




Rapid Simulation Based Inference (RSBI): Toward accelerated inference with differential simulators

Andre Scaffidi (ascaffid@sissa.it)




SISSA
DATA SCIENCE
Machine Learning for the Natural Sciences






Two Moons, 10^4 simulations. Large prior volume: $p(\theta) = \text{Uniform}(-10, 10)$


Differentiable simulators already used substantially in the SBI literature



Augmented loss functions that utilise gradients



Score-matching term added onto the NPE loss




Optimal summaries from score compression

In every case the gradient regularises an estimator fit to prior-wide draws.

Use it to decide where to simulate?


Can we make the differentiable simulator an integral part of the SBI pipeline?



each gradient step pulls $\hat{x}$ onto the observation

Two phase RSBI in large prior volumes...

- 1 Hone in on posterior modes
- 2 Refine with local NRE



Challenges in large prior volumes...

- 1 Hone in on posterior modes



We are doing ABC - Need to cover all modes...

Challenges in large prior volumes...


- 1 Hone in on posterior modes




We are doing ABC - Need to cover all modes...

Can use variational inference? RKL?


Phase 1: Surrogate guided zoom



Phase 1: Surrogate guided zoom




RKL is mode seeking!



Phase 1: Surrogate guided zoom

$$p(\theta | x_o) = \nu(\theta) \propto e^{-E(\theta)}$$




Adjoint Schrödinger Bridge Sampler

Guan-Horng Liu^{1,*}, Jaemoo Choi^{2,*}, Yongxin Chen², Benjamin Kurt Miller¹,
Ricky T. Q. Chen^{1,*}


¹FAIR at Meta, ²Georgia Institute of Technology, *Core contributors

Instead of learning variational density $\rightarrow$ diffusion!

Phase 1: Surrogate guided zoom




Phase 2: Posterior refinement via local NRE







the ratio need only be accurate where Phase 1 placed mass

$$\nabla_\theta E_r(\theta) = -\nabla_\theta \log p(\theta | x_o) = -\nabla_\theta \log \frac{p(x_o | \theta)}{q_r(x_o)} - \nabla_\theta \log p(\theta)$$







Two Moons, 10^4 simulations. Large prior volume: $p(\theta) = \text{Uniform}(-10, 10)$

Leakage mitigation!




Flow-based amortised methods place a substantial fraction of their mass outside the prior support.

Leakage mitigation!




Flow-based amortised methods place a substantial fraction of their mass outside the prior support.



RSBI with ASBS targets constrained energy function $\Rightarrow$ Mass remains in-support while recovering the modes

For results and more come see poster!

Poster 36



Lots of benchmarks!

GW ringdown!

Early warning astronomy...

