

A Statistical Perspective on SBI

Or: SBI for spherical cows

Paul Bürkner

TU Dortmund University, Germany

 <https://paul-buerkner.github.io/>

Our Team



Stefan T. Radev



Marvin Schmitt



Lars Kühmichel



Daniel Habermann



Simon Kucharsky



Lasse Elsemüller



Ullrich Köthe



Florence Bockting



Philipp Reiser



Jacob Grytzka



Soham Mukherjee



Maximilian Scholz

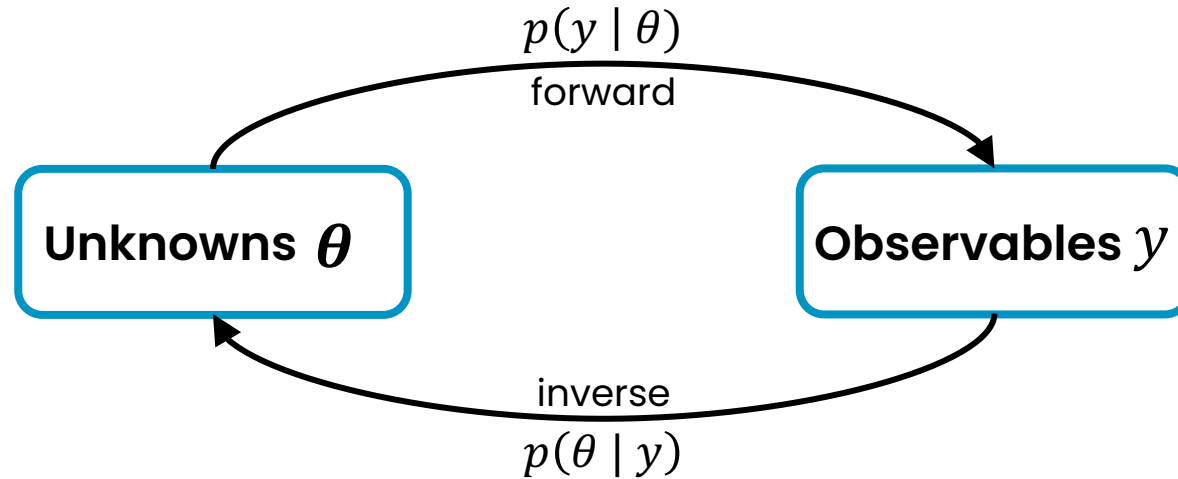


Luna Fazio



Javier Aguilar

Inverse Problems



Statistical modeling: **Parameters θ**

Data y

Epidemiology: Virus attributes

Infection curve (time series)

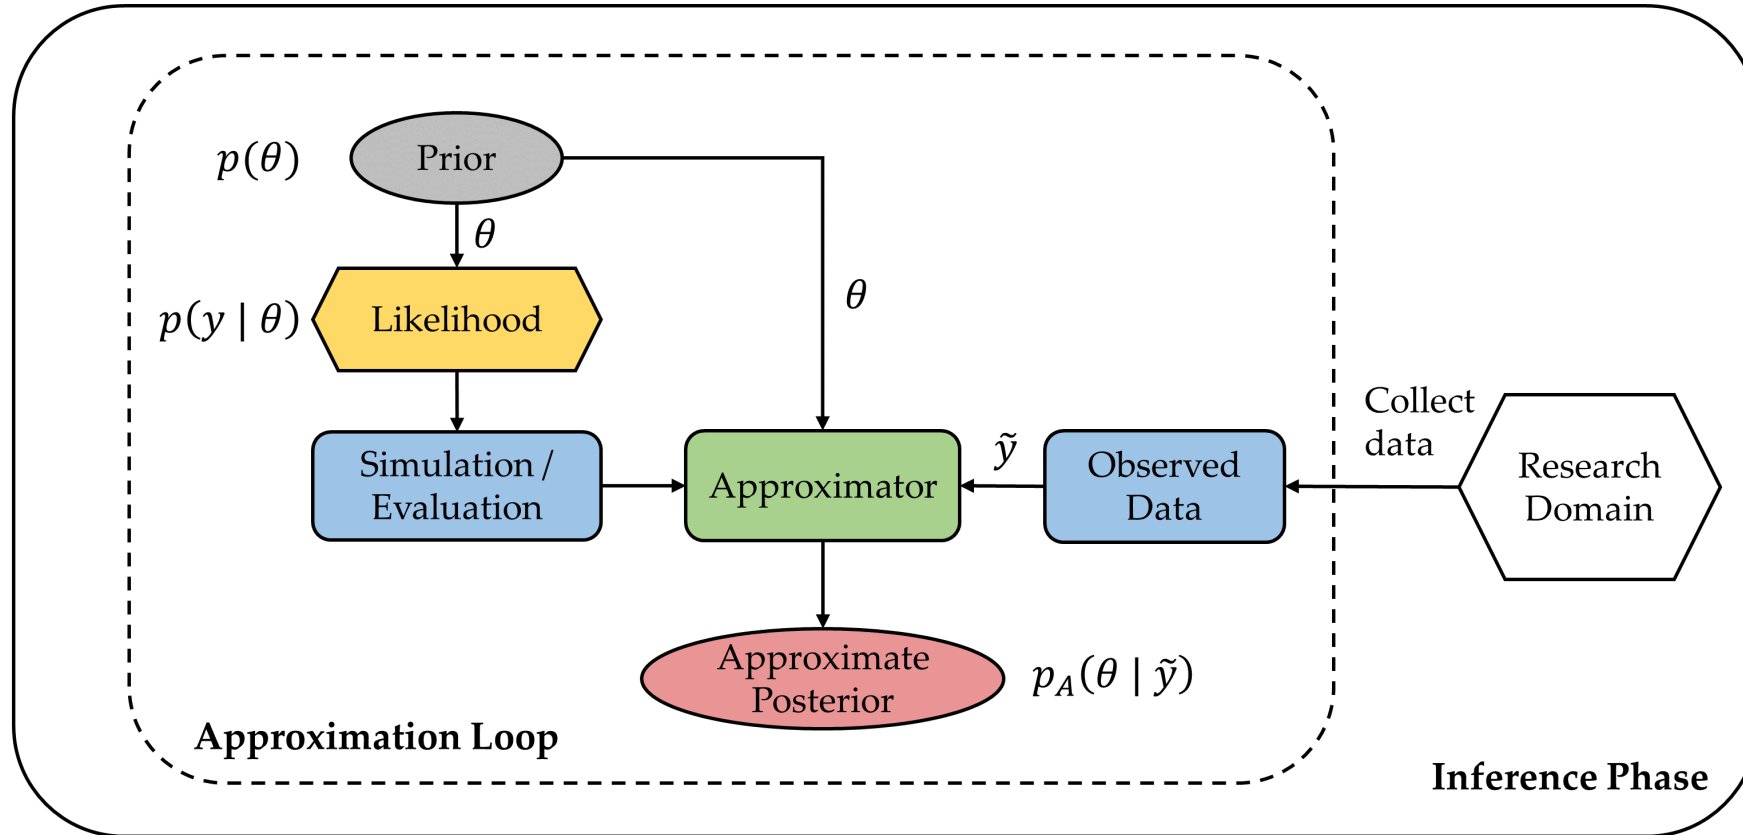
Image processing: Crisp image

Blurry image

Psychology: Cognitive parameters

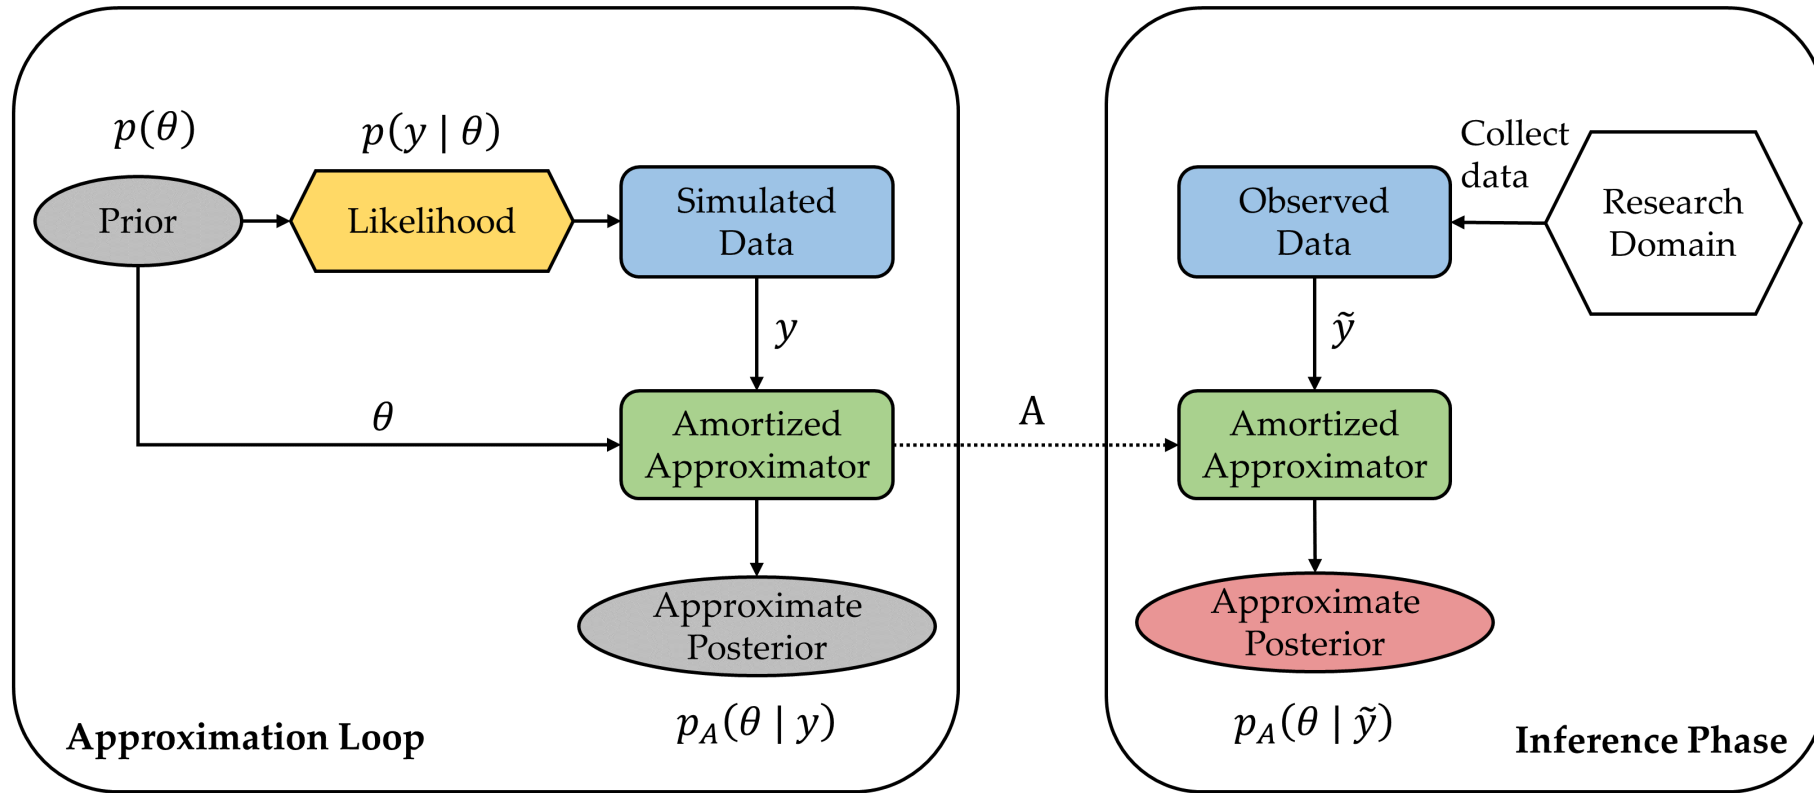
Reaction times

Non-amortized Bayesian inference



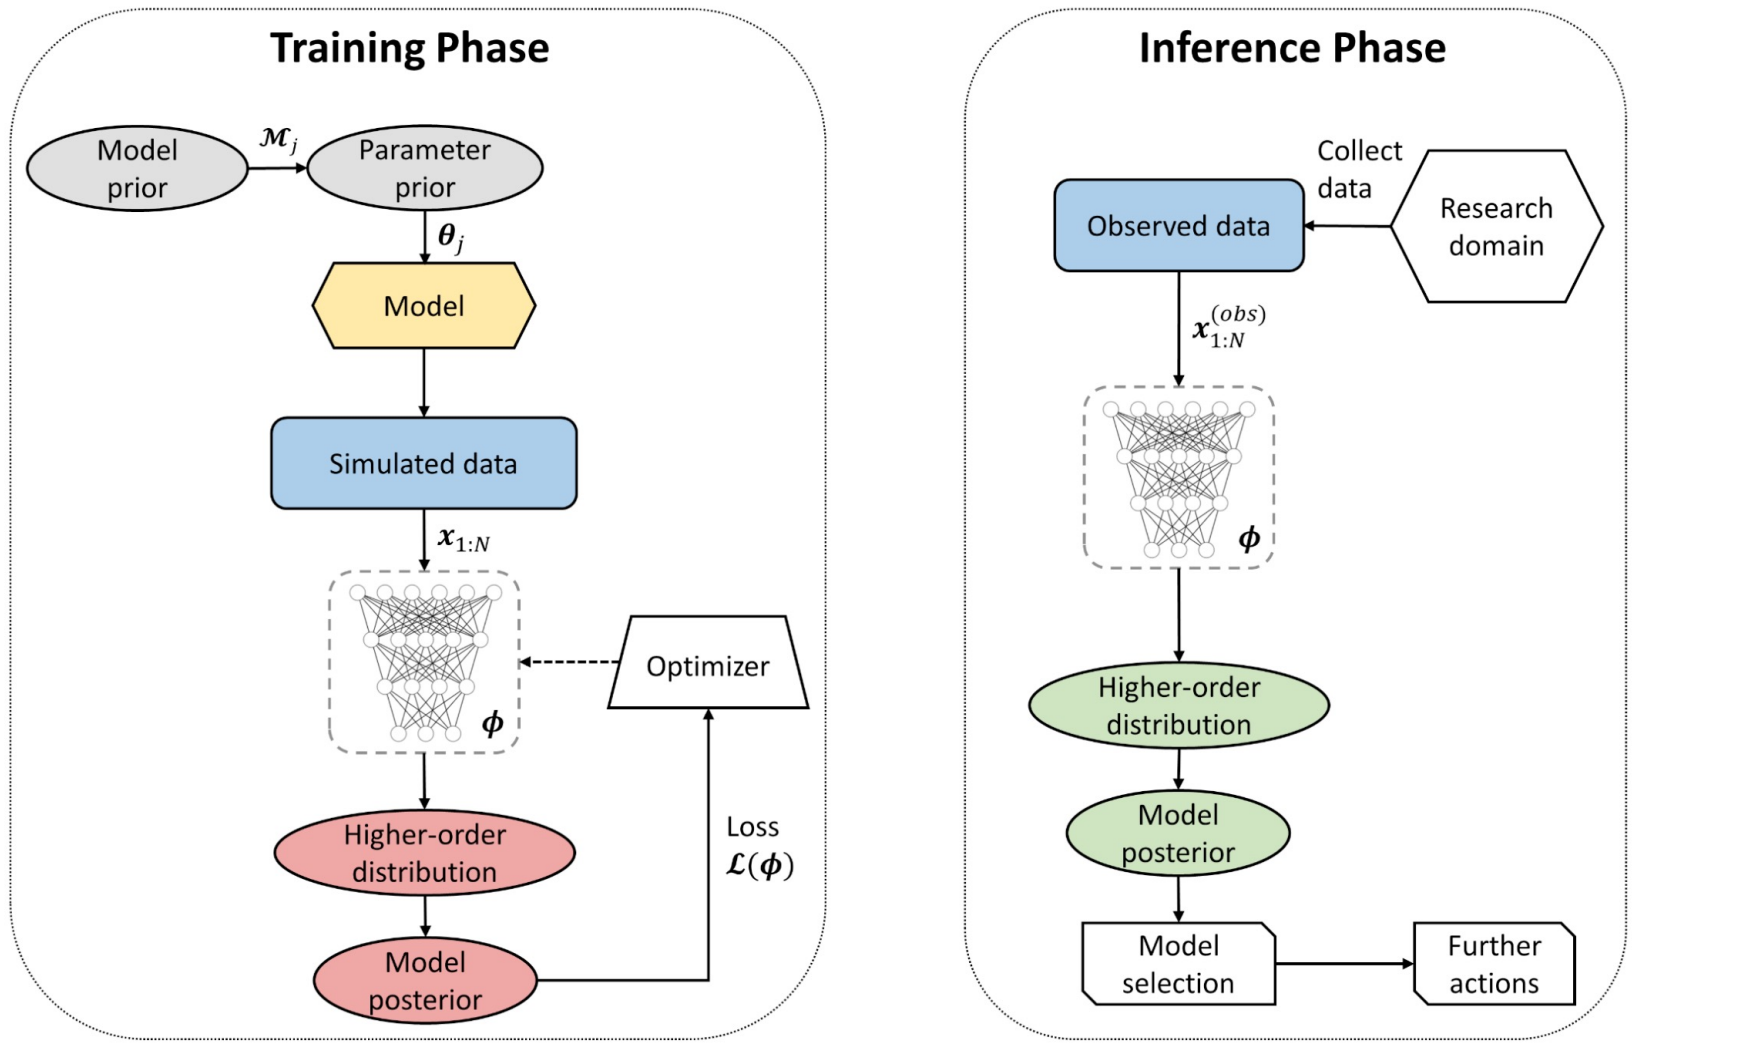
Approximation and inference are **coupled**. No resource pooling.

Amortized Bayesian inference (ABI)



Approximation and inference are **decoupled**. Pooling of resources.

Amortized Model Comparison



Loss Functions: Variational Inference

Backward KL divergence

$$\text{minimize } KL(p_A(\theta | \tilde{y}) || p(\theta | \tilde{y})) = \int \log \left(\frac{p_A(\theta | \tilde{y})}{p(\theta | \tilde{y})} \right) p_A(\theta | \tilde{y}) d\theta$$

$$\iff \text{maximize } \frac{1}{S} \sum_{s=1}^S \log \left(\frac{p(\theta^{(s)}, \tilde{y})}{p_A(\theta^{(s)} | \tilde{y})} \right) \quad \text{for } \theta^{(s)} \sim p_A(\theta | \tilde{y})$$

The ELBO just requires the joint density of the model

Loss Functions: Simulation-Based Inference

Forward KL divergence

$$\text{minimize } KL(p(\theta | \tilde{y}) || p_A(\theta | \tilde{y})) = \int \log \left(\frac{p(\theta | \tilde{y})}{p_A(\theta | \tilde{y})} \right) p(\theta | \tilde{y}) d\theta$$

$$\iff \text{maximize } \frac{1}{S} \sum_{s=1}^S \log p_A(\theta^{(s)} | \tilde{y}) \quad \text{for } \boxed{\theta^{(s)} \sim p(\theta | \tilde{y})}$$

Minimizing the expected KL over the whole data space

$$\text{minimize } E_{p(y)} [KL(p(\theta | y) || p_A(\theta | y))]$$

$$\iff \text{maximize } \frac{1}{S} \sum_{s=1}^S \log p_A(\theta^{(s)} | y^{(s)}) \quad \text{for } \boxed{(\theta^{(s)}, y^{(s)}) \sim p(\theta, y)}$$

How to obtain the posterior approximator?

Sample $z^{(s)} \sim \text{multinormal}(0, I)$

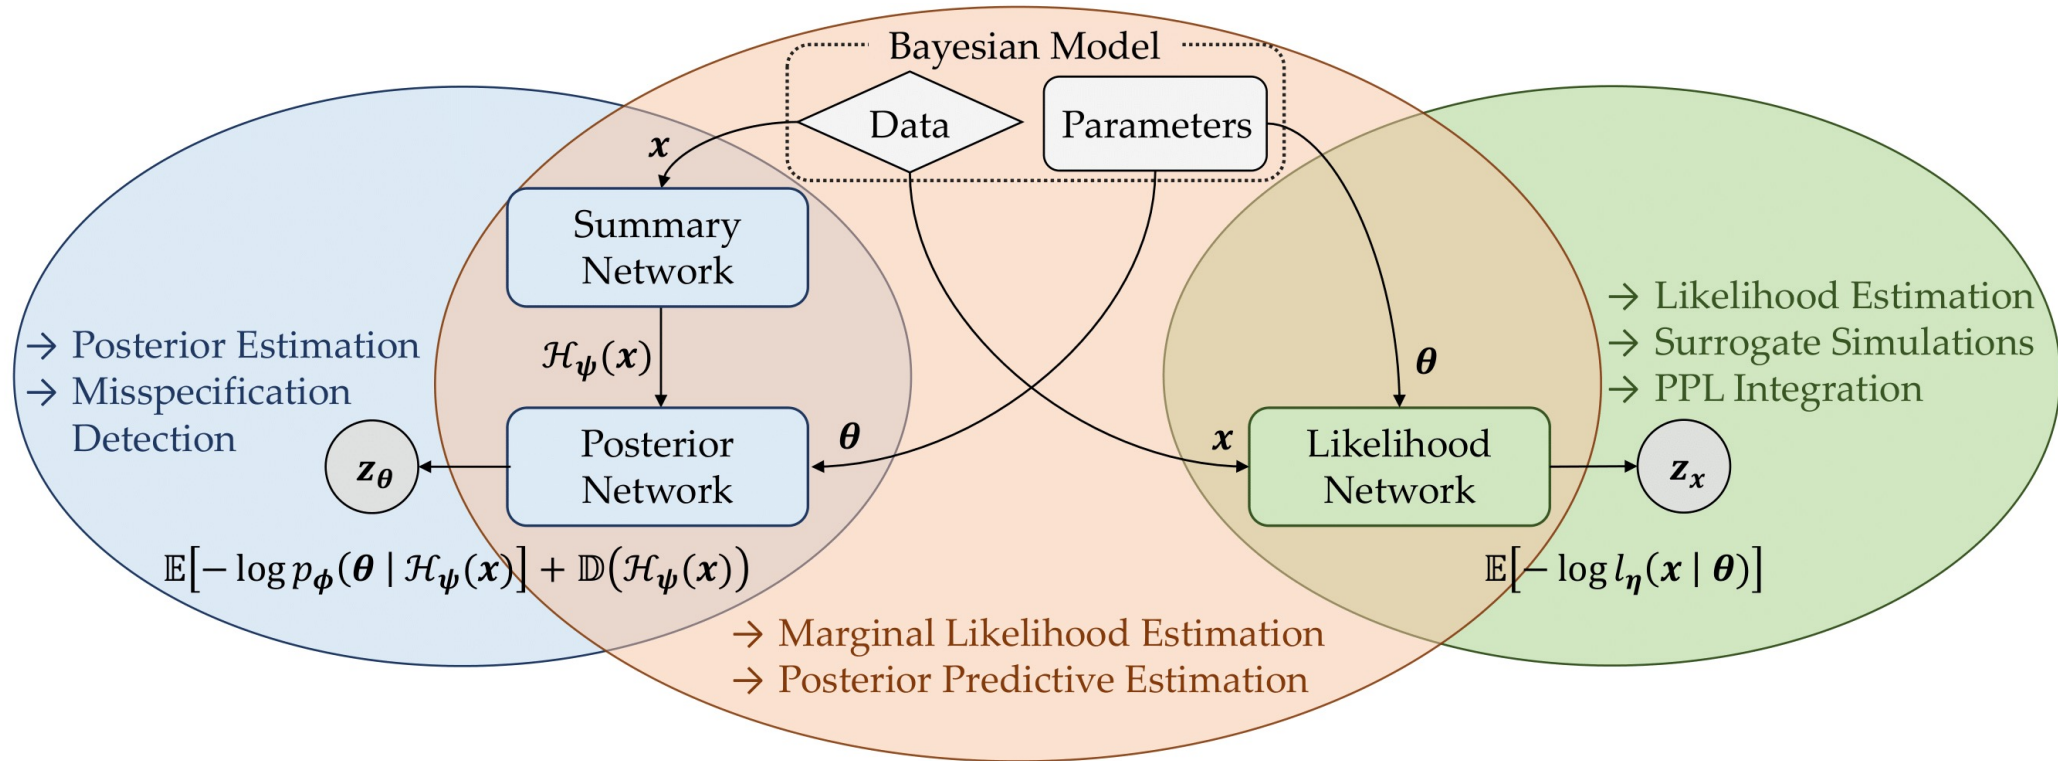
Transform to $\theta^{(s)} = f_\phi(z^{(s)} | y)$ with an invertible neural network

Obtain the approximator's density for training via expected KL divergence:


$$p_A(\theta | y) = p(z = f_\phi^{-1}(\theta | y)) \left| \det \left(\frac{\delta f_\phi^{-1}(\theta | y)}{\delta \theta} \right) \right|$$

Jointly amortized learning

- Jointly amortized neural approximation (JANA)



Isn't amortized inference wasteful?



Amortized methods perform on-par with non-amortized counterparts

Potential of Amortized Bayesian Inference

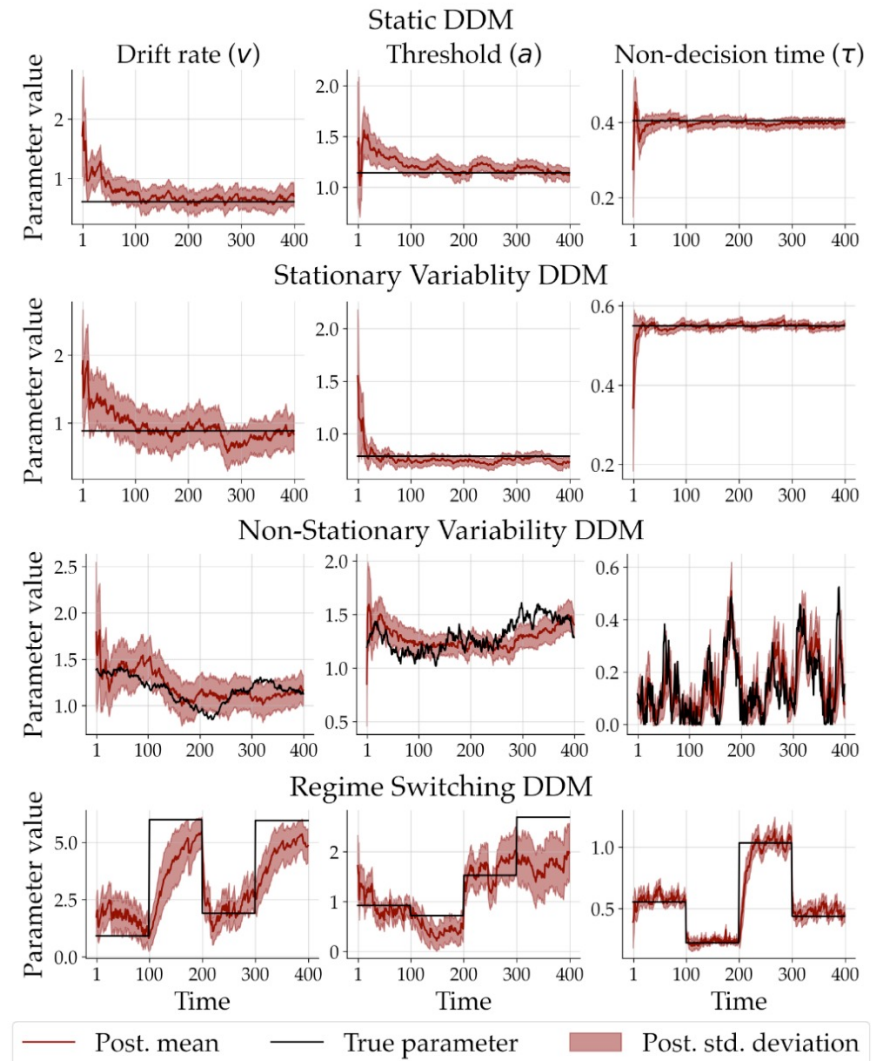
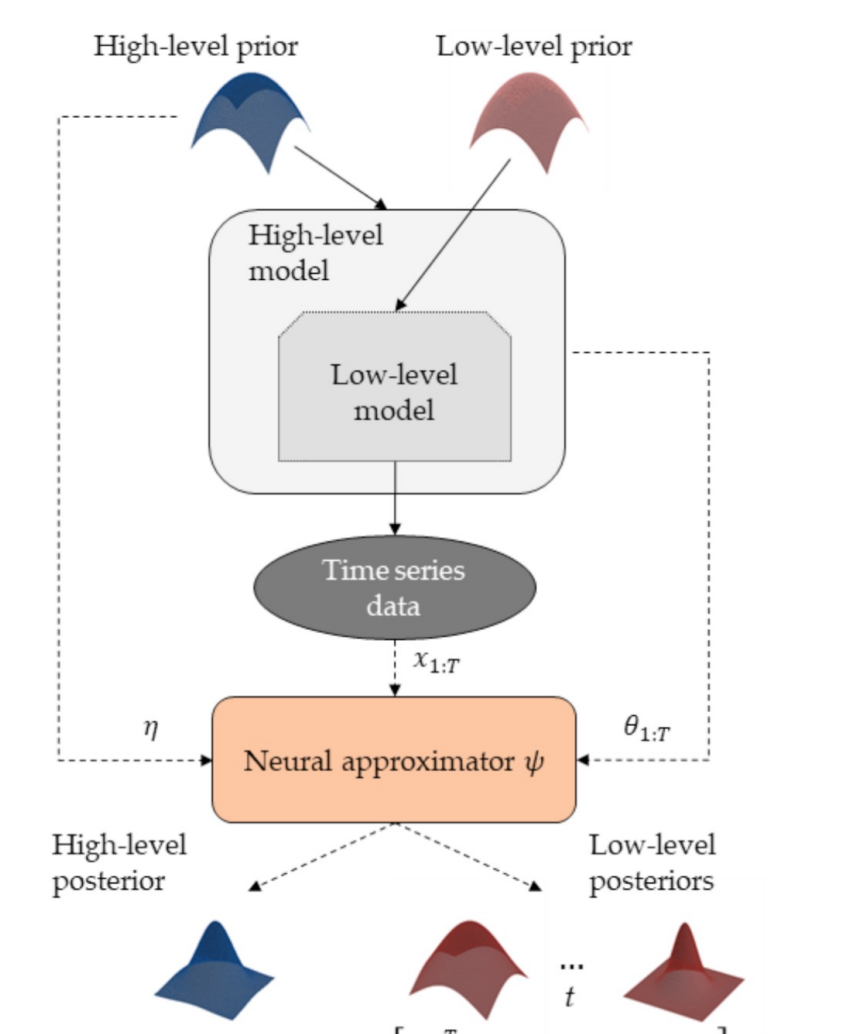
Massive number of inference repetitions

- Many data sets
- Cross-validation
- Sensitivity analyses, multiverse analyses

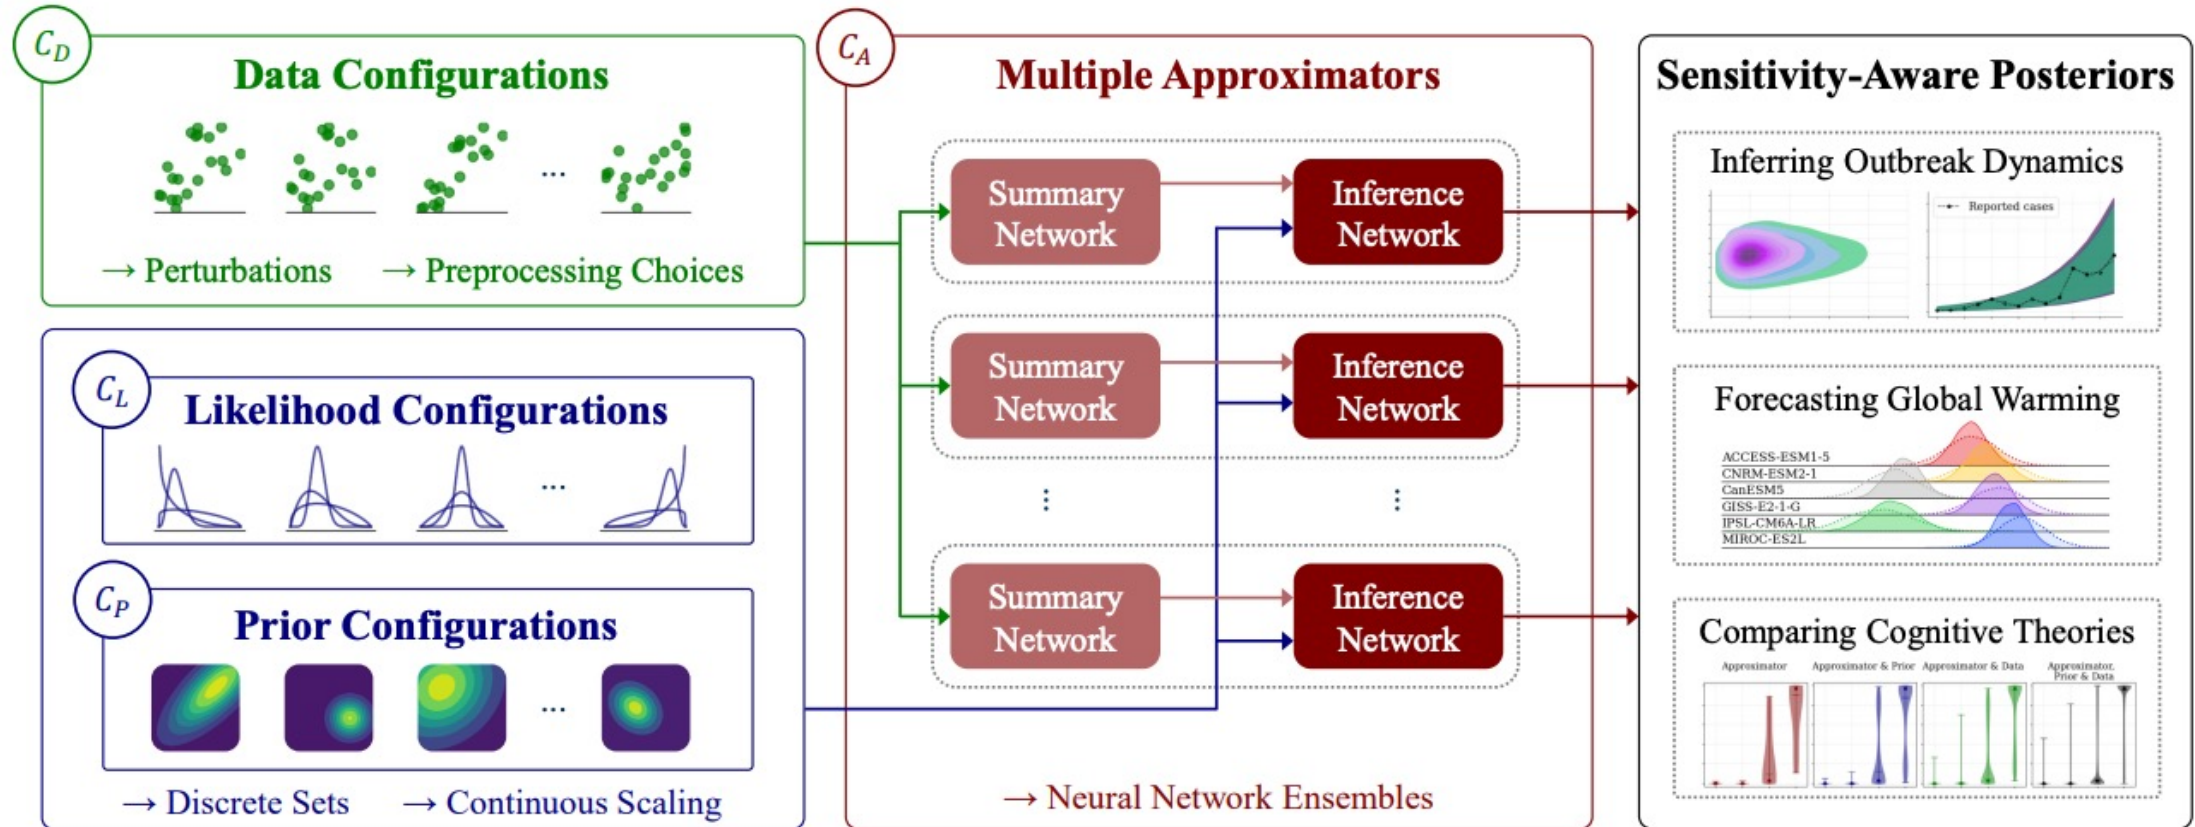
Real-time inference

- Neurological monitoring
- Adaptive experimental design
- Disease surveillance

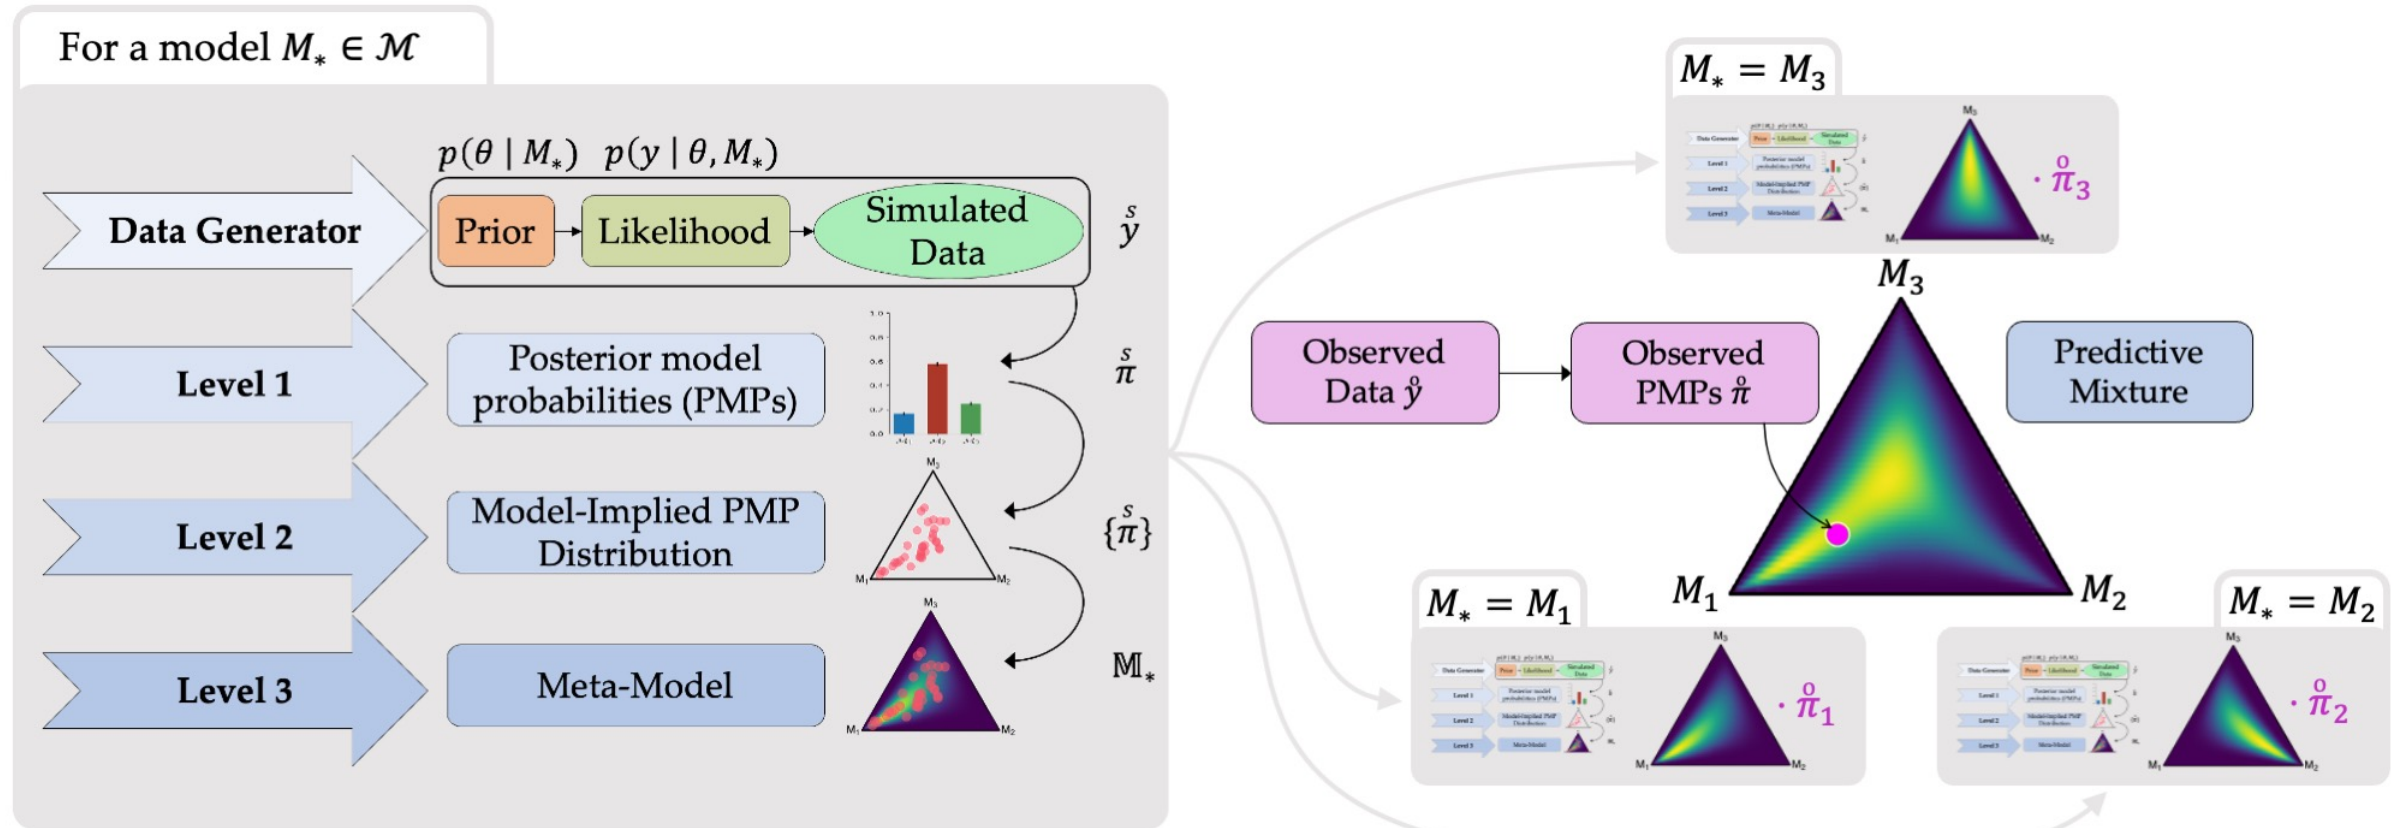
Dynamic Hierarchical Modeling



Amortized sensitivity analyses



Meta-uncertainty in Bayesian model comparison



Key challenges of ABI

- Neural networks have a bad user experience
- Heaps of simulated training data necessary
- Constrained neural network architecture of normalizing flows
- Priors are more important for training than I would like
- Model misspecification invalidates simulation-based training

ABI library: BayesFlow

BayesFlow

Tests **passing** License **MIT** JOSS **10.21105/joss.05702** contributions **welcome**



Welcome to our BayesFlow library for efficient simulation-based Bayesian workflows! Our library enables users to create specialized neural networks for *amortized Bayesian inference*, which repay users with rapid statistical inference after a potentially longer simulation-based training phase.

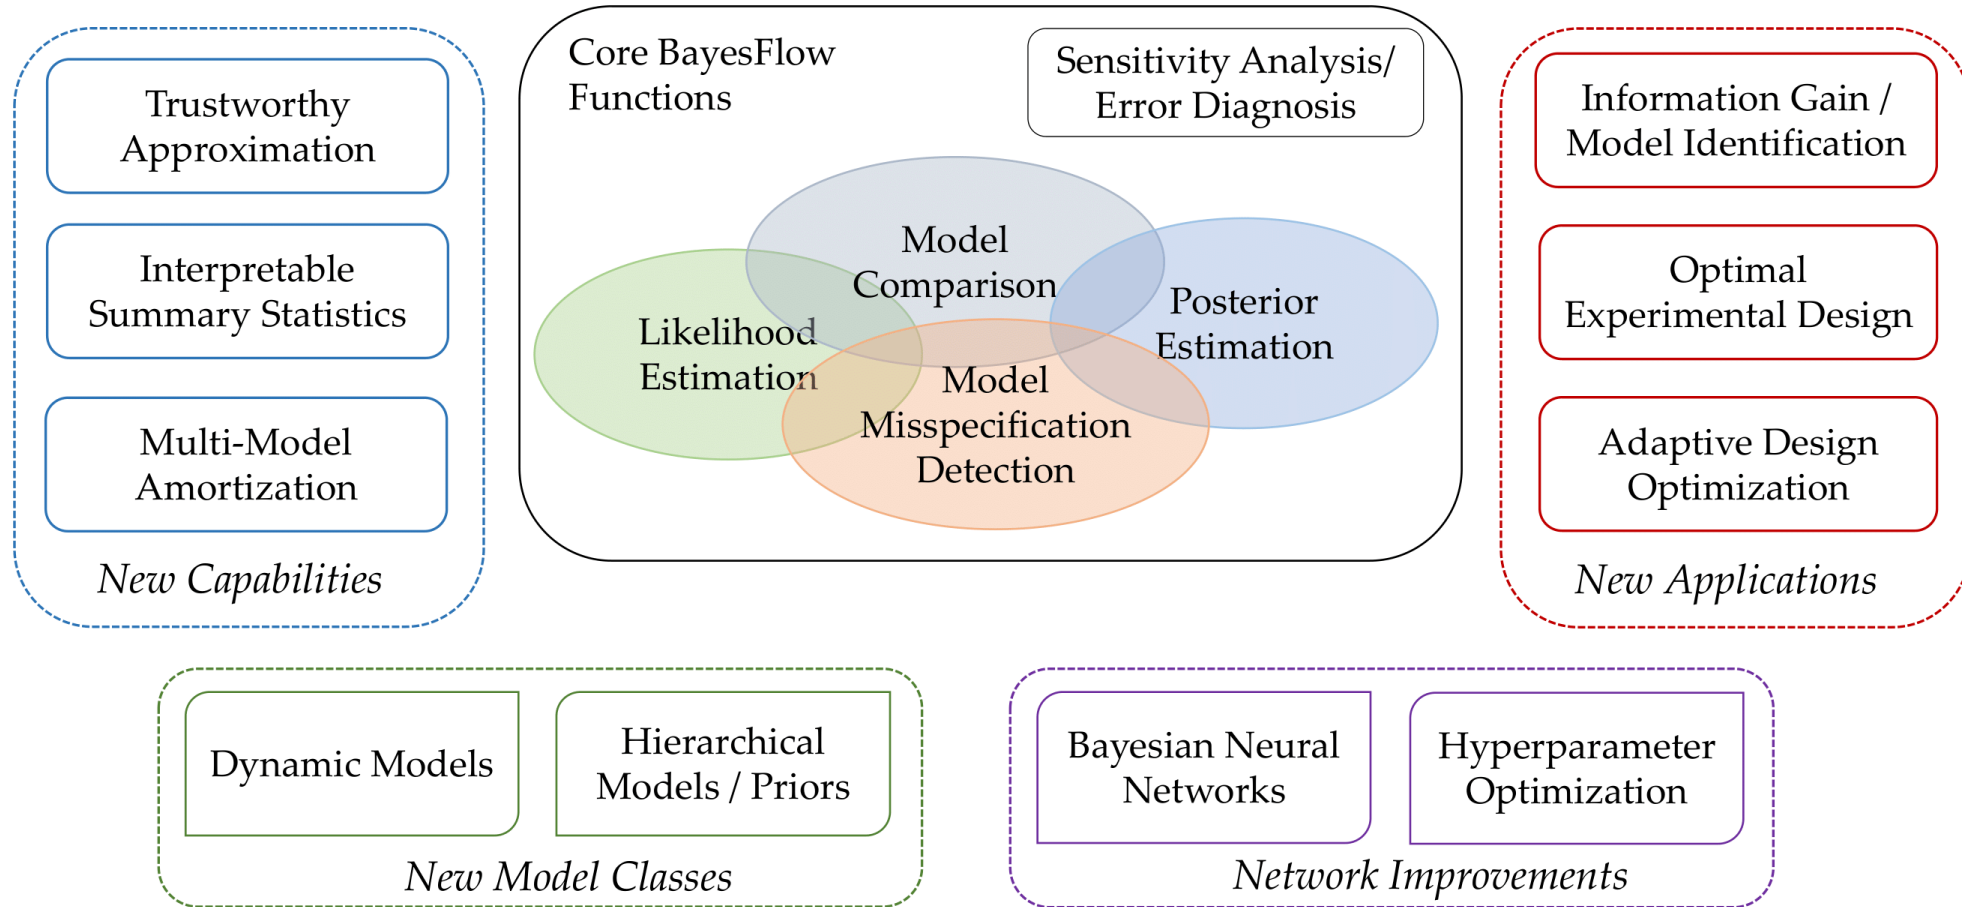
Installation

```
pip install bayesflow
```

Documentation + Support

- www.bayesflow.org
- discuss.bayesflow.org

The BayesFlow framework



Our new BayesFlow interface

- Backend agnostic via Keras 3 (Tensorflow, Pytorch, JAX)
- Modular and extensible
- Improved and simplified user interface
- Sensible network and training defaults

Current WIP branch:

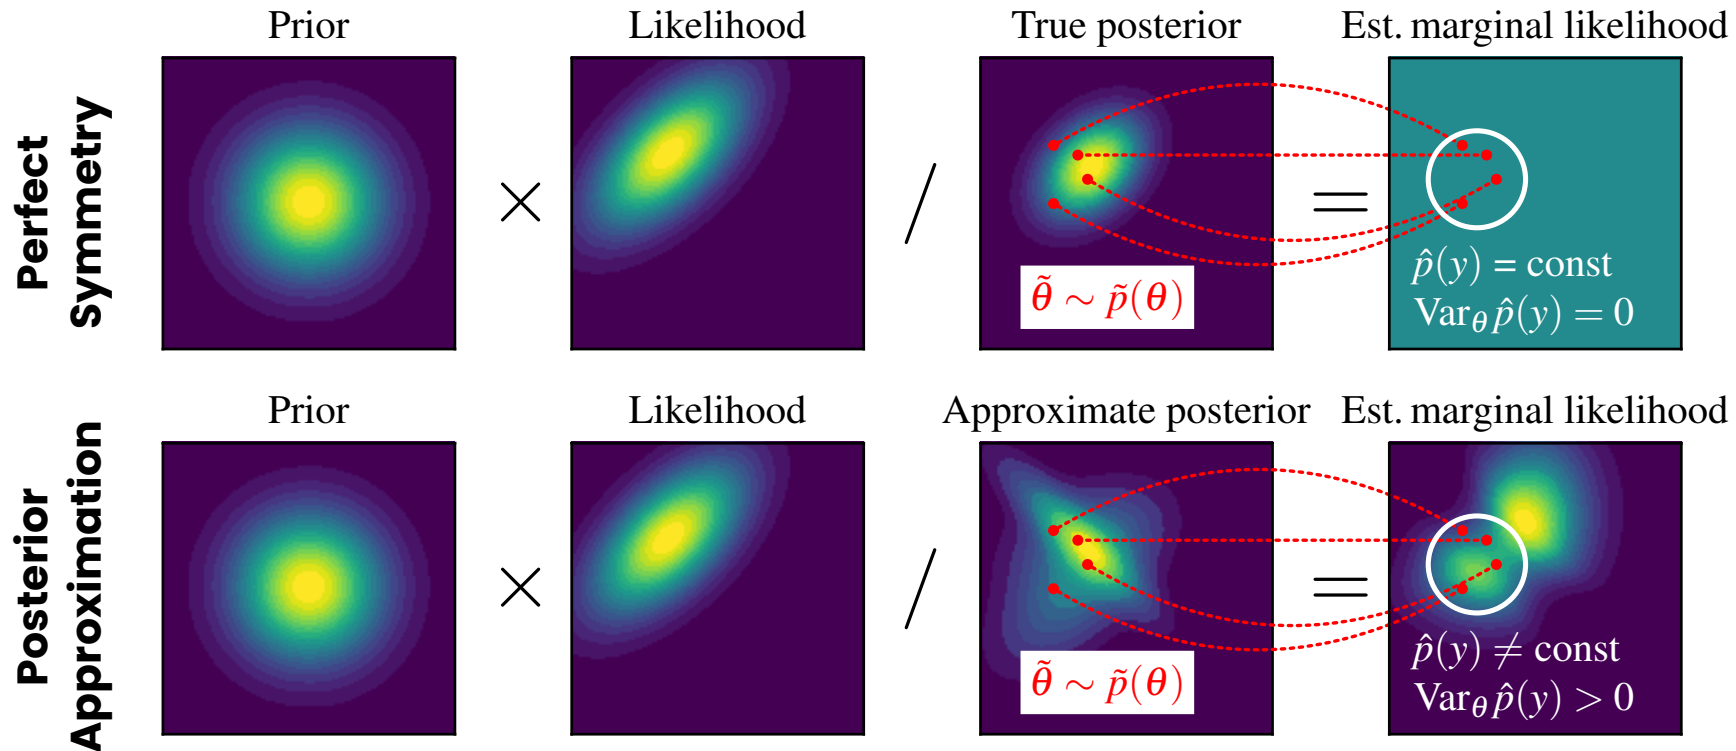
<https://github.com/stefanradev93/BayesFlow/tree/streamlined-backend>

- We will merge the new interface soon
- We would love you to use bayesflow and help us further improve it



Low data \rightarrow self-consistency

$$p(y) = p(\theta) p(y | \theta) / p(\theta | y) \implies \frac{p(\tilde{\theta}_1) p(y | \tilde{\theta}_1)}{p(\tilde{\theta}_1 | y)} = \dots = \frac{p(\tilde{\theta}_K) p(y | \tilde{\theta}_K)}{p(\tilde{\theta}_K | y)} \quad \tilde{\theta}_1, \dots, \tilde{\theta}_K \in \Theta$$



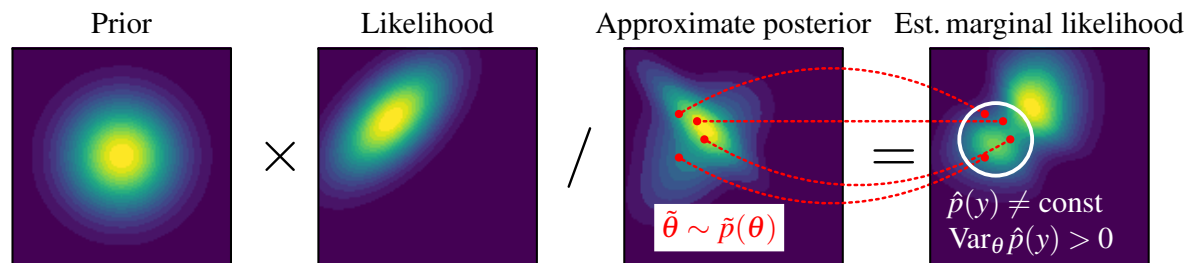
Low data \rightarrow self-consistency

- Idea: Violations of self-consistency property as loss function


$$\mathcal{L}_{\text{SC}} := \mathbb{E}_{p(y)} \left[\text{Var}_{\tilde{\theta} \sim \tilde{p}(\theta)} \left(\log p(\tilde{\theta}) + \log p(y | \tilde{\theta}) - \log q_{\phi}(\tilde{\theta} | y) \right) \right]$$

- Integration into standard neural posterior estimation loss


$$\mathcal{L}_{\text{NPE-SC}} := \mathbb{E}_{p(y)} \left[\underbrace{\mathbb{E}_{p(\theta | y)} [-\log q_{\phi}(\theta | y)]}_{\text{NPE loss}} + \underbrace{\lambda \text{Var}_{\tilde{\theta} \sim \tilde{p}(\theta)} \left(\log p(\tilde{\theta}) + \log p(y | \tilde{\theta}) - \log q_{\phi}(\tilde{\theta} | y) \right)}_{\text{self-consistency loss } \mathcal{L}_{\text{SC}} \text{ with weight } \lambda \in \mathbb{R}_+} \right]$$



Neural network constraints \rightarrow consistency models



Model misspecification → detection



Summary

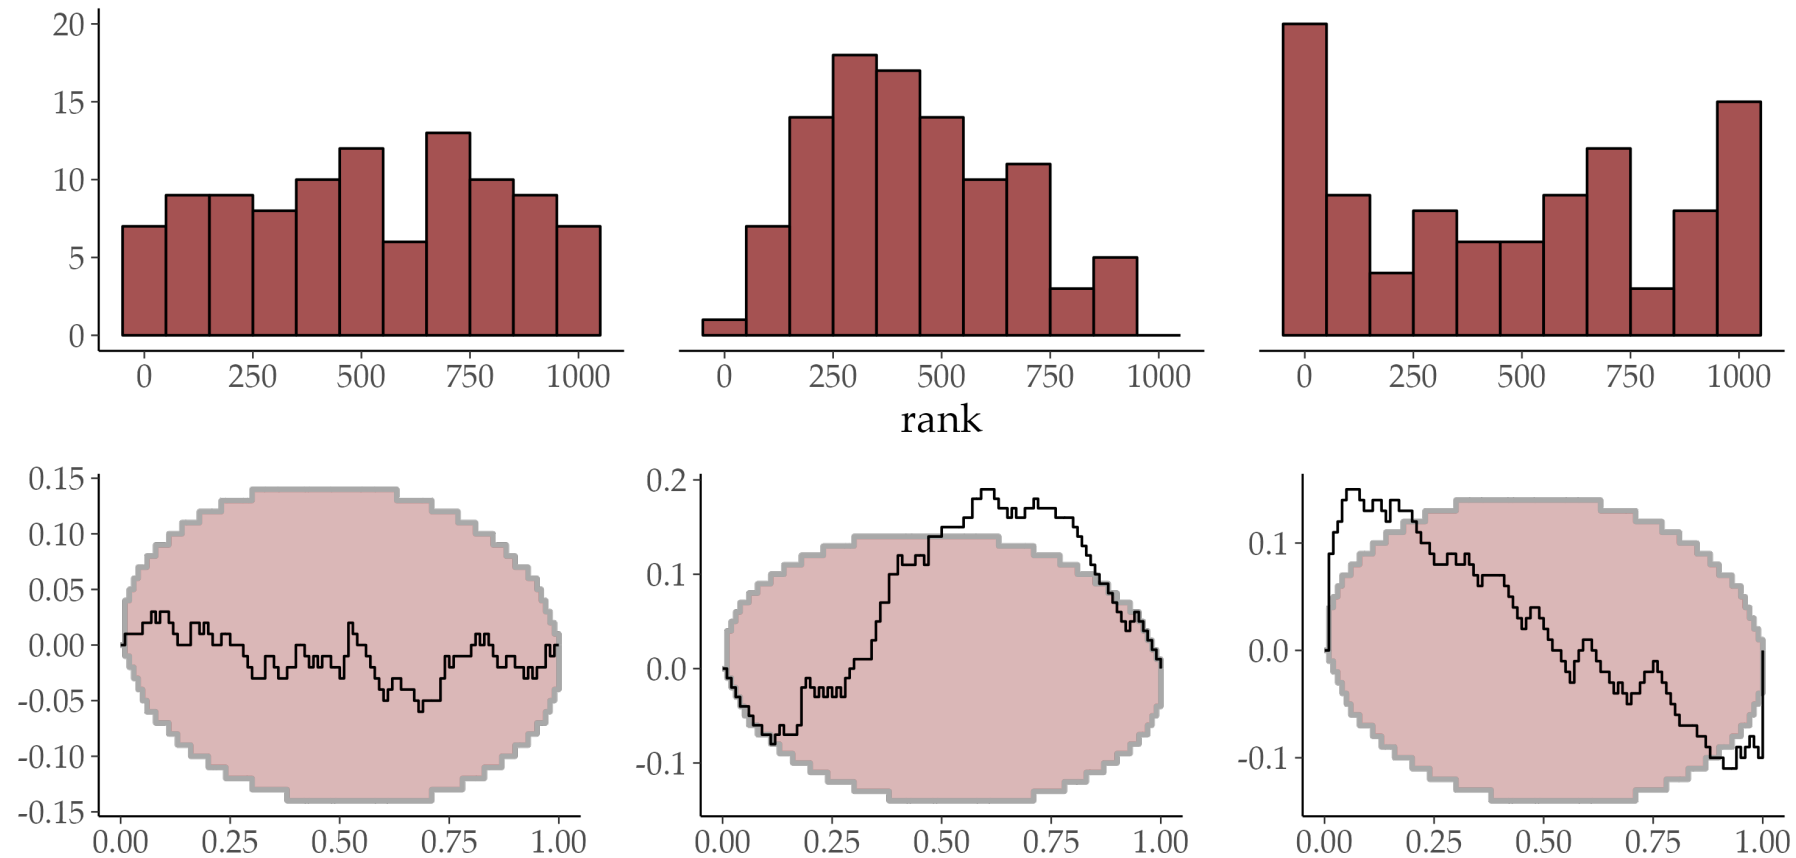
Amortized Bayesian inference

- High potential for large scale applicability
- Biggest issues: Reliability and trustworthiness
- Some questions have been tackled
- A lot of questions remain open

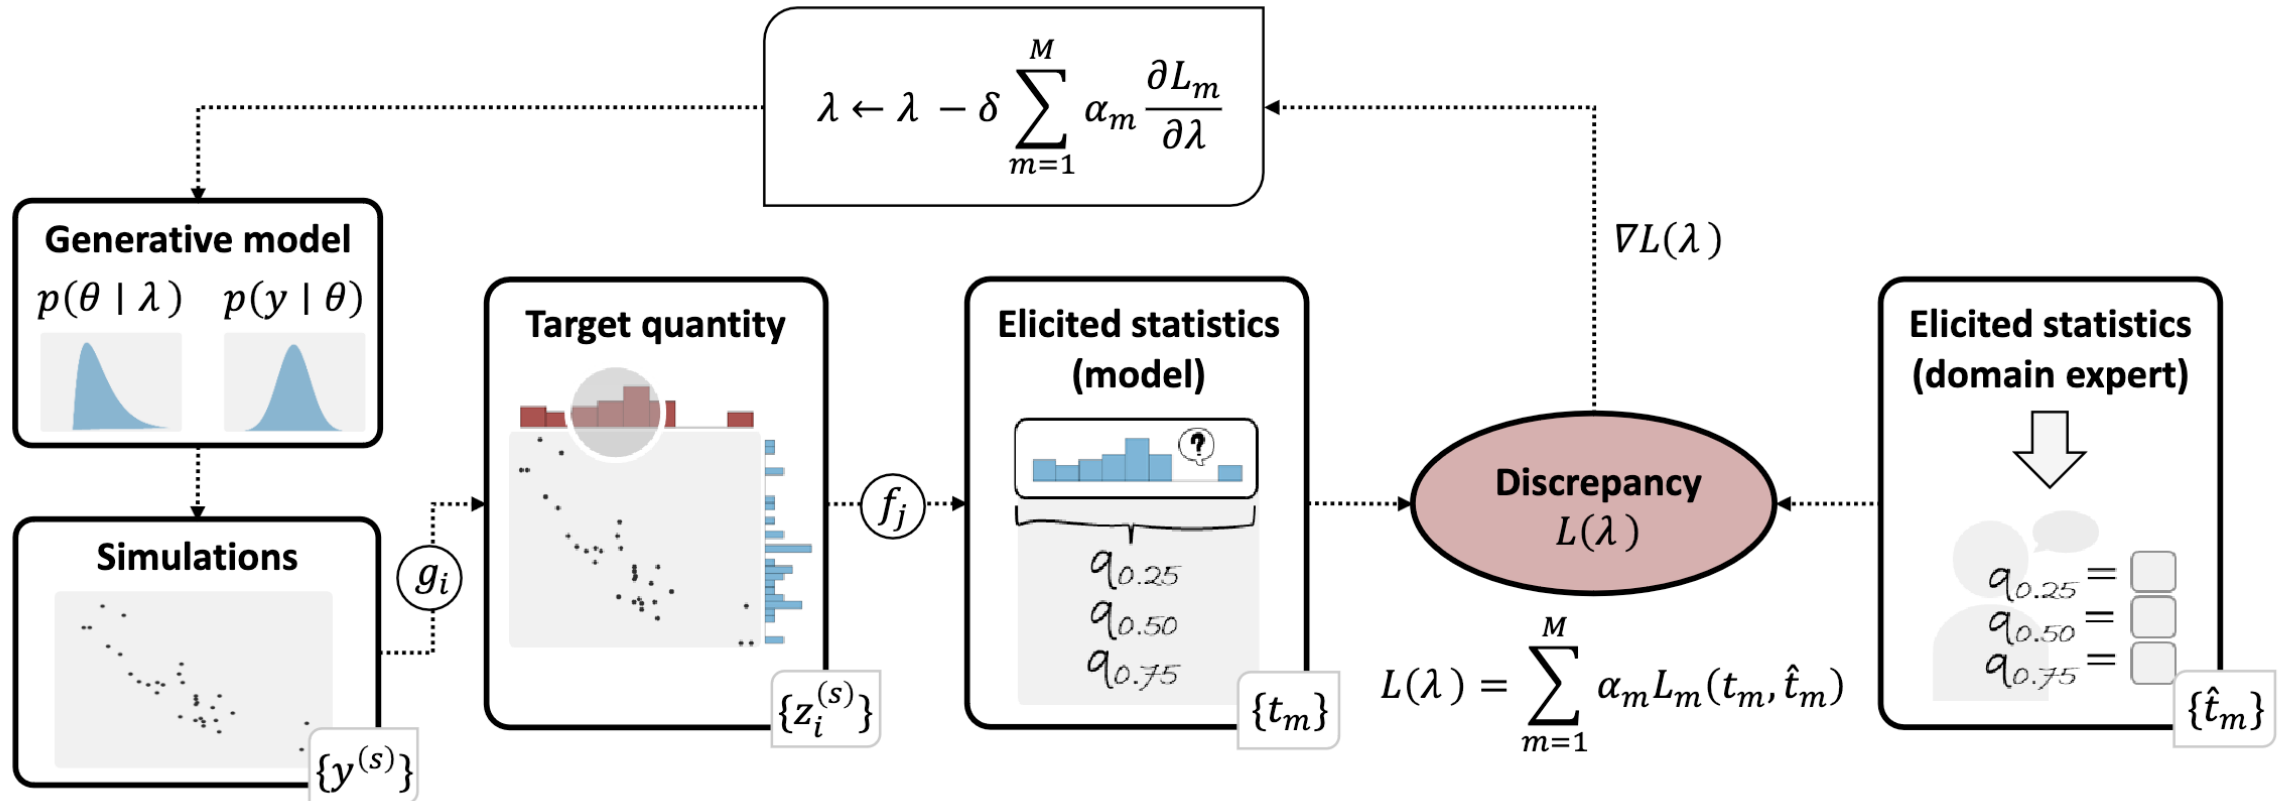
If you are interested in working with us, please reach out!

Appendix

Simulation-based calibration as convergence check



Simulation-Based Prior Elicitation



Fusing heterogeneous data sources

