

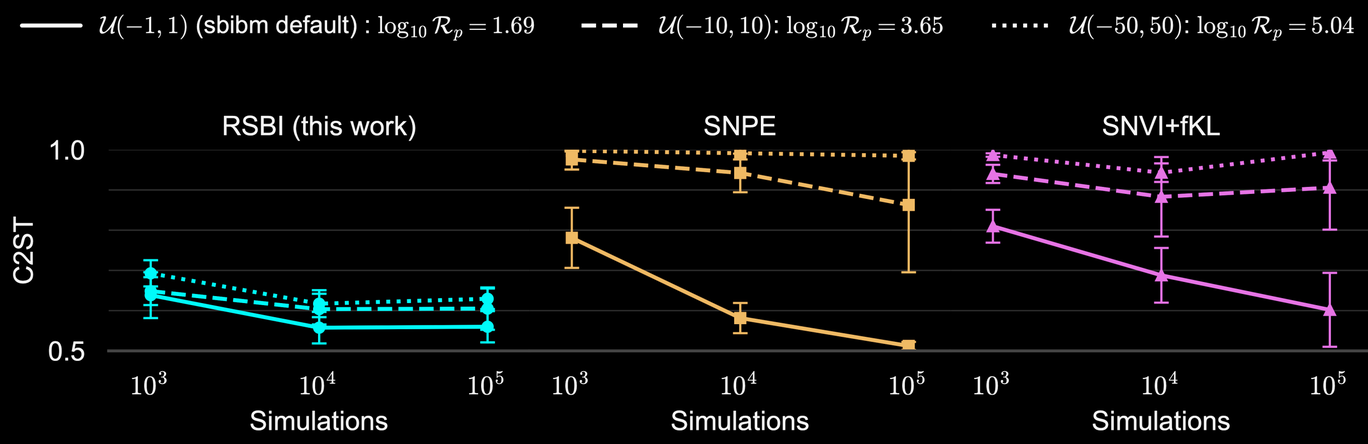
@ NerIPS 2026

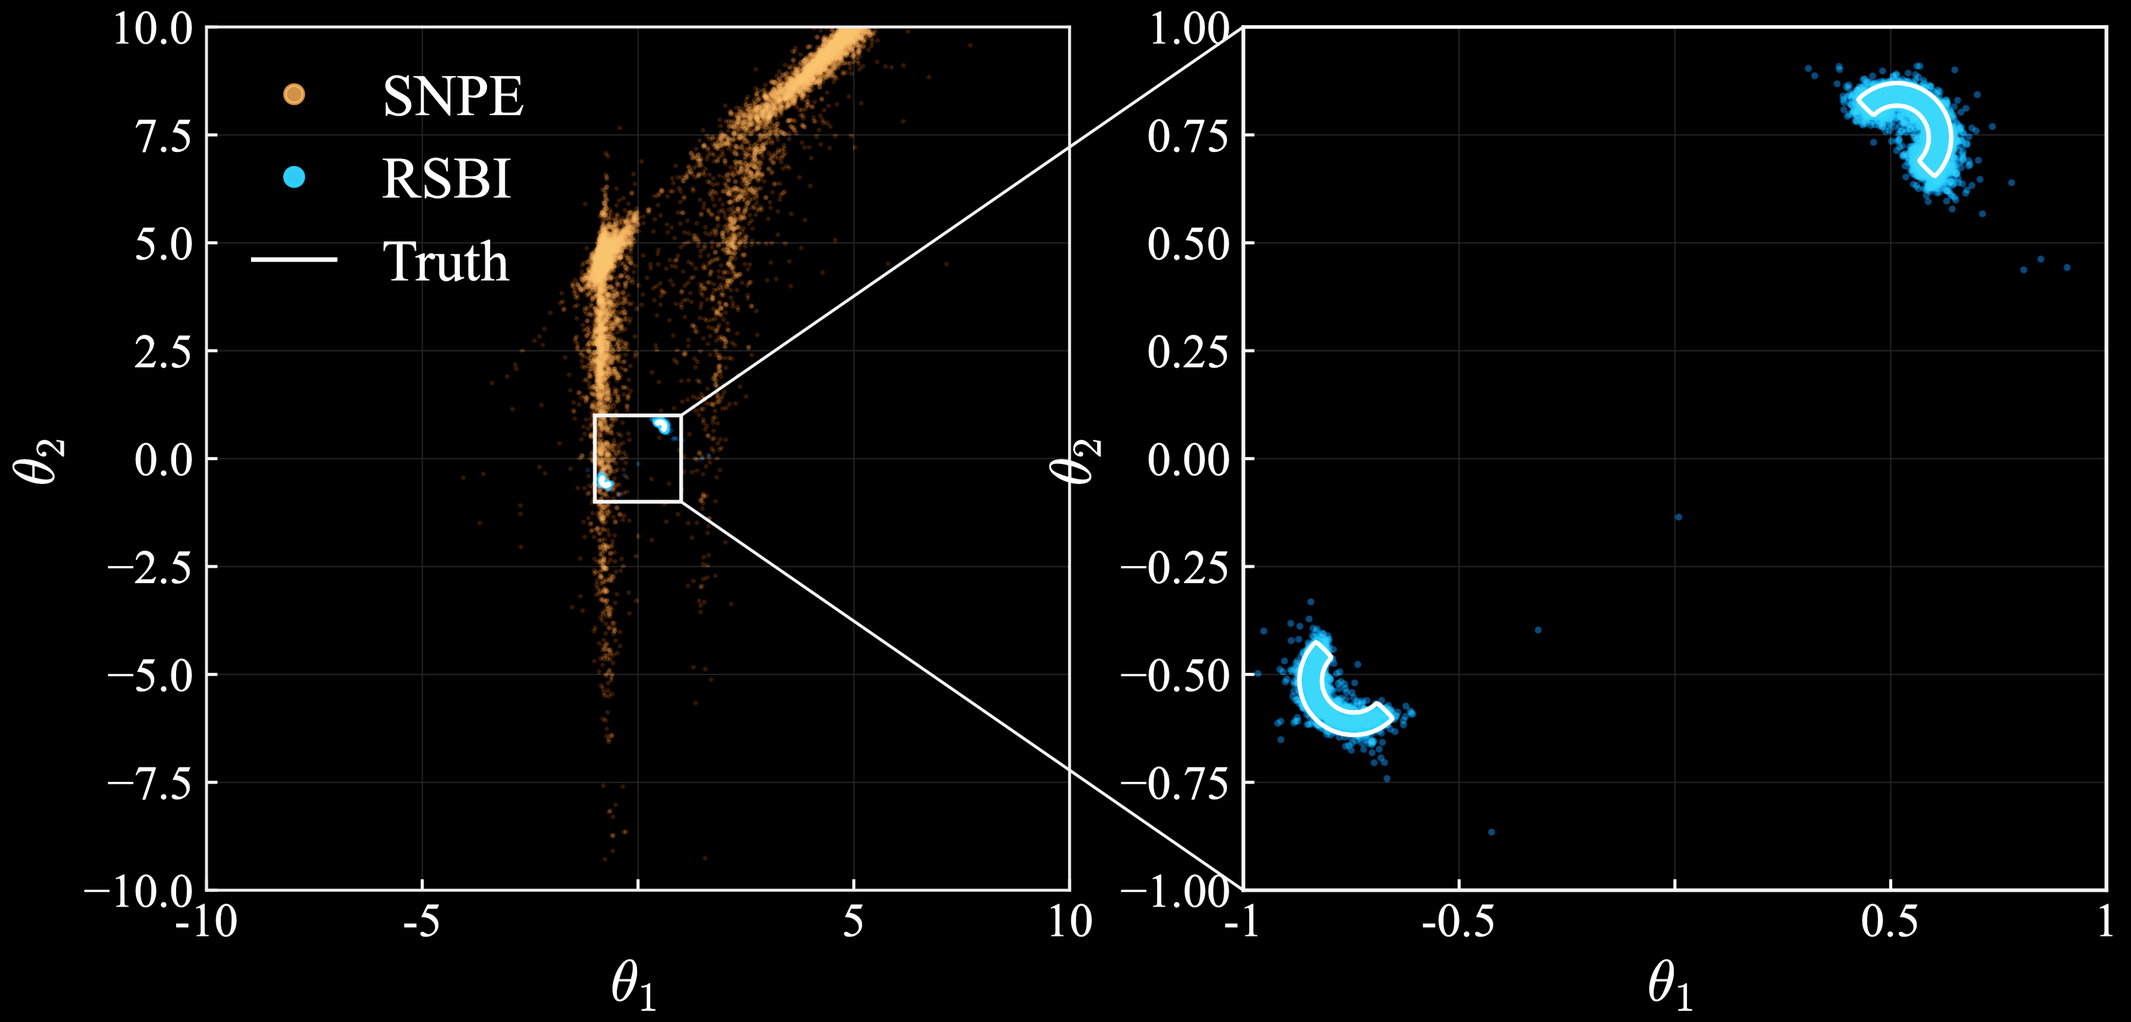
Rapid Simulation Based Inference (RSBI): Toward accelerated inference with differential simulators

Andre Scaffidi (ascaffid@sissa.it)



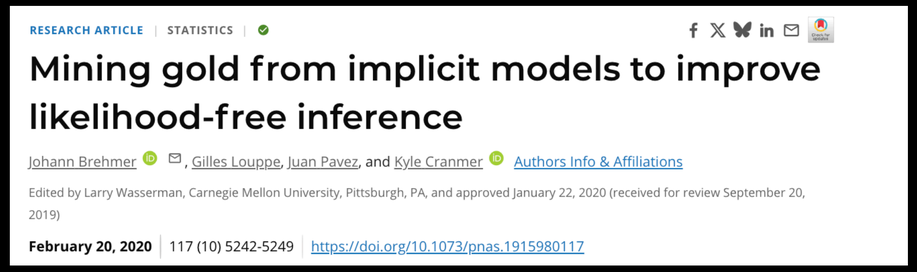






Two Moons, 10^4 simulations. Large prior volume: $p(\theta) = \text{Uniform}(-10, 10)$

Differentiable simulators already used substantially in the SBI literature



RESEARCH ARTICLE | STATISTICS |

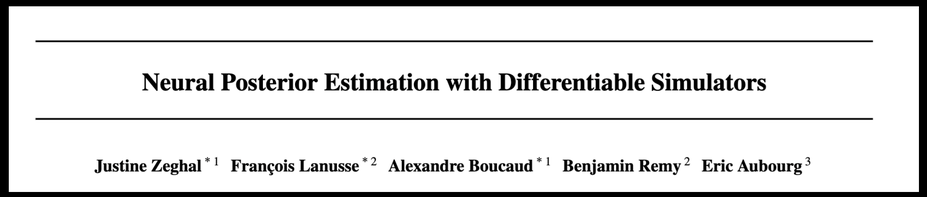
Mining gold from implicit models to improve likelihood-free inference

Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer Authors Info & Affiliations

Edited by Larry Wasserman, Carnegie Mellon University, Pittsburgh, PA, and approved January 22, 2020 (received for review September 20, 2019)

February 20, 2020 | 117 (10) 5242-5249 | https://doi.org/10.1073/pnas.1915980117

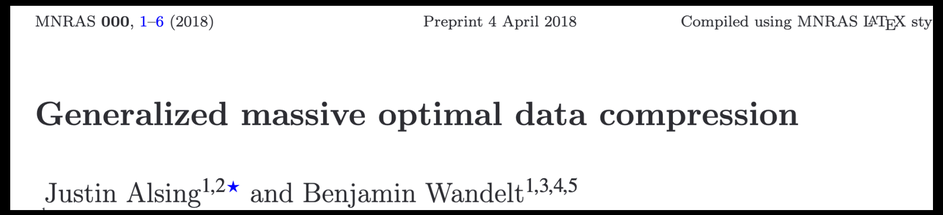
Augmented loss functions that utilise gradients



Neural Posterior Estimation with Differentiable Simulators

Justine Zeghal [*1] François Lanusse [*2] Alexandre Boucaud [*1] Benjamin Remy [2] Eric Aubourg [3]

Score-matching term added onto the NPE loss



MNRAS 000, 1–6 (2018) Preprint 4 April 2018 Compiled using MNRAS LATEX sty

Generalized massive optimal data compression

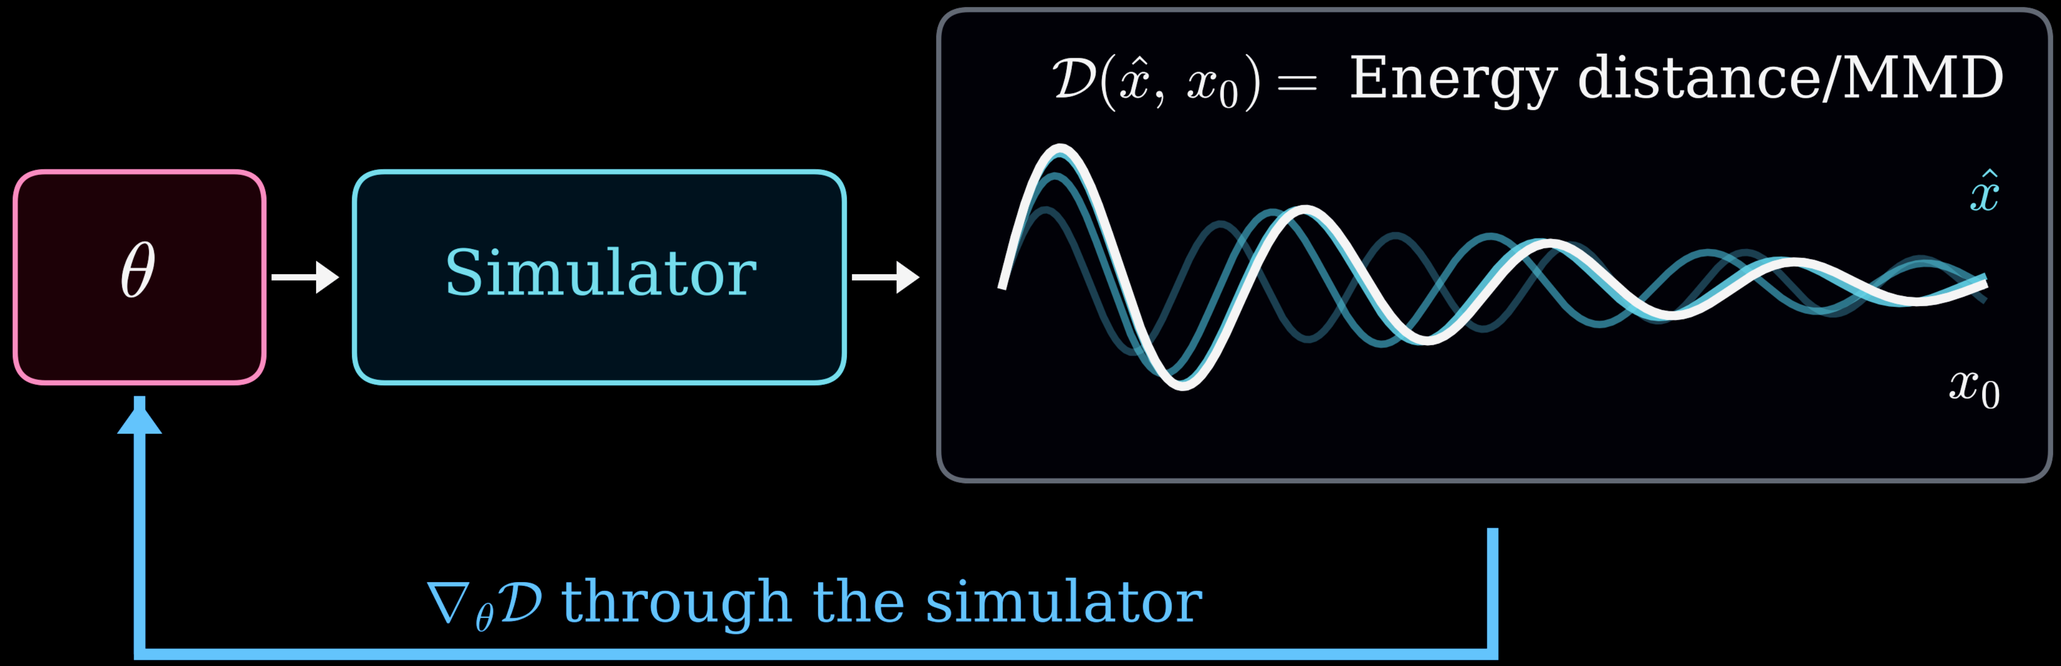
Justin Alsing[1,2★] and Benjamin Wandelt[1,3,4,5]

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?

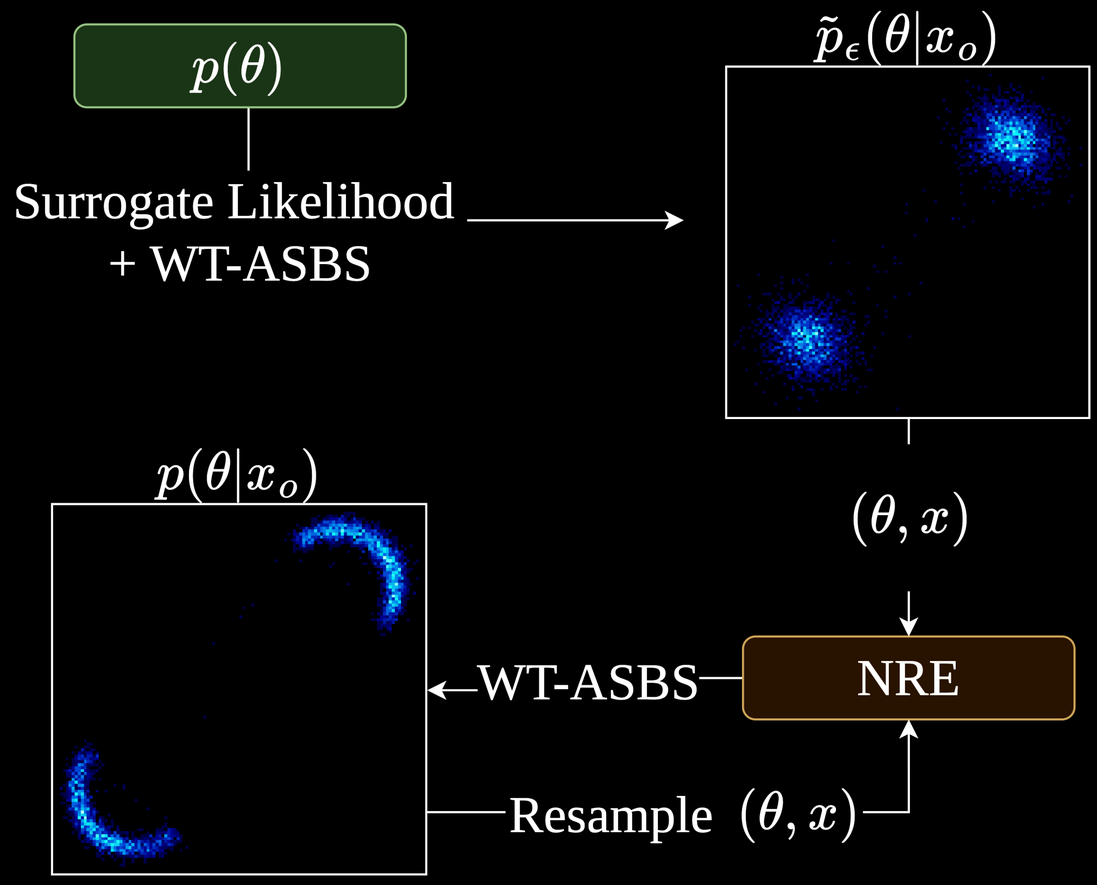


each gradient step pulls $\hat{x}$ onto the observation

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...

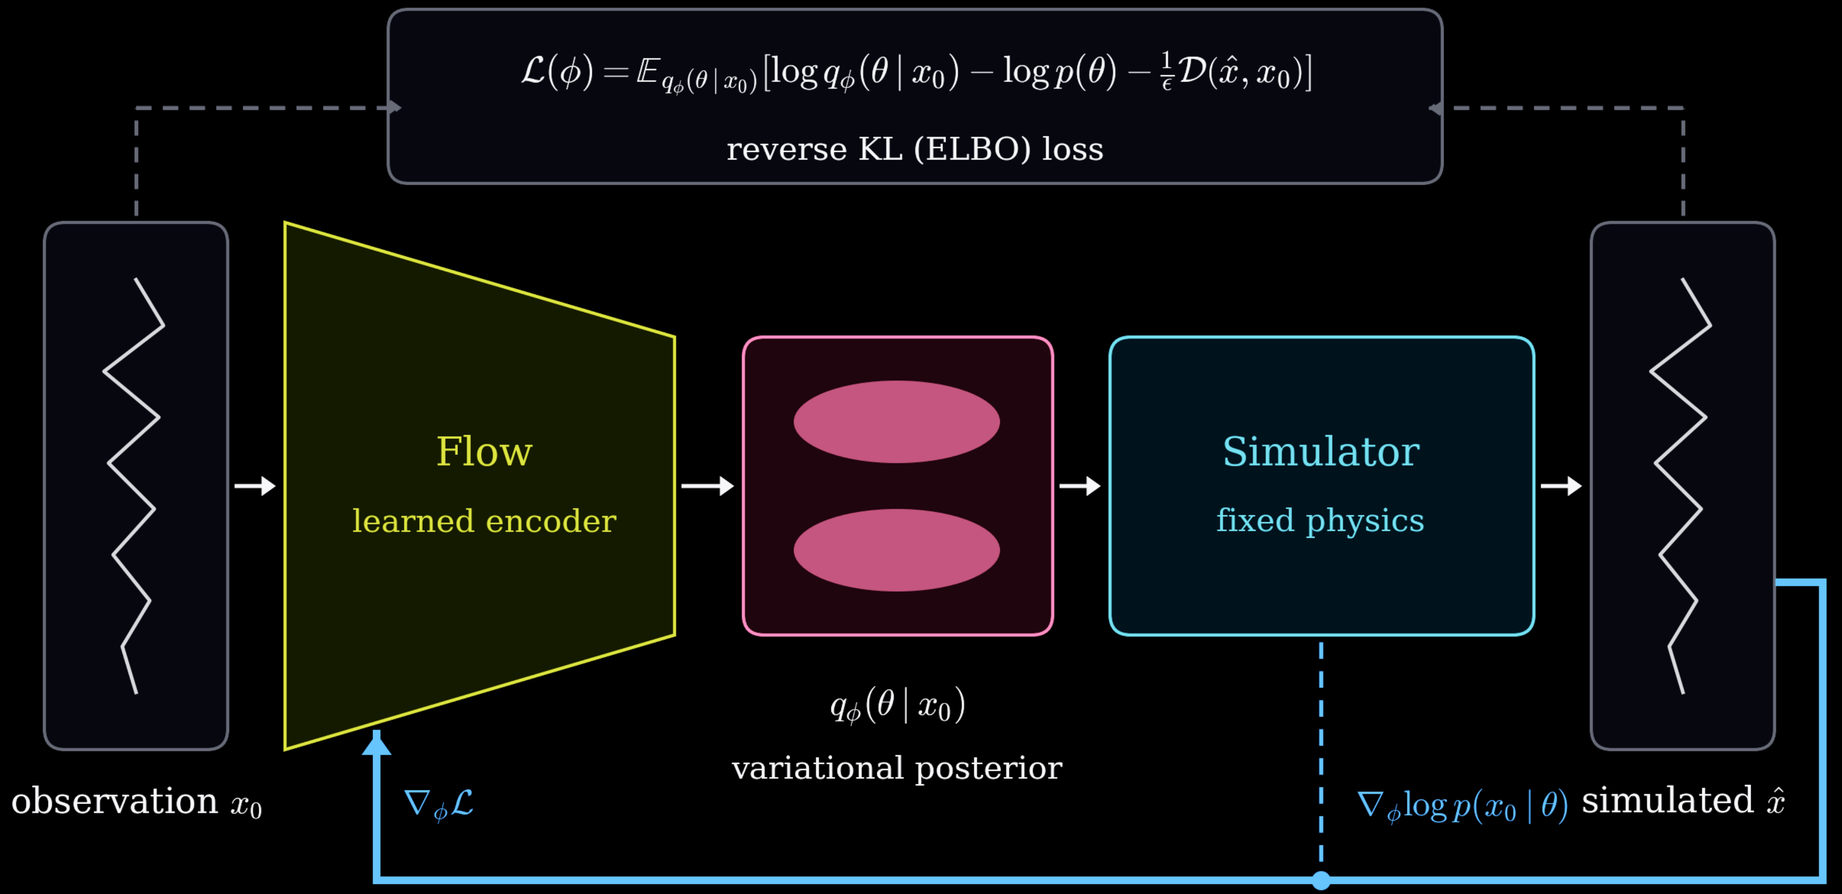
Challenges in large prior volumes...

1. Hone in on posterior modes

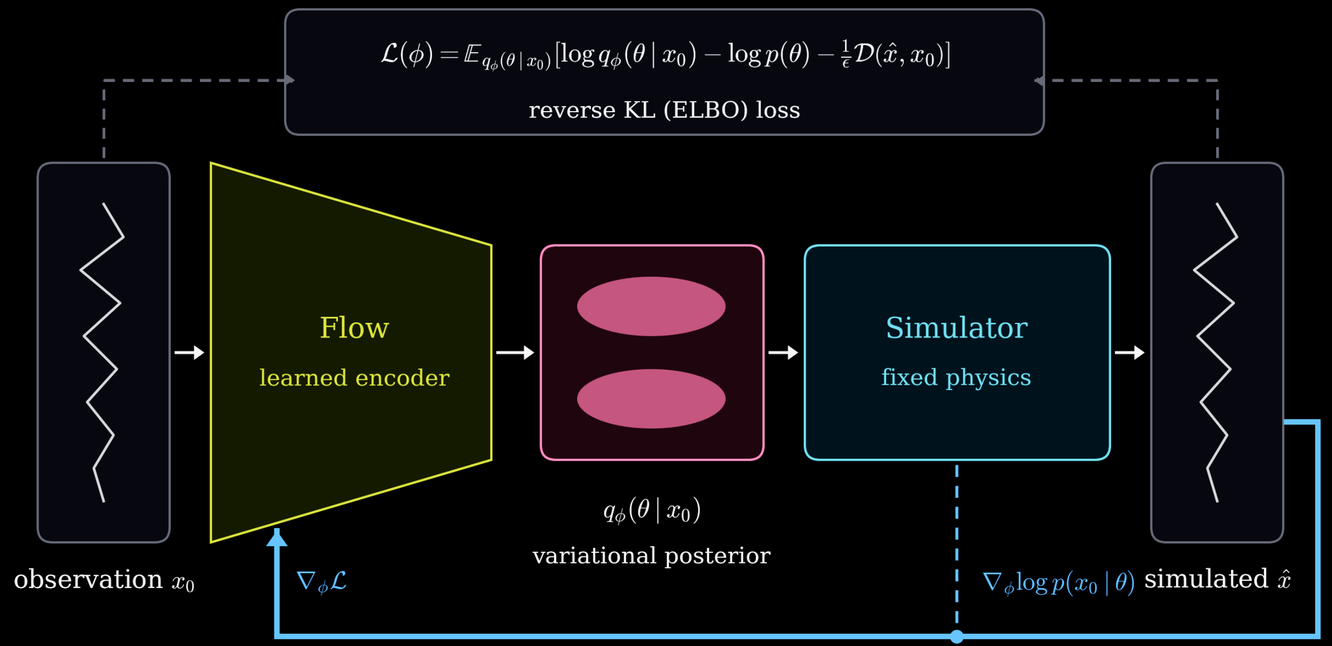
We are doing ABC - Need to cover all modes...

Can use variational inference? RKL?

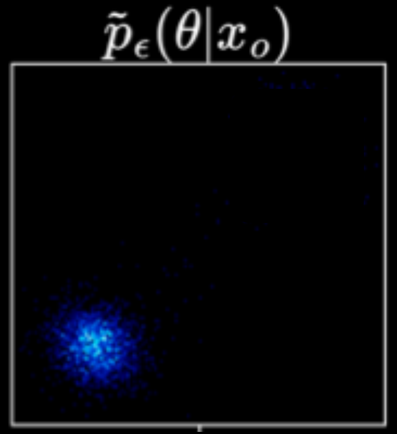
Phase 1: Surrogate guided zoom



Phase 1: Surrogate guided zoom

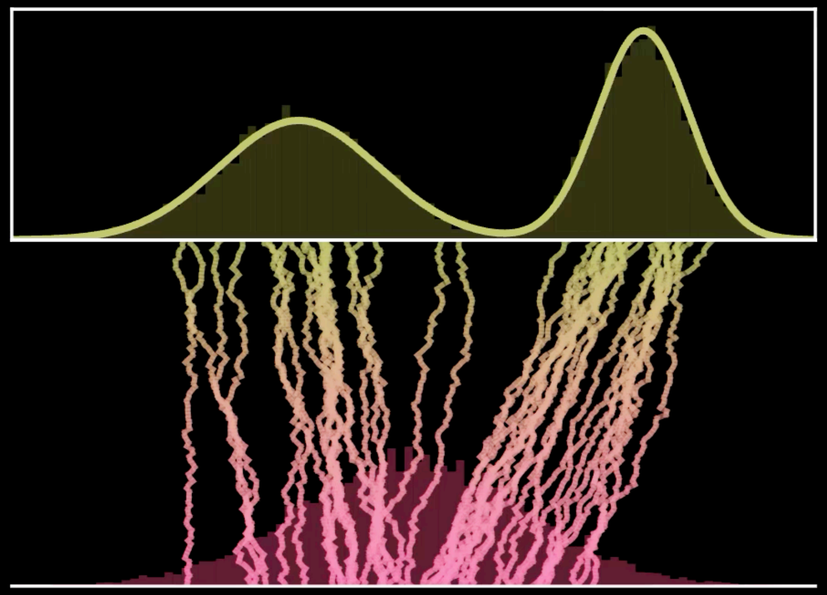


RKL is mode seeking!



Phase 1: Surrogate guided zoom

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



$p(\theta)$

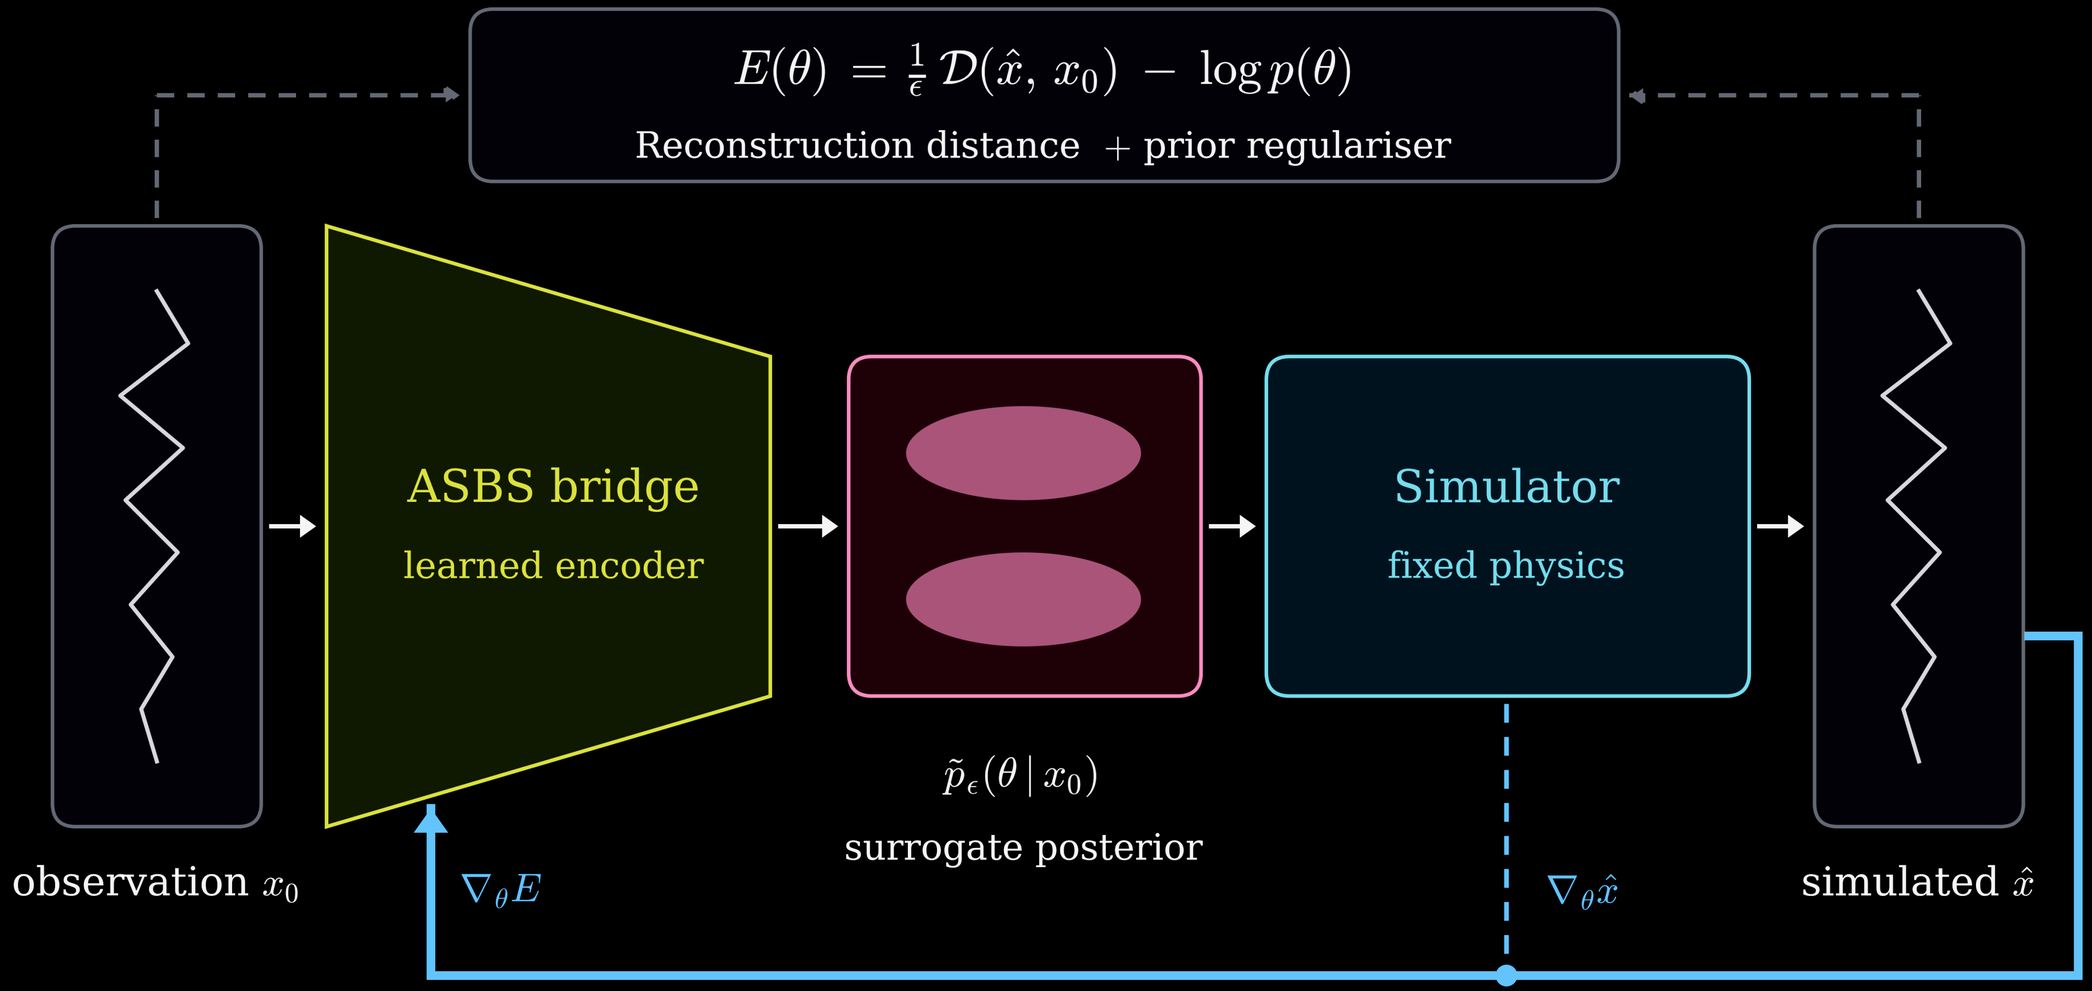
Adjoint Schrödinger Bridge Sampler

Guan-Horng Liu[1,*], **Jaemoo Choi**[2,*], **Yongxin Chen**[2], **Benjamin Kurt Miller**[1], **Ricky T. Q. Chen**[1,*]

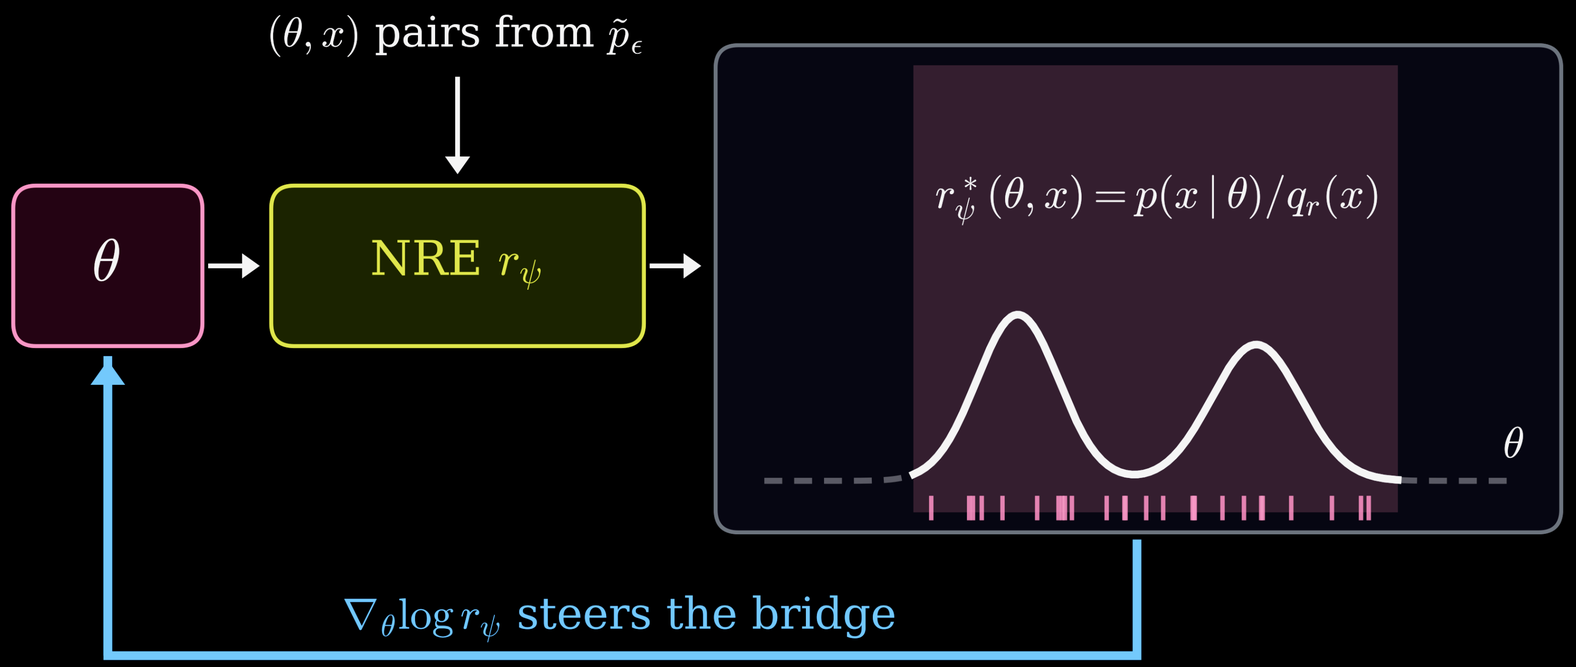
[1]FAIR at Meta, [2]Georgia Institute of Technology, *Core contributors

Instead of learning variational density → diffusion!

Phase 1: Surrogate guided zoom

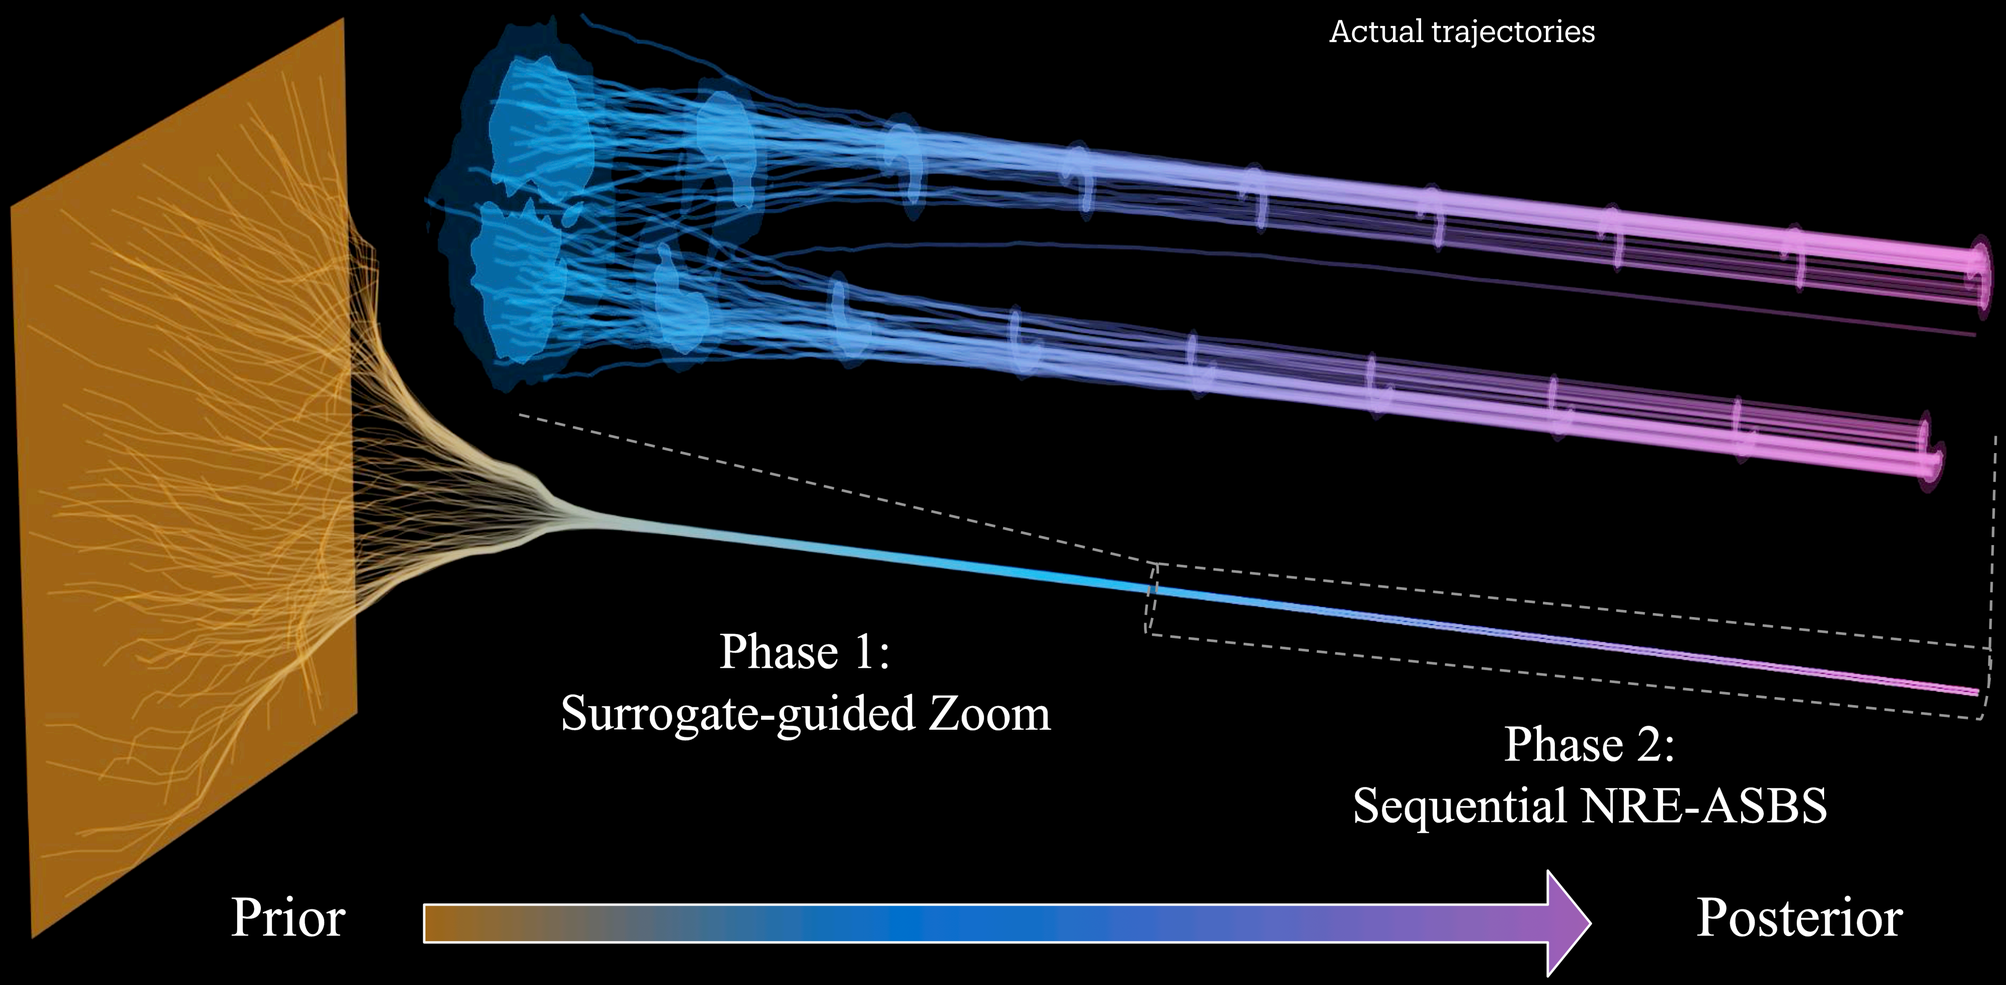


Phase 2: Posterior refinement via local NRE



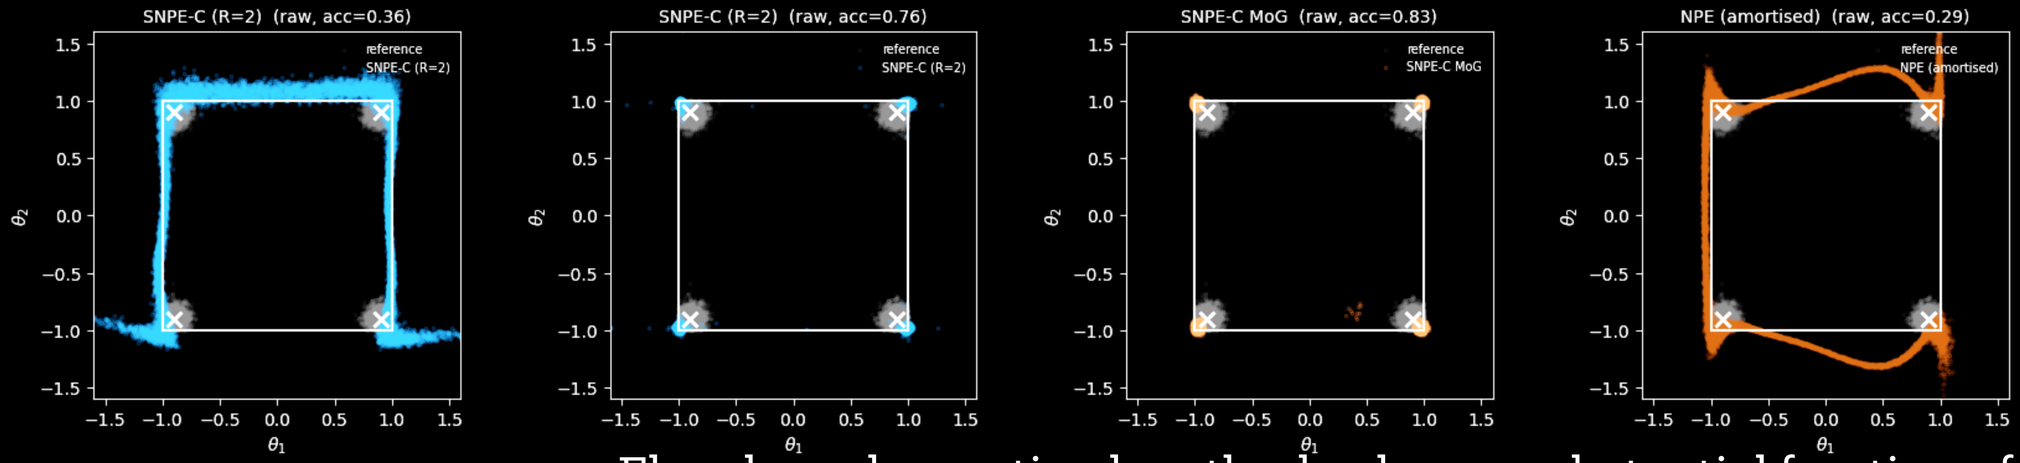
the ratio need only be accurate where Phase 1 placed mass

$$\nabla_\theta E_r(\theta) = -\nabla_\theta \log p(\theta \mid x_o) = -\nabla_\theta \log \frac{p(x_0 \mid \theta)}{q_r(x_0)} - \nabla_\theta \log p(\theta)$$



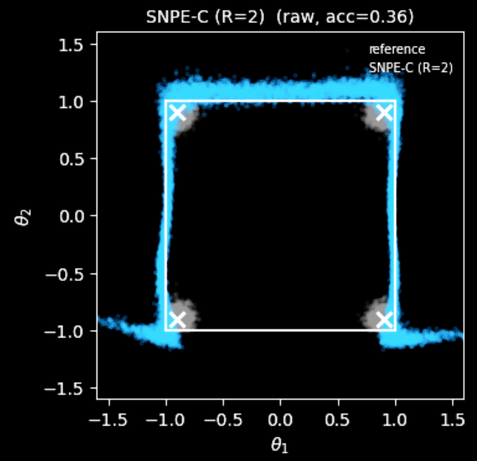
Two Moons, 10^4 simulations. Large prior volume: $p(\theta) = \text{Uniform}(-10, 10)$

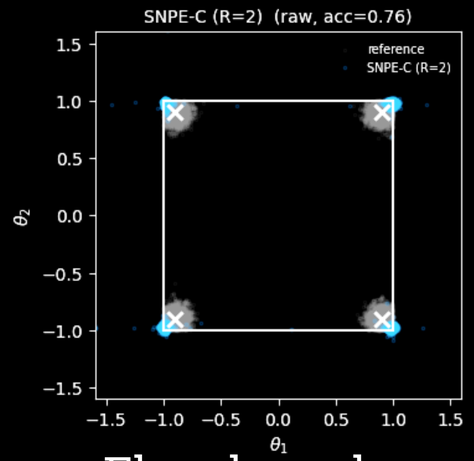
Leakage mitigation!

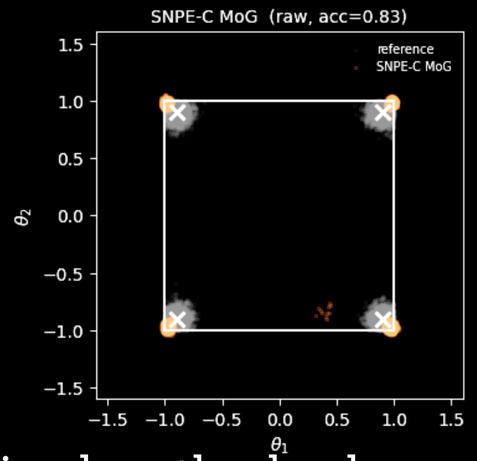


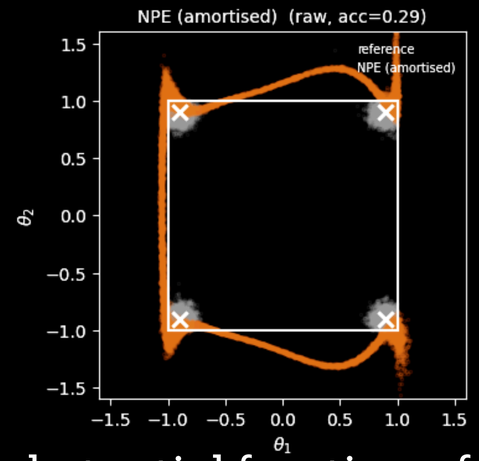
Flow-based amortised methods place a substantial fraction of their mass outside the prior support.

Leakage mitigation!

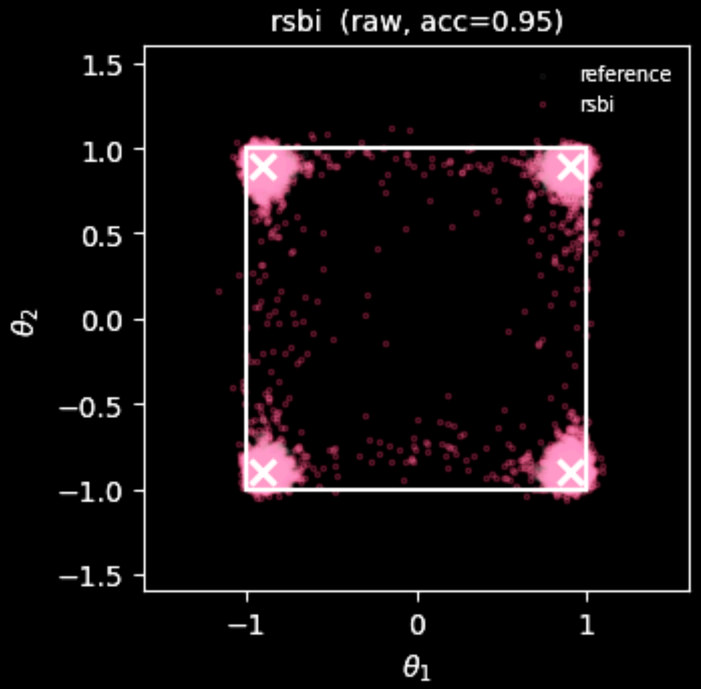




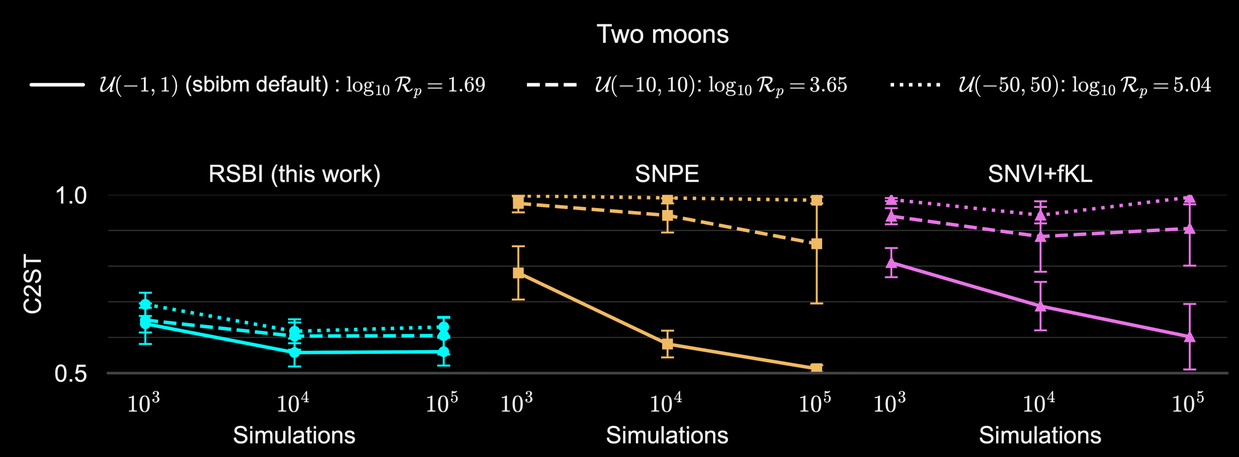


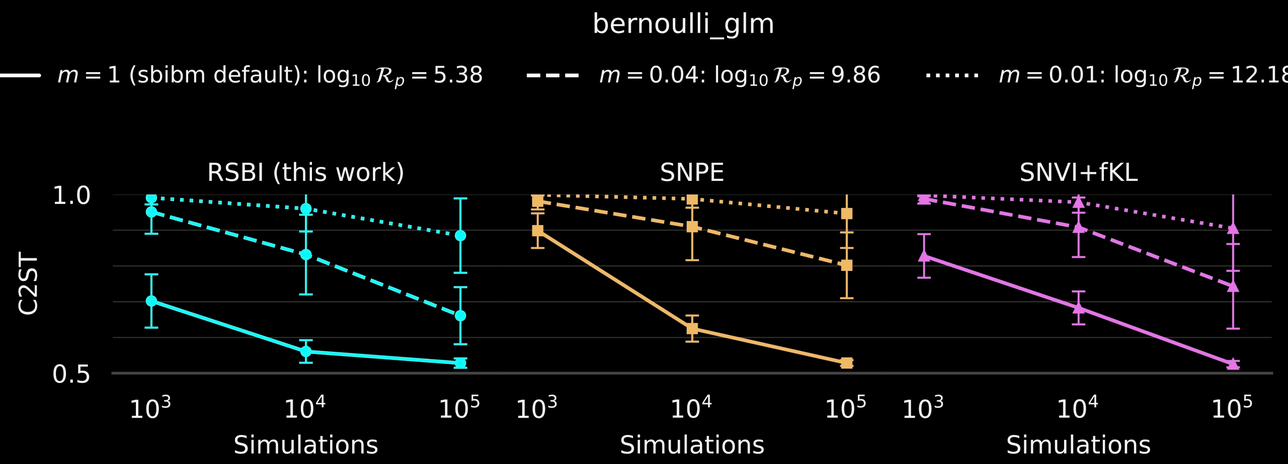


Flow-based amortised methods place a substantial fraction of their mass outside the prior support.

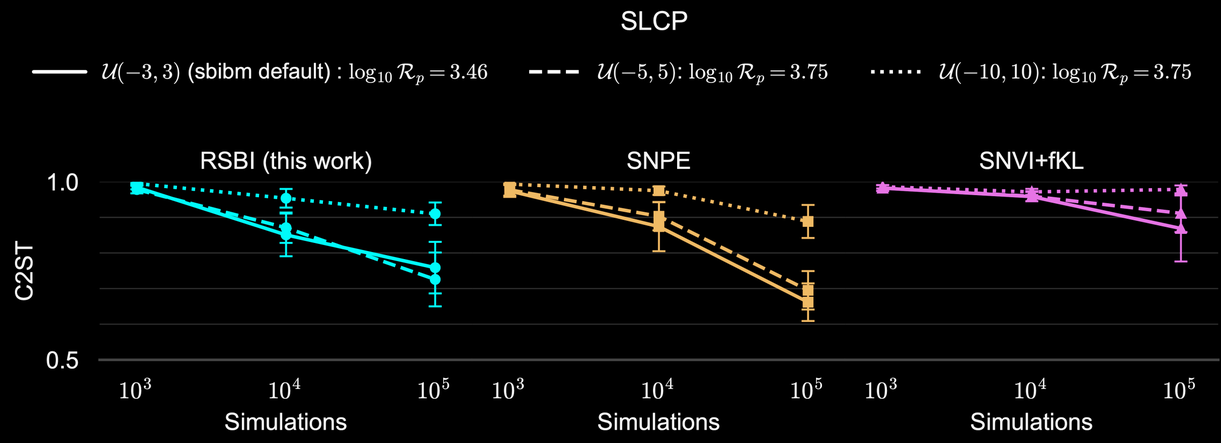


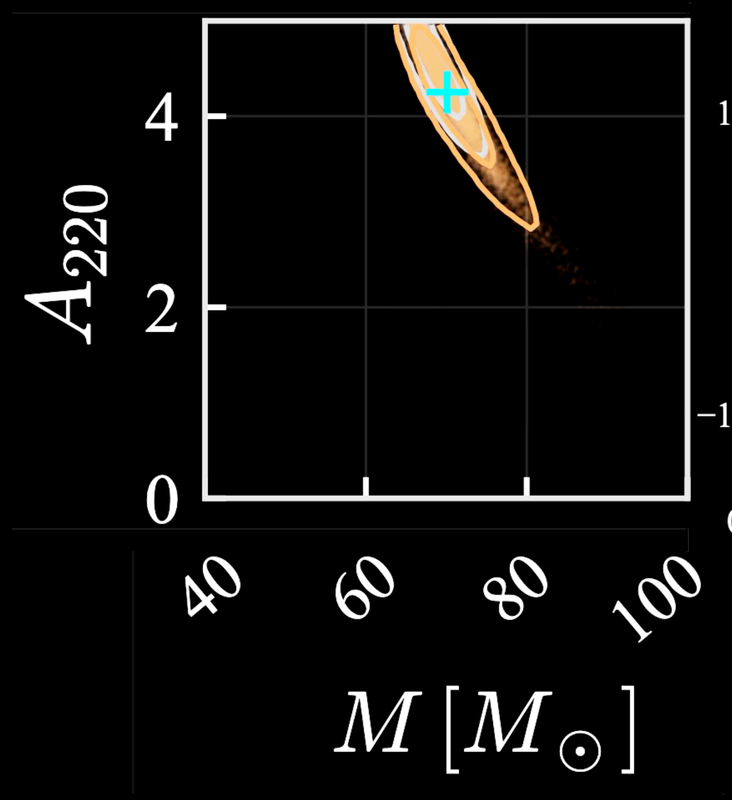
RSBI with ASBS targets constrained energy function ⇒ Mass remains in-support while recovering the modes

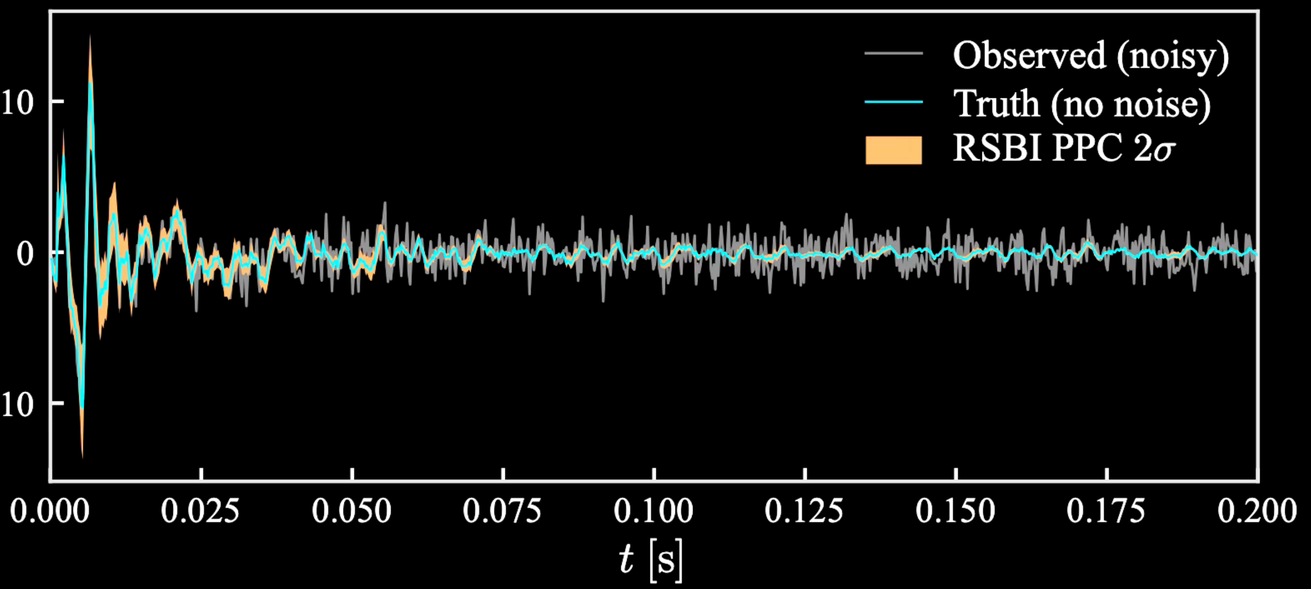




Lots of benchmarks!







Method	default ($\log_{10}\mathcal{R}_p \approx 6.4$)	wide ($\log_{10}\mathcal{R}_p \approx 6.7$)
SNVI-fKL	0.958 (0.032)	0.977 (0.027)
SNPE-C	0.935 (0.071)	0.982 (0.017)
RSBI (ours)	**0.834** (**0.103**)	**0.894** (**0.054**)

GW ringdown!

Many potential follow ups...online early warning astro?