NPE-PF: Simulation-based inference with foundation models

Julius Vetter, Manuel Gloeckler, Daniel Gedon, Jakob H. Macke

Simulation-based inference (SBI)

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}) \qquad x \sim p(x|\theta)$$

prior → simulator → NPE → posterior

Sequentially adapt proposal
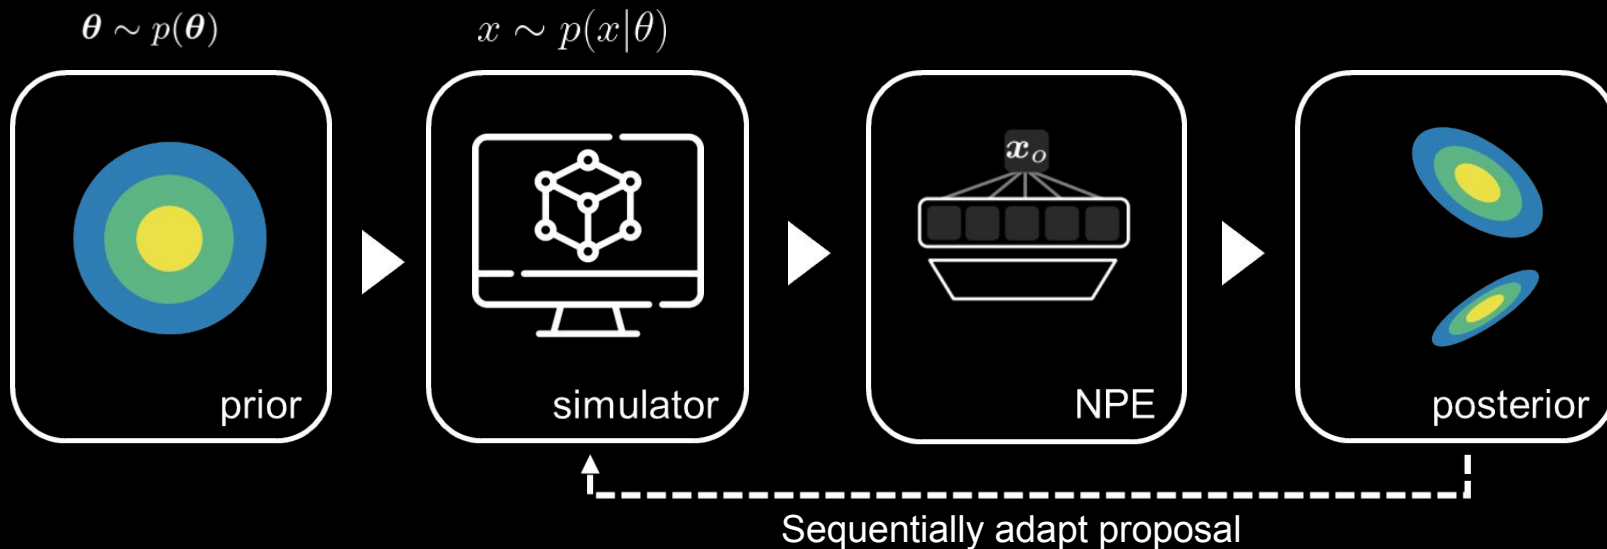
Pros:

- Likelihood-free i.e. simulator as black box.
- Fast amortized inference after training.

Cons:

- Needs a lot of simulations
- Needs hyperparameter tuning
- Needs (re-)training for any pair of prior and simulator

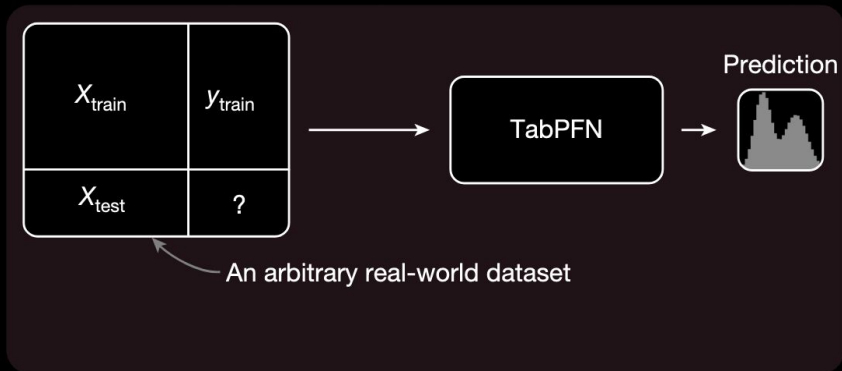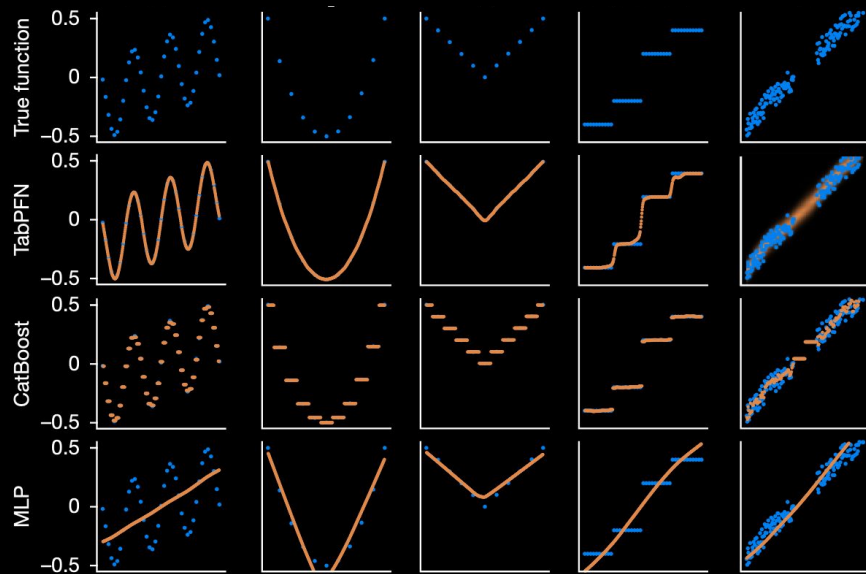# TabPFN: In-context learning for tabular regression/classification


Figure from Hollmann et al. (2025)

TabPFN: In-context estimator for one-dimensional densities!
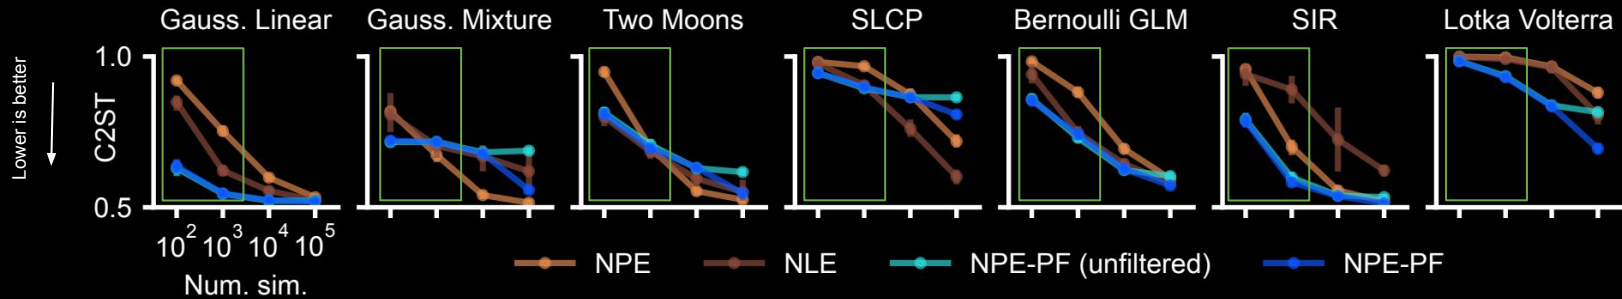
NPE-PF: SBI using TabPFN

How to use TabPFN for SBI?

- Autoregressively sample from posterior

- Predict parameter for new observation

  using in-context learning.

- Increase context by filtering simulations

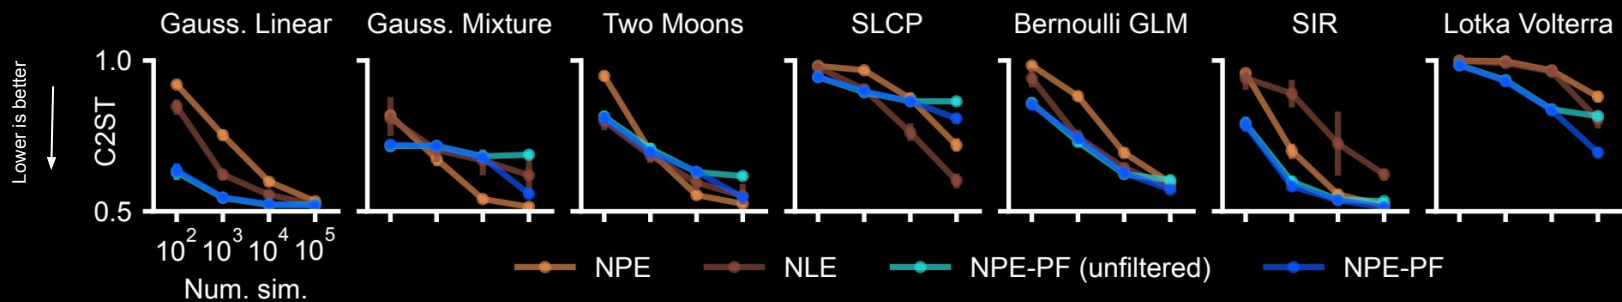- Sequential SBI: TabPFN classifier for fast density ratios
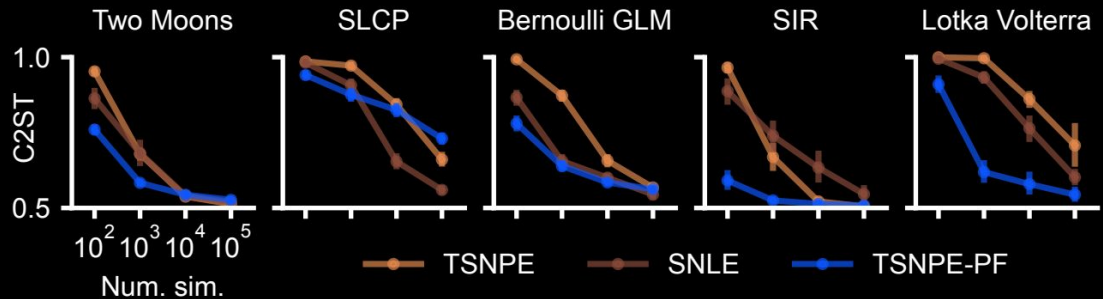- No training!
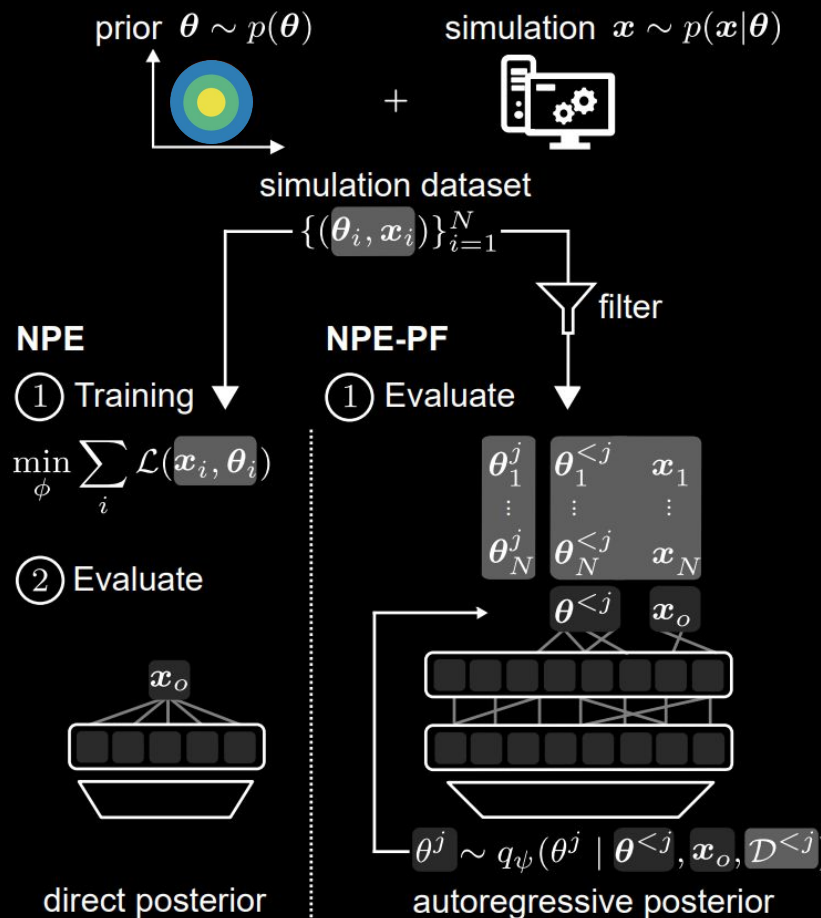
Results: SBI Benchmark

- Solid performance across standard benchmark tasks.

- Strong performance for a small number of simulations
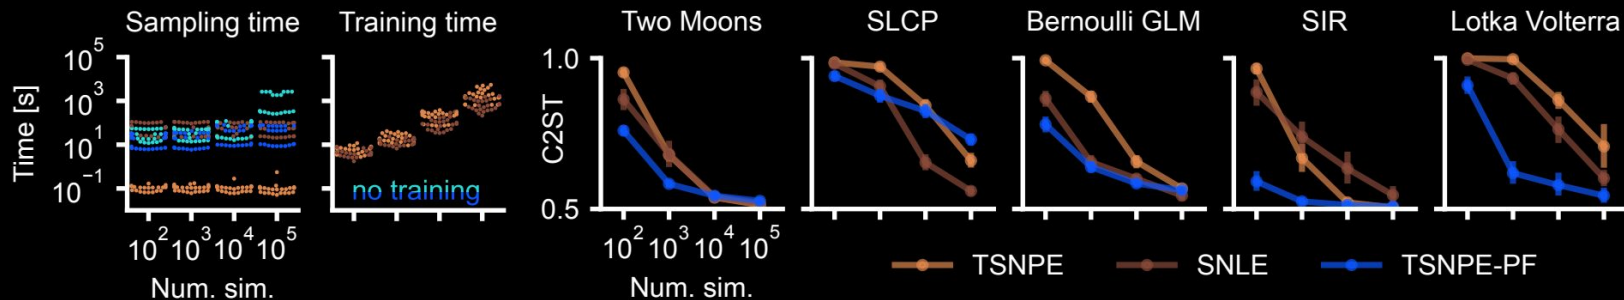
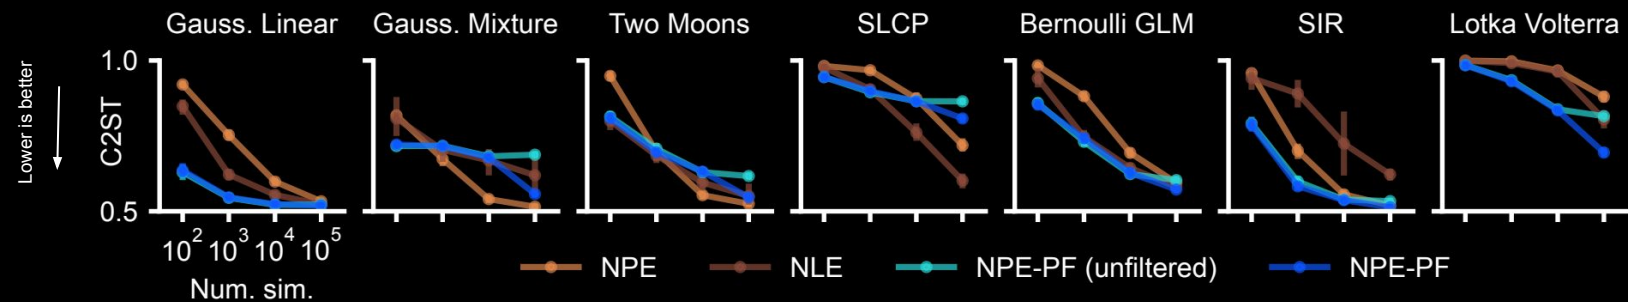▶ Quite **simulation-efficient** !

Results: SBI Benchmark

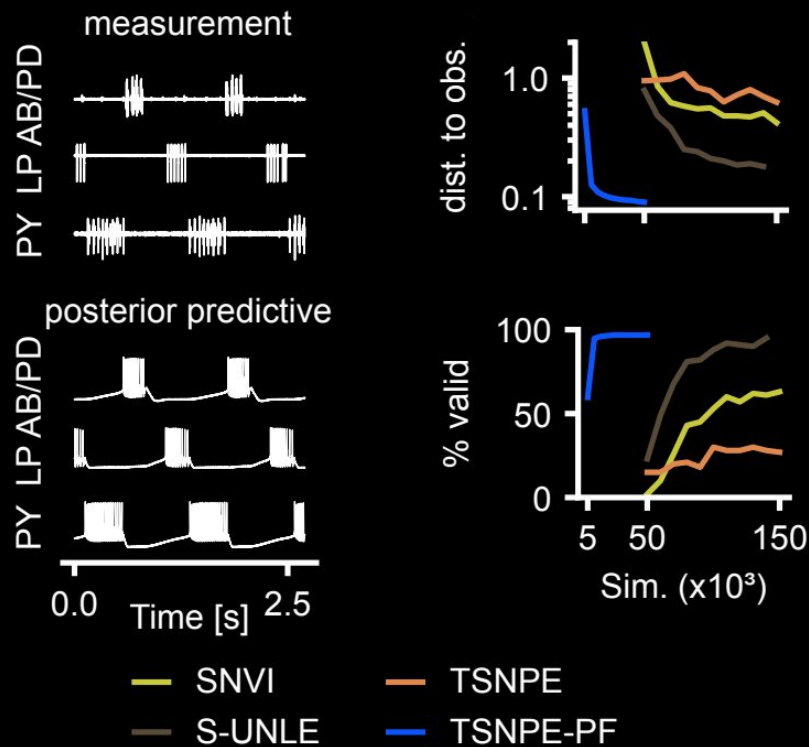Benefits sequential algorithms that **actively** propose relevant simulations.

Results: SBI Benchmark

Results: Crab pyloric network simulator

- High-dimensional, three neuron network.

- Challenge: sparsity of valid simulations.

- NPE-PF is highly simulation efficient.

Take away

- NPE-PF repurposes TabPFN for in-context SBI

- NPE-PF is easy-to-use and simulation-efficient

- Sequential variant TSNPE-PF offers even higher simulation-efficiency

- And much more (unconditional densities, misspecification, …)

Thank you

Julius Vetter

Manuel Gloeckler

Daniel Gedon

Jakob H. Macke

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

machine learning
new perspectives
for science

The Mackelab

arXiv:2504.17660