

Probabilistic Learning Theory

Why there is no free lunch, but the lunch we get is really tasty

Paul Bürkner (SimTech, University of Stuttgart)

The Black Swan of True Uncertainty

True data-generating process:

$$y \sim \text{Reality}$$

Assumed data-generating process for learning:

$$y \sim \text{Model}(x, \theta)$$

Predicted data-generating process for new data:

$$\hat{y}^{\text{new}} \sim \text{Model}(x^{\text{new}}, \hat{\theta})$$

Does \hat{y}^{new} know its true uncertainty?

Okay, let's talk about computation

The Innocent Bayes Theorem


$$p(\theta | y) = \frac{p(y | \theta) p(\theta)}{p(y)}$$

The Guilty Marginal Likelihood

$$p(y) = \int p(y | \theta) p(\theta) d\theta$$

Inference in High-Dimensional Spaces

How can we bridge the gap between learning regimes?




$$O(n^3) \text{ vs. } O(nm^d + m^d)$$

Riutort-Mayol et al. (in review)

Global vs. Local Regularization


How global is local regularization in high-dimensional spaces?



Aguilar & Bürkner (in review)

Efficient Cross-Validation

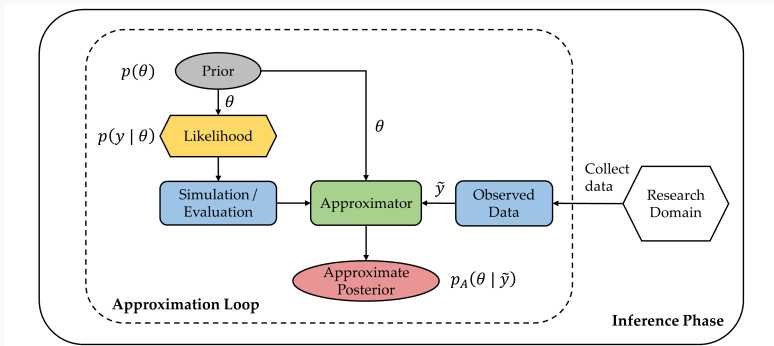
How can we reduce the computational burden of cross-validation?



Bürkner et al. (2020)

Non-Amortized (Standard) Inference

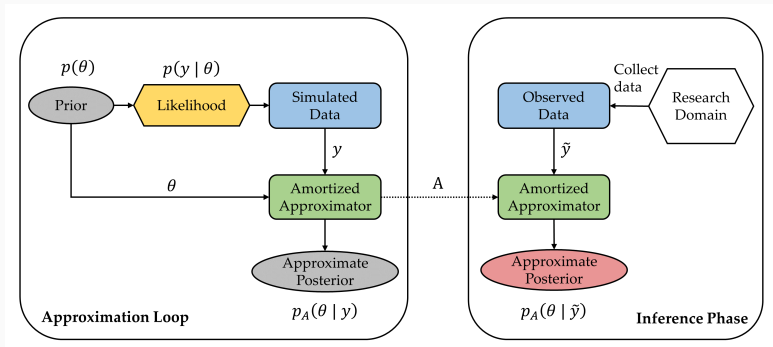
How can we break out of the standard inference setting?



Bürkner et al. (in review)

Amortized Inference


How far can we scale amortized inference?



Bürkner et al. (in review)

Uncertainty of Uncertainty (Meta-Uncertainty)

How can we combine Bayesian and frequentist uncertainties?



Schmitt et al. (in review soon)