

Reliable Simulation-Based Inference with Balanced Neural Ratio Estimation

PHYSTAT

February 8, 2023

Gilles Louppe
g.louppe@uliege.be



Kyle Cranmer




Johann
Brehmer



Joeri
Hermans



Antoine
Wehenkel



Norman Marlier



Siddharth
Mishra-
Sharma



Christoph
Weniger



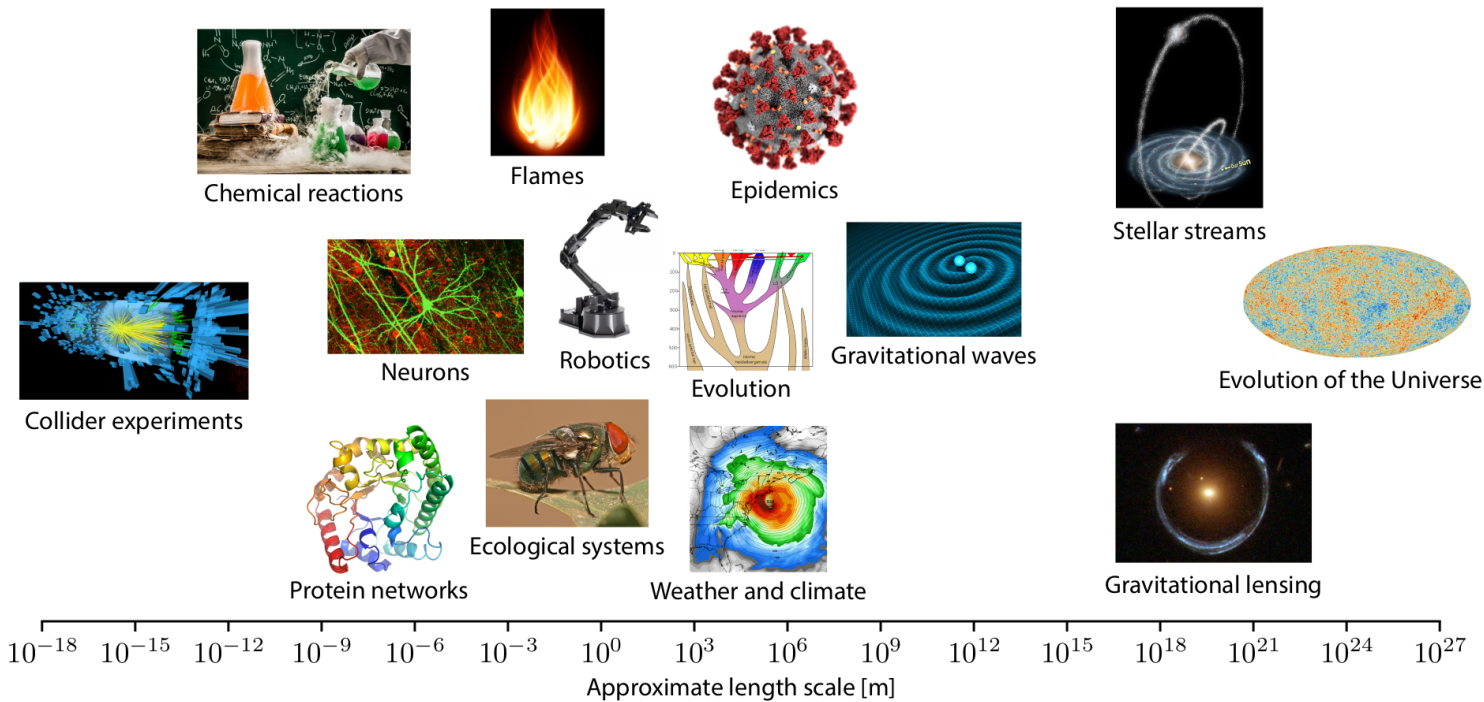
Arnaud
Delaunoy




Malavika
Vasist



Francois Rozet







$$v_x = v \cos(\alpha), \quad v_y = v \sin(\alpha),$$


$$\frac{dx}{dt} = v_x, \quad \frac{dy}{dt} = v_y, \quad \frac{dv_y}{dt} = -G.$$



```
def simulate(v, alpha, dt=0.001):  
    v_x = v * np.cos(alpha) # x velocity m/s  
    v_y = v * np.sin(alpha) # y velocity m/s  
    y = 1.1 + 0.3 * random.normal()  
    x = 0.0  
  
    while y > 0: # simulate until ball hits floor  
        v_y += dt * -G # acceleration due to gravity  
        x += dt * v_x  
        y += dt * v_y  
  
    return x + 0.25 * random.normal()
```






The computer simulator defines the likelihood function $p(x|\theta)$ implicitly.



What parameter values θ are the most plausible?


$$p(\theta | x_{\text{obs}}) = \frac{p(x_{\text{obs}} | \theta) p(\theta)}{p(x_{\text{obs}})}$$






Neural ratio estimation (NRE)

The likelihood-to-evidence $r(x|\theta) = \frac{p(x|\theta)}{p(x)} = \frac{p(x,\theta)}{p(x)p(\theta)}$ ratio can be learned, even if neither the likelihood nor the evidence can be evaluated:



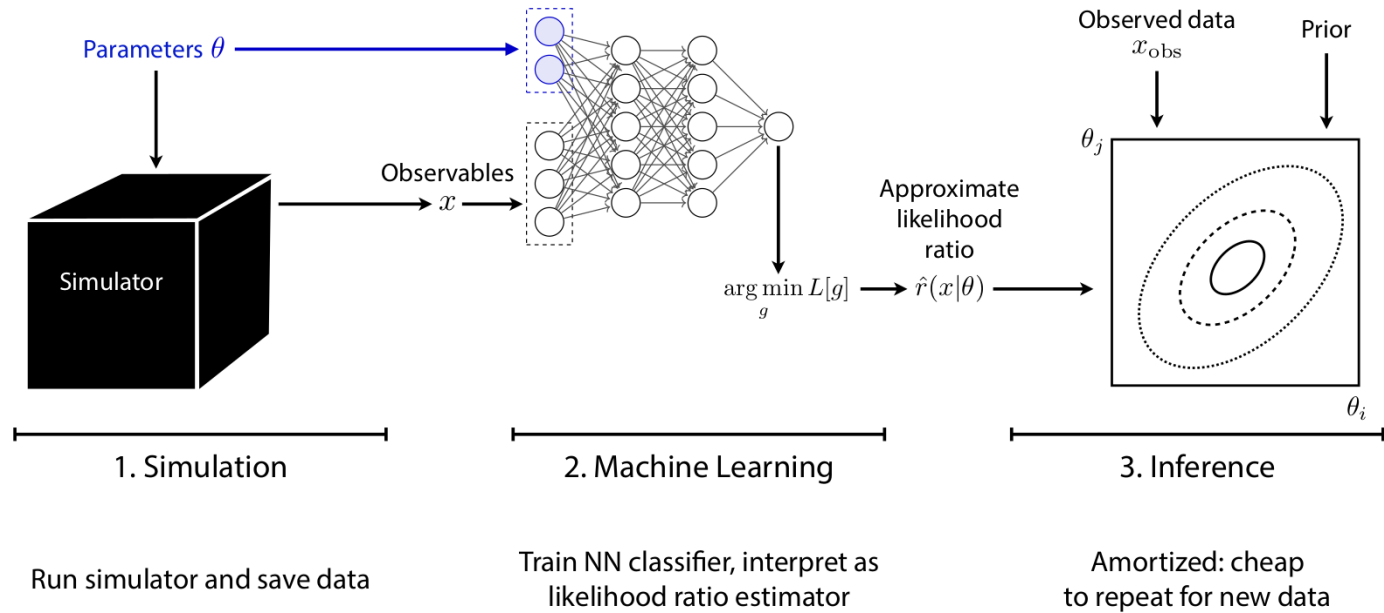



The solution d found after training approximates the optimal classifier

$$d(x, \theta) \approx d^*(x, \theta) = \frac{p(x, \theta)}{p(x, \theta) + p(x)p(\theta)}.$$

Therefore,

$$r(x|\theta) = \frac{p(x|\theta)}{p(x)} = \frac{p(x, \theta)}{p(x)p(\theta)} \approx \frac{d(x, \theta)}{1 - d(x, \theta)} = \hat{r}(x|\theta).$$



Constraining dark matter with stellar streams (Bayesian)



Palomar 5 (Pal5) stream
Pal5 was discovered in 2001 as the first thin stream formed from a globular cluster. Its current orbit takes it far over the galactic center.


Globular clusters
These hives typically hold 100,000 stars or fewer and give rise to long, thin streams.


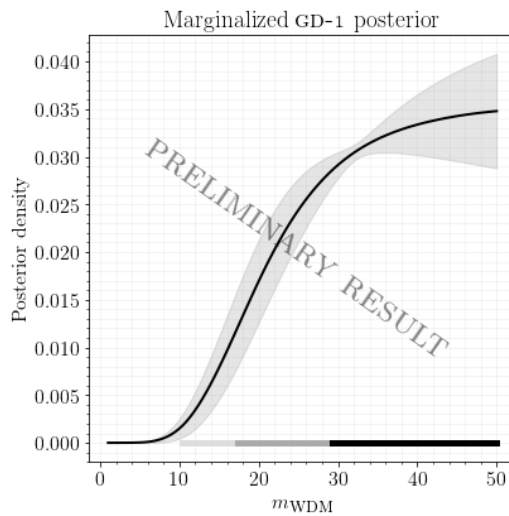
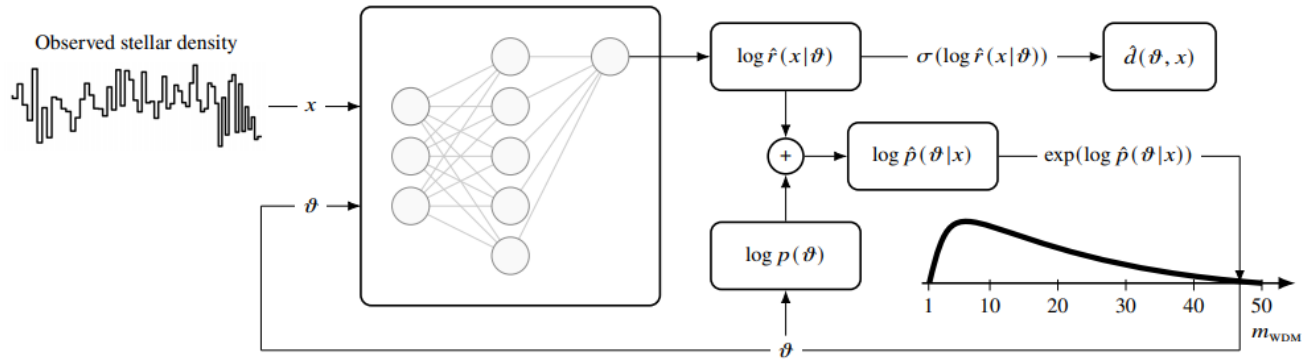
← Gap


Sun

Milky Way

GD1 stream
Discovered in 2006, GD1 is the longest known thin stream, stretching across more than half the northern sky. It contains a gap that could be the scar of a dark matter collision 500 million years ago.












Preliminary results for GD-1 suggest a **preference for CDM over WDM.**



Neural Posterior Estimation (NPE)



$$\min_{q_\phi} \mathbb{E}_{p(x)} [\text{KL}(p(\theta|x) || q_\phi(\theta|x))]$$

Exoplanet atmosphere characterization (Bayesian)






Computational faithfulness

$$\hat{p}(\theta|x) = \text{sbi}(p(x|\theta), p(\theta), x)$$

We must make sure our approximate simulation-based inference algorithms can (at least) actually realize faithful inferences on the (expected) observations.




How do we know this is good enough?



Mode convergence:

The maximum a posteriori estimate converges towards the nominal value θ^* for an increasing number of independent and identically distributed observables $x_i \sim p(x|\theta^*)$:

$$\begin{aligned} & \lim_{N \rightarrow \infty} \arg \max_{\theta} p(\theta | \{x_i\}_{i=1}^N) \\ &= \lim_{N \rightarrow \infty} \arg \max_{\theta} p(\theta) \prod_{x_i} r(x_i | \theta) = \theta^* \end{aligned}$$







A common observation at the root of several other diagnostics is to check for the **self-consistency** of the Bayesian joint distribution,


$$p(\theta) = \int p(\theta')p(x|\theta')p(\theta|x)d\theta' dx.$$

Coverage diagnostic:

- For $x, \theta \sim p(x, \theta)$, compute the $1 - \alpha$ credible interval based on $\hat{p}(\theta|x)$.
- If the fraction of samples for which θ is contained within the interval is larger than the nominal coverage probability $1 - \alpha$, then the approximate posterior $\hat{p}(\theta|x)$ has coverage.







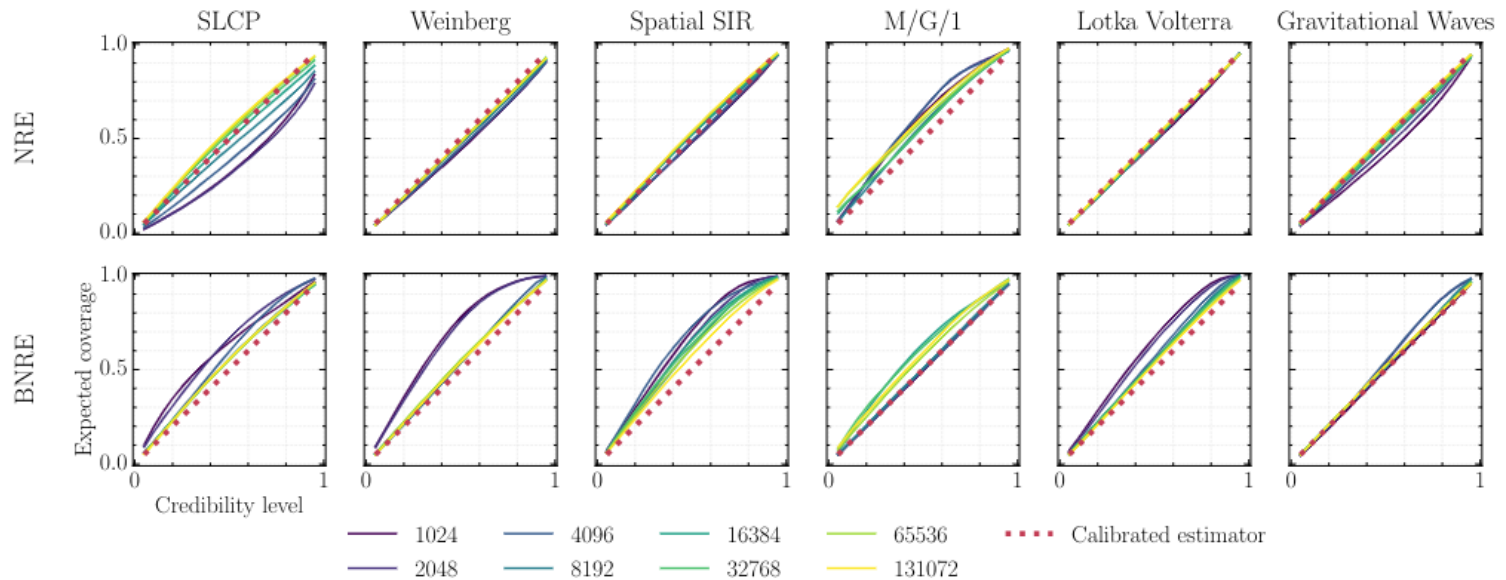
What if diagnostics fail?


Balanced NRE



Enforce neural ratio estimation to be **conservative** by using binary classifiers \hat{d} that are balanced, i.e. such that

$$\mathbb{E}_{p(\theta, x)} \left[\hat{d}(\theta, x) \right] = \mathbb{E}_{p(\theta)p(x)} \left[1 - \hat{d}(\theta, x) \right].$$





Summary

Simulation-based inference is a major evolution in the statistical capabilities for science, enabled by advances in machine learning.

Need to reliably and efficiently evaluate the quality of the posterior approximations.

The end.