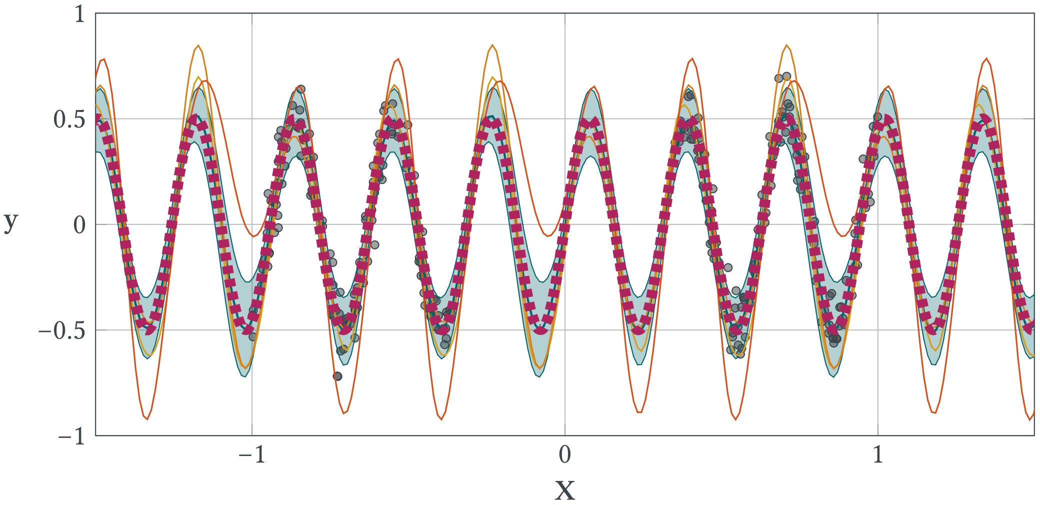

Marginal Uncertainties and Samples

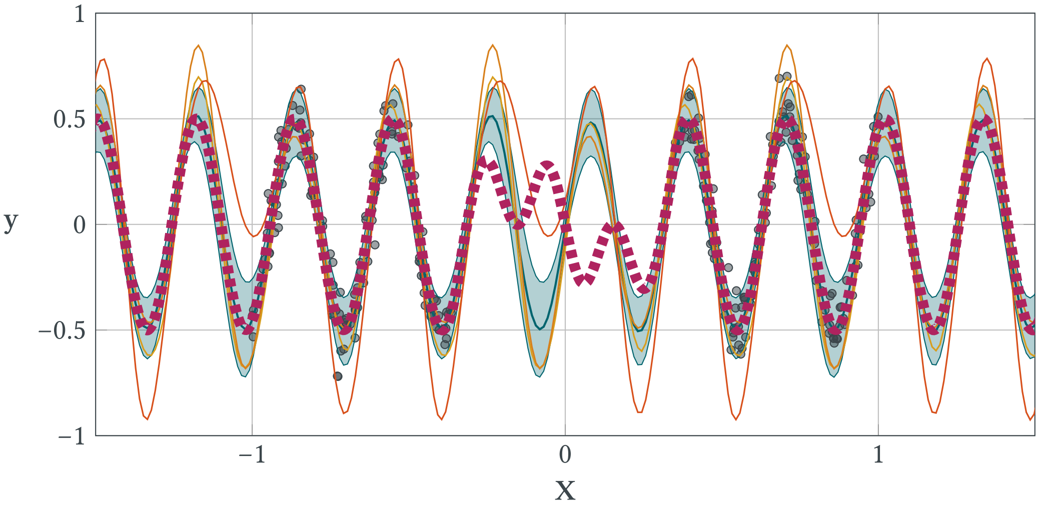

Marginal Uncertainties and Samples

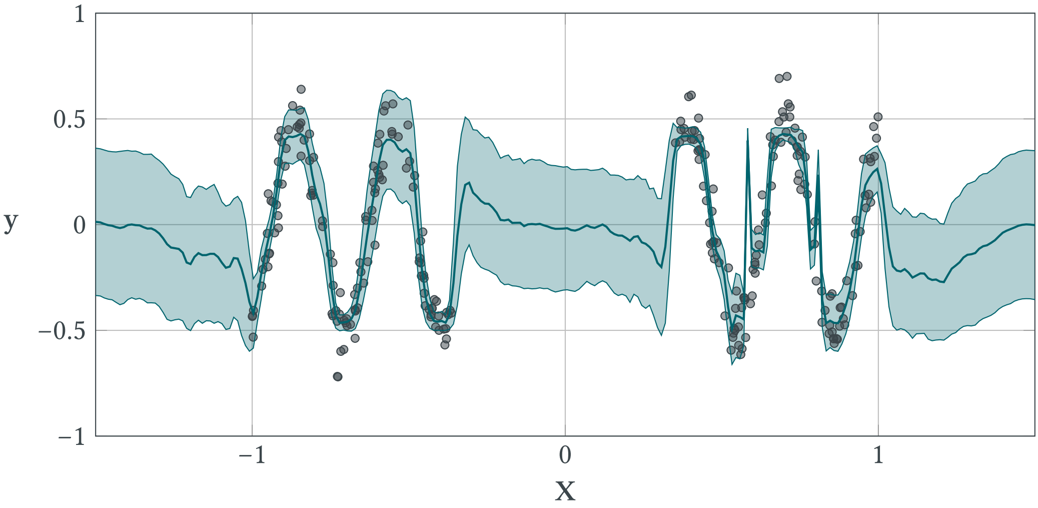

Marginal Uncertainties and Samples

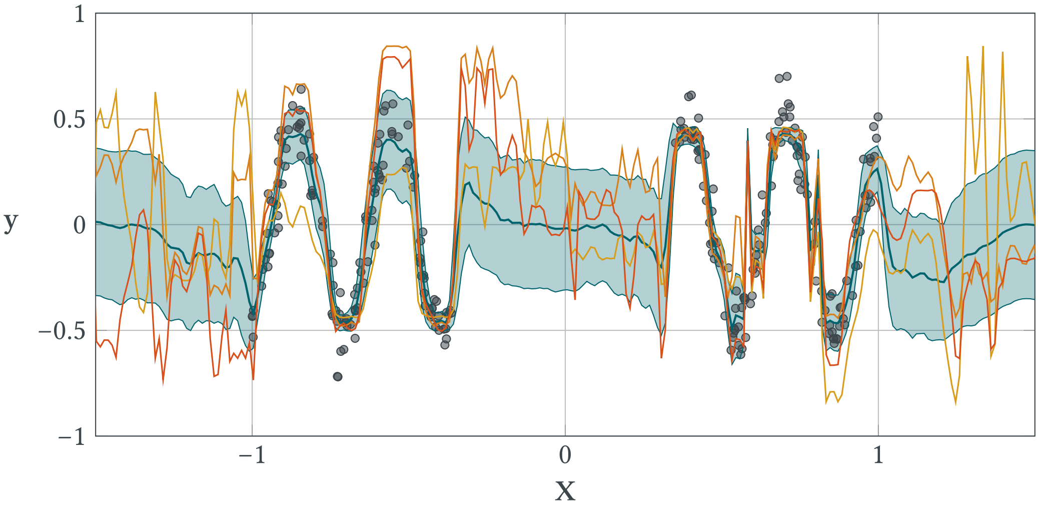

Case study: Wind propagation[1]

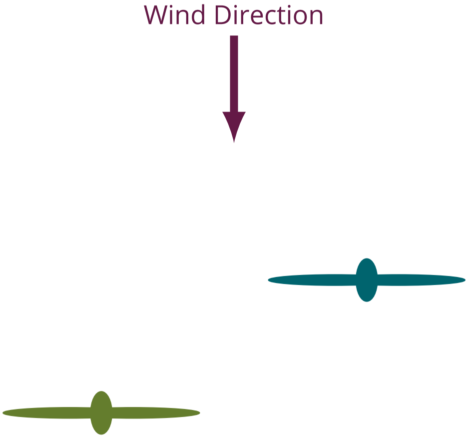

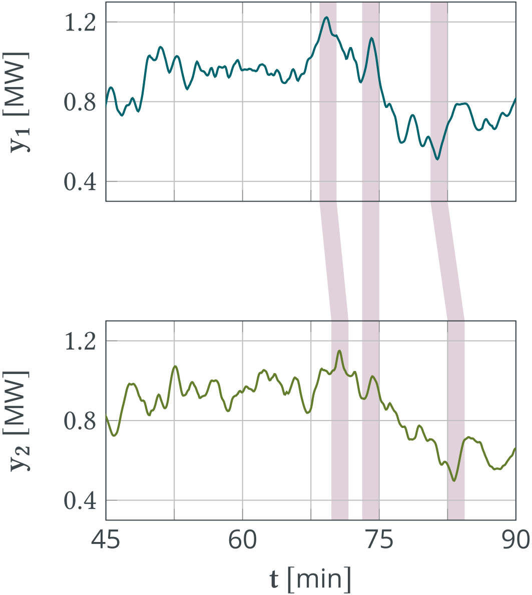

[1] Kaiser et al. 2018.

Case study: Wind propagation

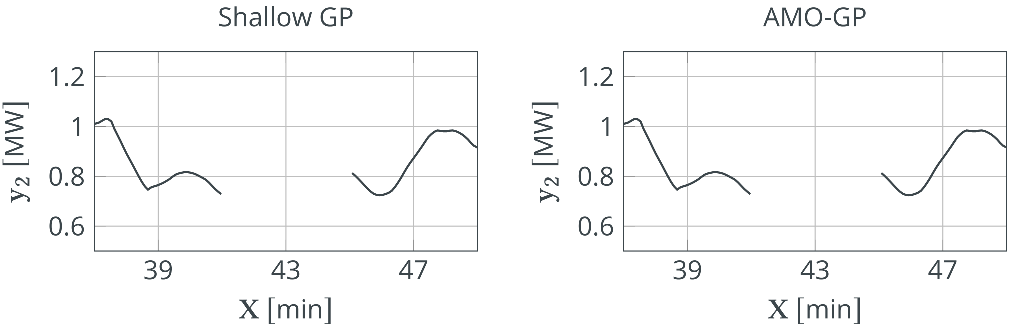

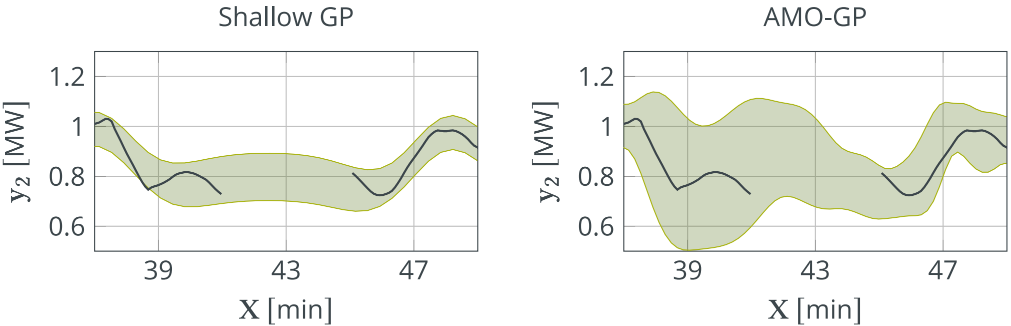
Shallow GP
AMO-GP
1.2
1
0.8
0.6
y_2 [MW]
39
43
47
X [min]

Case study: Wind propagation

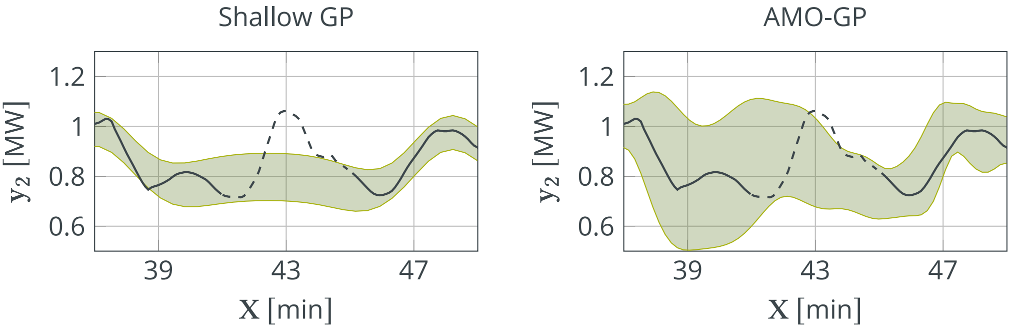

Case study: Wind propagation

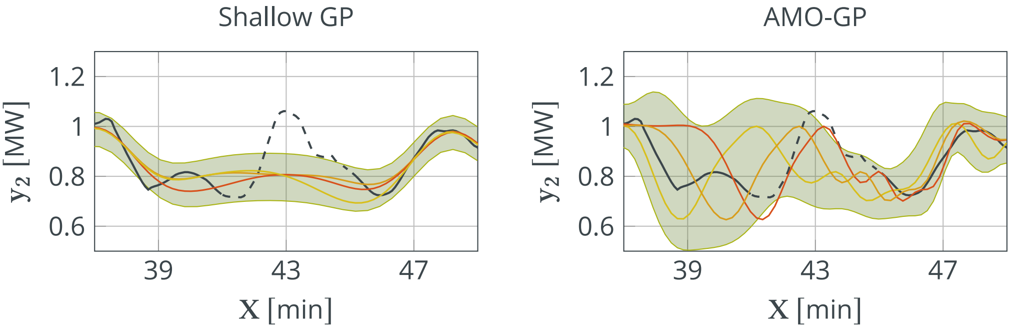

Reasoning about (un-)supervised learning is hard!

- **What makes a model good?**
- The marginal likelihood is not enough
- We can find good models **with respect to some prior**

Reasoning about (un-)supervised learning is hard!

- **What makes a model good?**
- The marginal likelihood is not enough
- We can find good models **with respect to some prior**

- Why do we use badly understood priors then?
- Why do we ask for good performance relative to wrong priors?

Reasoning about (un-)supervised learning is hard!

- **What makes a model good?**
- The marginal likelihood is not enough
- We can find good models **with respect to some prior**

- Why do we use badly understood priors then?
- Why do we ask for good performance relative to wrong priors?

Do we use priors as proxies for tasks?

Case Study: Bayesian Optimization[2]

Bayesian Optimization

- **Task: Find the minimum** of some function f given few observations

$$\mathbf{x}_* \in \operatorname*{argmin}_{\mathbf{x}} f(\mathbf{x})$$

- Assume Bayesian **prior** $f \sim GP(\cdot, \cdot)$ and derive a **posterior**
- Use some **acquisition function** to translate to belief about minimum

[2]Bodin et al. 2019.

Case Study: Bayesian Optimization[2]

Bayesian Optimization

- **Task: Find the minimum** of some function f given few observations

$$\mathbf{x}_* \in \operatorname*{argmin}_{\mathbf{x}} f(\mathbf{x})$$

- Assume Bayesian **prior** $f \sim GP(\cdot, \cdot)$ and derive a **posterior**
- Use some **acquisition function** to translate to belief about minimum

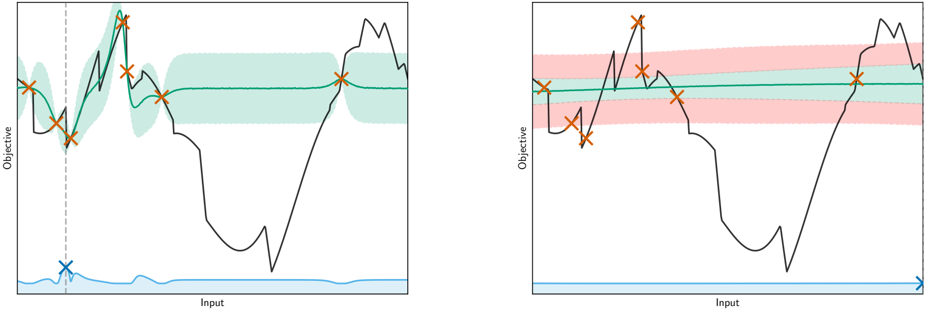

[2]Bodin et al. 2019.

Case Study: Bayesian Probabilistic Numerics[3]

Bayesian Quadrature

- **Task: Approximate a definite integral** over f

$$Q(f) = \int_a^b f(t)\,\mathrm{d}t$$

- Assume Bayesian **prior** for f, gather J **observations** of f
- Derive a **posterior** belief about $Q(f)$

[3] Oates and Sullivan 2019.

Case Study: Bayesian Probabilistic Numerics[3]

Bayesian Quadrature

- **Task: Approximate a definite integral** over f

$$Q(f) = \int_a^b f(t)\,\mathrm{d}t$$

- Assume Bayesian **prior** for f, gather J **observations** of f
- Derive a **posterior** belief about $Q(f)$

Assume $f \sim GP(0, \min(\cdot,\cdot))$. Then the posterior marginal for $Q(f)$ is

$$\mathcal{N}\left(\underbrace{\frac{1}{2}\sum_{j=1}^{J-1}(z_{j+1}+z_j)(t_{j+1}-t_j)}_{\textbf{Trapezoidal rule}}, \frac{1}{12}\sum_{j=1}^{J-1}(t_{j+1}-t_j)^3\right)$$

[3]Oates and Sullivan 2019.

Quantities of Interest

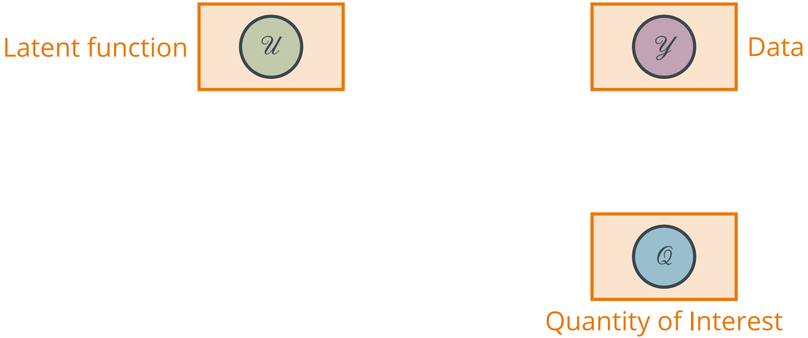

Quantities of Interest

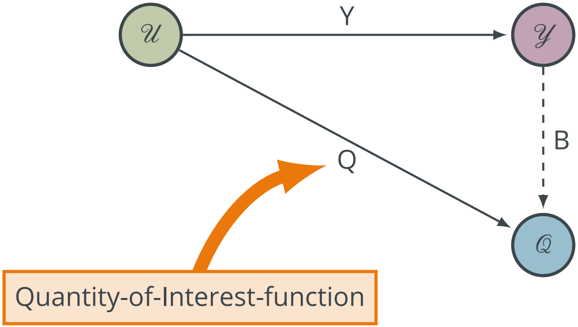

Quantities of Interest

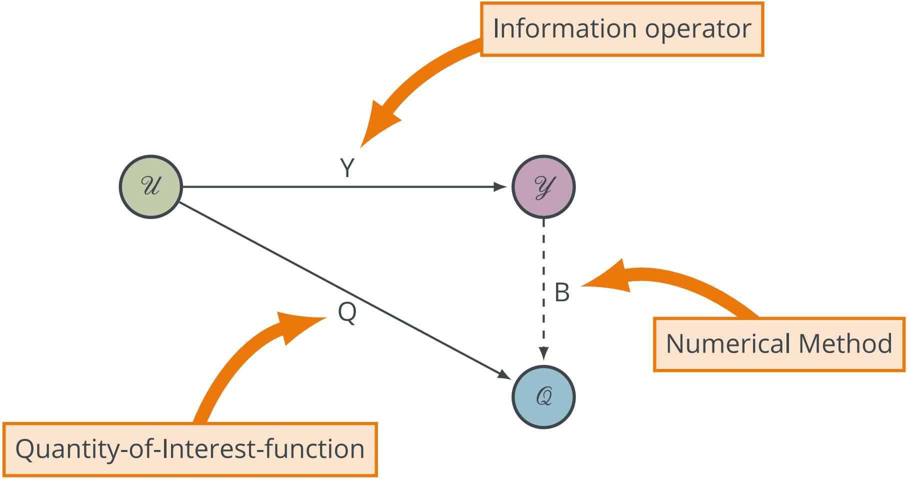

Quantities of Interest

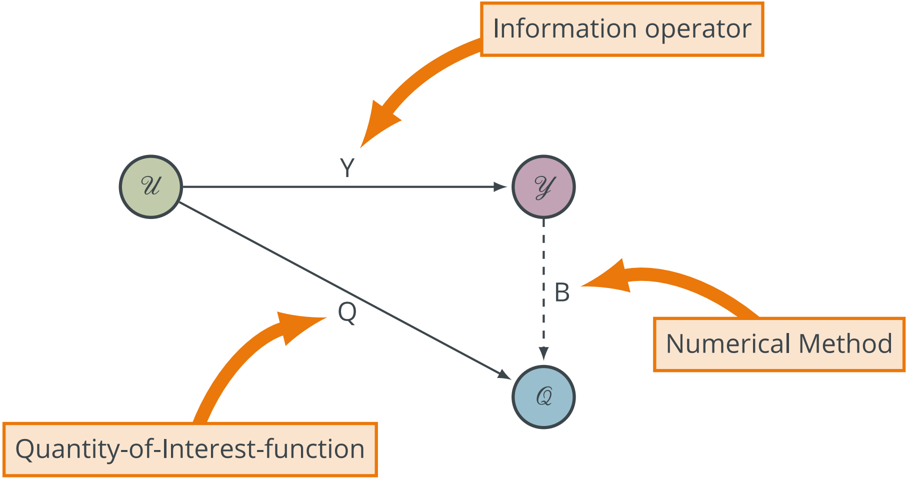

Bayesian belief about $\mathcal{U}$ and $\mathcal{Y}$ can be translated in belief about $\mathcal{Q}$.

Reinforcement Learning

- **Task: Find a policy π_* with maximum value** wrt. a system f and reward r

$$\pi_* \in \operatorname*{argmax}_{\pi} \mathbb{E}[J^{\pi}(\mathbf{s}_0)] = \sum_{t=1}^{T} \gamma^t \, \mathbb{E}_{p(\mathbf{s}_t)}[r(\mathbf{s}_t)]$$

- Assume Bayesian **prior** $f \sim GP(\cdot, \cdot)$ and derive a **posterior**
- Use **(stochastic) roll outs** in the optimization problem

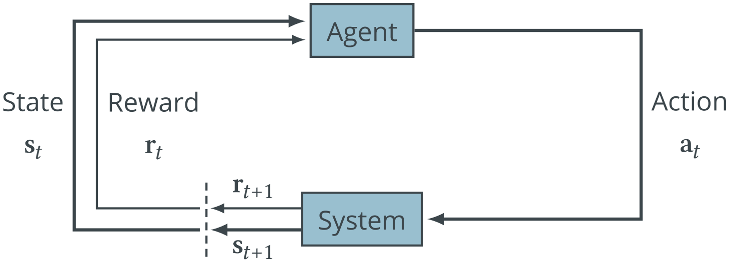

Graphical Model for RL

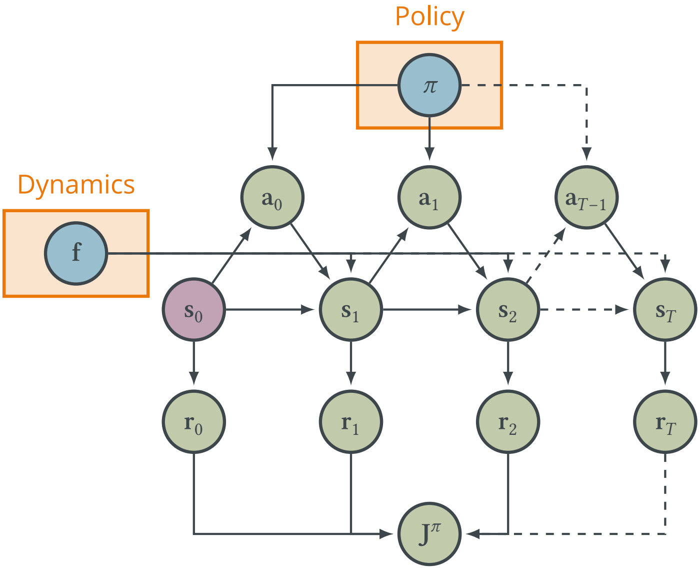

Graphical Model for RL

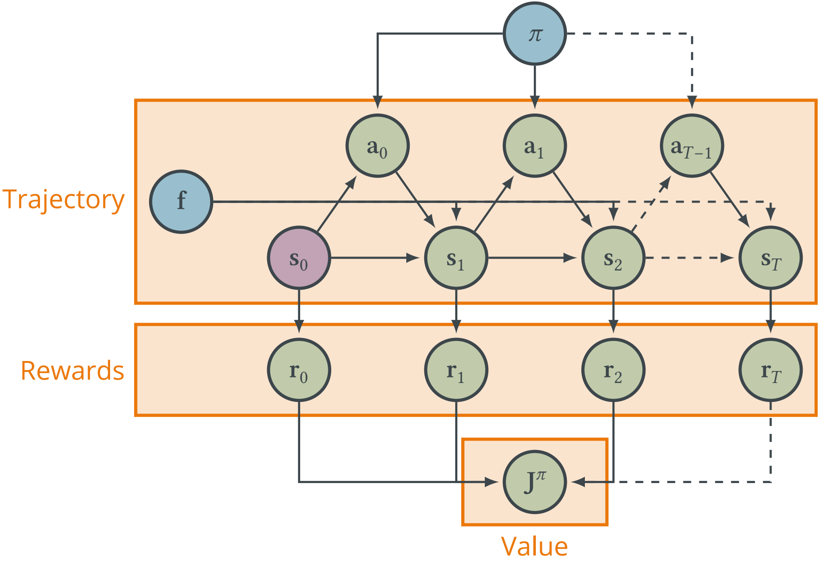

Graphical Model for RL

$$\pi_* \in \operatorname*{argmax}_{\pi} \mathbb{E}[\mathbf{J}^{\pi}]$$

$$\mathbb{E}[\mathbf{J}^{\pi}] = \sum_{t=0}^{T} \gamma^t \, \mathbb{E}_{p(\mathbf{s}_t|\pi)}[\mathbf{r}_t]$$
$$\approx \frac{1}{P} \sum_{p=1}^{P} \sum_{t=0}^{T} \gamma^t r_t^{(p)}$$

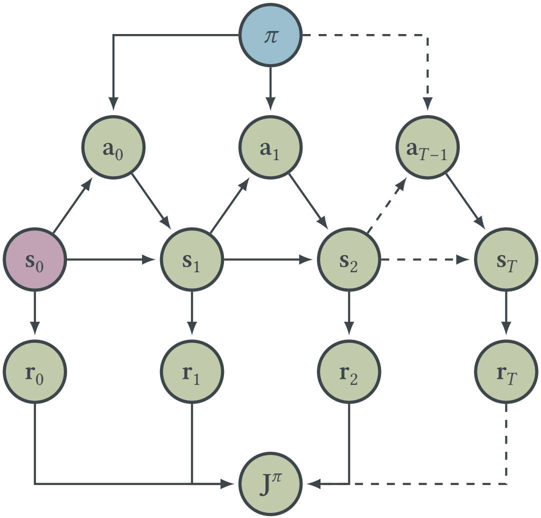

Graphical Model for RL

$$p(J_*) = ?$$

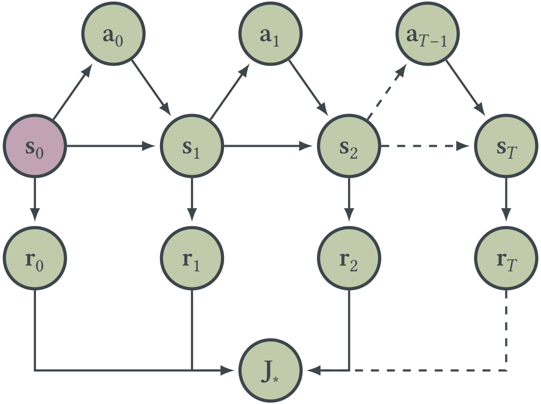

Graphical Model for RL

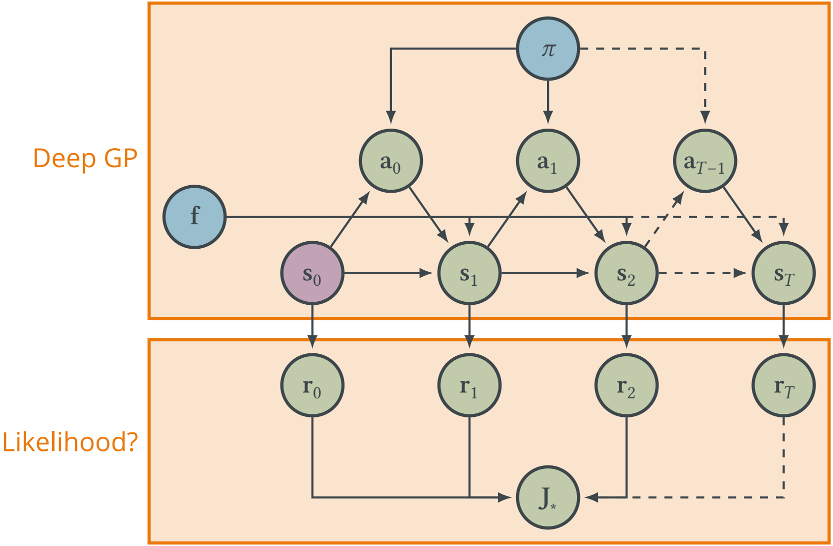

Quantities of Interest in Reinforcement Learning

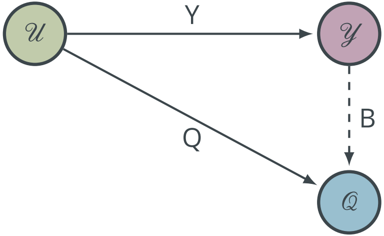

	Probabilistic Numerics	Reinforcement Learning
$\mathscr{U}$	Latent function	True system-dynamics
$\mathscr{Q}$	Definite Integral	Optimal value
$\mathscr{Y}$	Function evaluations	Batch/Online data
$Q : \mathscr{U} \longrightarrow \mathscr{Q}$	Integration	Bellman principle
$Y : \mathscr{U} \longrightarrow \mathscr{Y}$	Observation	Exploration
$B : \mathscr{Y} \longrightarrow \mathscr{Q}$	Quadrature	Policy search

A lower bound for the optimal value

Bounds from both sides

Assuming that $\max_{\mathbf{s}} r(\mathbf{s}) = 0$, then

$$\forall \mathbf{s} \forall \boldsymbol{\pi} \; : \; J^{\pi}(\mathbf{s}) \leq J_*(\mathbf{s}) \leq J_{\text{max}} := \sum_{t=0}^{T} \gamma^t \cdot 0 = 0$$

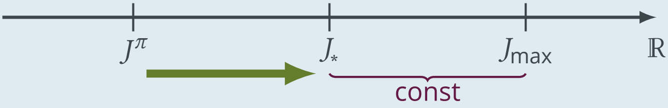

A lower bound for the optimal value

Bounds from both sides

Assuming that $\max_s r(s) = 0$, then

$$\forall s \forall \pi \;:\; J^{\pi}(s) \le J_*(s) \le J_{\max} := \sum_{t=0}^{T} \gamma^t \cdot 0 = 0$$

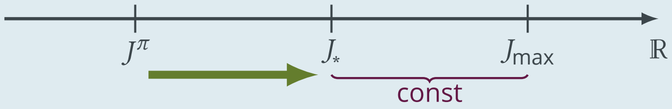

$$p(J_* \mid s_0) = \int \underbrace{p(J_* \mid J)}_{\text{Likelihood}} \underbrace{p(J \mid \pi_*, s_0, f)}_{\text{Trajectory}} \underbrace{p(\pi_*, f)}_{\text{System}} \, dJ \, d\pi_* \, df \, ds_0,$$

$$\ge \int p(J_{\max} \mid J) \, p(J \mid \pi_*, s_0, f) \, p(\pi_*, f) \, dJ \, d\pi_* \, df \, ds_0,$$

Deep GP variational bound

$$\log p(J_* \mid s_0) \geq \log \int \underbrace{p(J_{max} \mid J)}_{\text{Likelihood}} \underbrace{p(J \mid \pi_*, s_0, f)\, p(\pi_*, f)}_{\textbf{Deep GP}} \, dJ \, d\pi_* \, df \, ds_0$$

$$\geq \mathbb{E}_{q(s_0,\ldots,s_T)} \left[\log \int p(J_{max} \mid J) \underbrace{p(J \mid s_0, \ldots, s_T)}_{\text{Reward}} \, dJ \right] - \text{klterm}$$

$$= \mathbb{E}_{q(J)}[\log p(J_{max} \mid J)] - \text{klterm},$$

A variational bound for RL

Deep GP variational bound

$$\log p(J_* \mid s_0) \geq \log \int \underbrace{p(J_{max} \mid J)}_{\text{Likelihood}} \underbrace{p(J \mid \pi_*, s_0, f)\, p(\pi_*, f)}_{\textbf{Deep GP}} \, dJ \, d\pi_* \, df \, ds_0$$

$$\geq \mathbb{E}_{q(s_0, \ldots, s_T)}\left[\log \int p(J_{max} \mid J) \underbrace{p(J \mid s_0, \ldots, s_T)}_{\text{Reward}} \, dJ \right] - \text{klterm}$$

$$= \mathbb{E}_{q(J)}[\log p(J_{max} \mid J)] - \text{klterm},$$

With the exponential likelihood and $\lambda = 1$

$$p(J_{max} \mid J) := \lambda \exp(-\lambda(J_{max} - J)) = \lambda \exp(\lambda J)$$

we have

$$\mathbb{E}_{q(J)}[\log p(J_{max} \mid J)] - \text{klterm} = \mathbb{E}_{q(J)}[J] - \text{KL}(q(\pi_*) \,\|\, p(\pi_*)) + \text{const},$$

Case Study: MountainCar

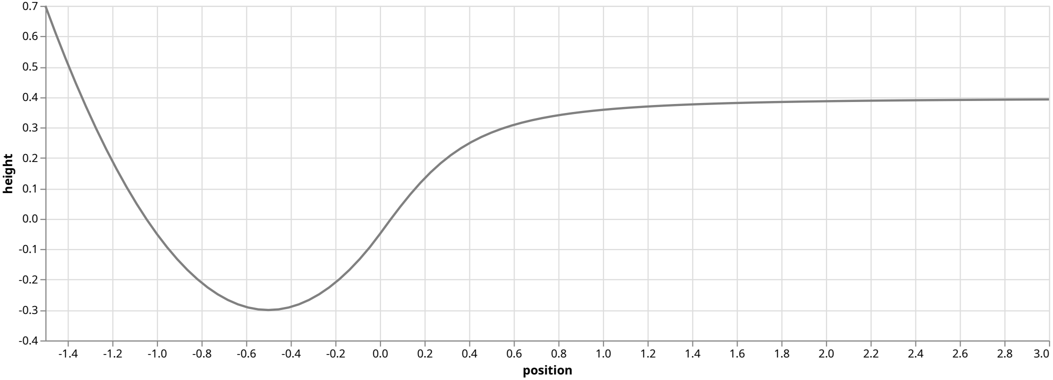

Case Study: MountainCar

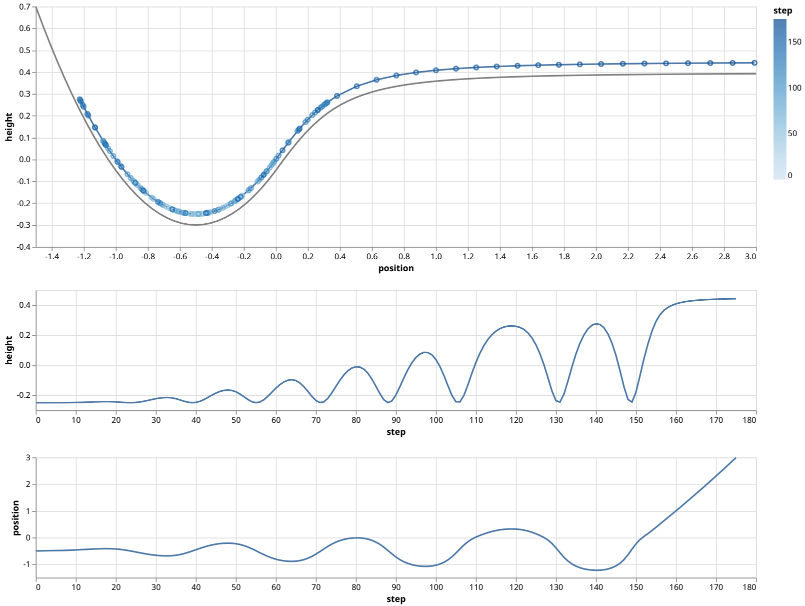

Case Study: MountainCar

After 1 step

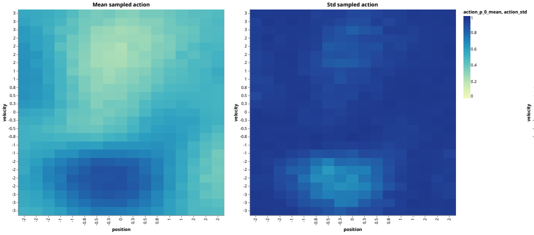

After 5 steps

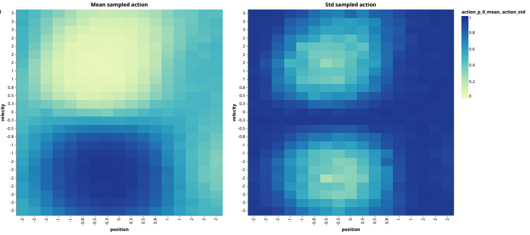

After 10 steps

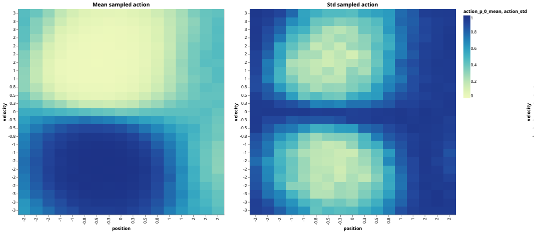

After 50 steps

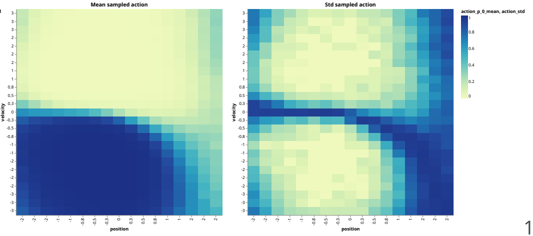

After 50 steps

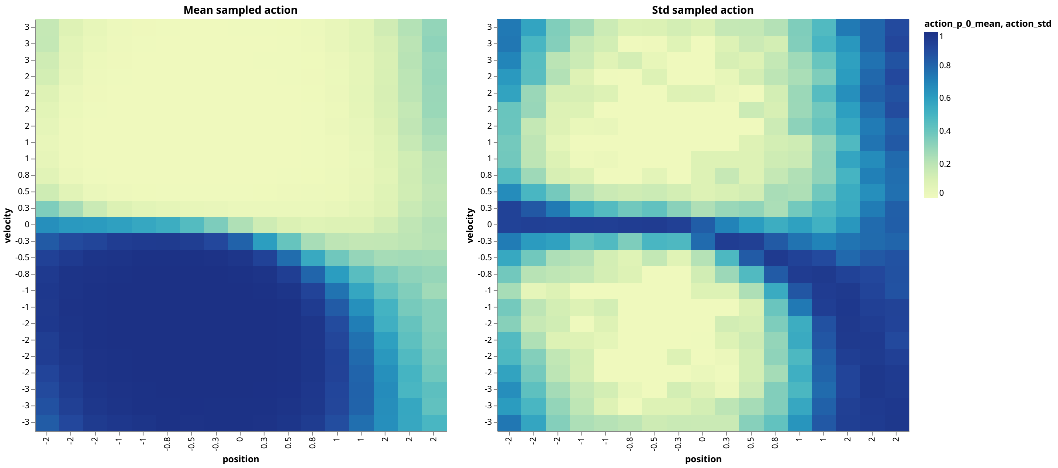

Summary

Summary

- I have a hard time reasoning about (un-)supervised learning
- Task-based uncertainties might be a way out
- Probabilistic Numerics formulates a nice framework
- Let's apply it to all the things!

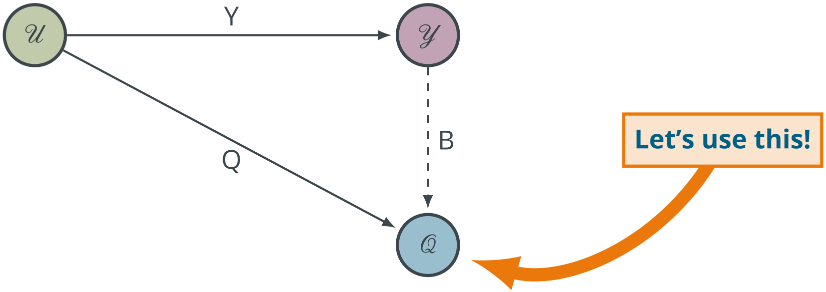

Bodin, Erik et al. (June 26, 2019). “Modulated Bayesian Optimization Using Latent Gaussian Process Models”. In: arXiv: `1906.11152 [cs, stat]`.

Kaiser, Markus et al. (2018). “Bayesian Alignments of Warped Multi-Output Gaussian Processes”. In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio et al. Curran Associates, Inc., pp. 6995–7004.

Oates, C. J. and T. J. Sullivan (Jan. 14, 2019). “A Modern Retrospective on Probabilistic Numerics”. In: arXiv: `1901.04457 [math, stat]`.