Learned harmonic mean estimation of the marginal likelihood with normalizing flows

Alicja Polanska, Matthew A. Price,
Alessio Spurio Mancini and Jason D. McEwen
University College London. Correspondence: alicja.polanska.22@ucl.ac.uk
Software: https://github.com/astro-informatics/harmonic.git

Overview

Problem To discover the science underlying physical phenomena we need to find out which model is the most plausible in the light of data. Bayesian evidence is a crucial quantity in this task, but it is challenging to compute in practice. An existing method, the harmonic mean estimator fails catastrophically.

Solution Present *learned harmonic mean estimator* to revise original estimator. New target distribution is learned using normalizing flows that are concentrated so that target is contained within the posterior. This addresses problems with the original method, allowing for accurate and robust Bayesian evidence estimates.

Code Learned harmonic mean estimator is implemented in the `harmonic` software package (https://github.com/astro-informatics/harmonic.git).

Learned harmonic mean

Bayesian model selection:

▶ Posterior distribution of parameters $\theta \in \Theta$ given data $\boldsymbol{y}$

$$P(\theta|\boldsymbol{y}, M) = \frac{\overbrace{P(\boldsymbol{y}|\theta, M)}^{\text{Likelihood}}\,\overbrace{P(\theta|M)}^{\text{Prior}}}{\underbrace{P(\boldsymbol{y}|M)}_{\text{Bayesian Evidence}}} = \frac{\mathcal{L}(\theta)\pi(\theta)}{z}$$

where z is the Bayesian evidence

$$z = \mathrm{P}(\boldsymbol{y}|M) = \int \mathrm{d}\theta\, \mathcal{L}(\theta)\pi(\theta).$$

▶ Estimate evidence for model comparison.

Original harmonic mean [1]:

▶ Harmonic mean estimator:

$$\rho = \mathbb{E}_{P(\theta|\boldsymbol{y})}\left[\frac{1}{\mathcal{L}(\theta)}\right] = \frac{1}{z}.$$

▶ Agnostic to sampling method.

▶ Can fail catastrophically due to large variance.

Re-targeted harmonic mean [2]:

▶ Introduce alternative target distribution $\varphi(\theta)$ with thinner tails than posterior:

$$\rho = \mathbb{E}_{P(\theta|\boldsymbol{y})}\left[\frac{\varphi(\theta)}{\mathcal{L}(\theta)\pi(\theta)}\right].$$

▶ Interpret as importance sampling, with sampling density given by posterior and target φ.

Learned target distribution [3]:

▶ Optimal target is the posterior but requires z to be known, which is precisely quantity estimating.

▶ Train a machine learning model of the target distrubtion ψ^{ML} on samples of the posterior:

$$\psi^{\mathrm{ML}} \approx \psi^{\mathrm{optimal}}(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{z}.$$

▶ Don't require accurate approximation but do require learned model to be contained within posterior → bespoke optimisation problem to learn target while minimising variance of the estimator.

▶ Extended to simulation-based inference (SBI), when an explicit likelihood is unavailable of infeasible [4].

References

[1] Newton M. A. and Raftery A. E. 1994, Approximate Bayesian inference with the weighted likelihood bootstrap
[2] Gelfand A. E. and Dey D. K. 1994, Bayesian model choice: asymptotics and exact calculations
[3] McEwen et al. 2021, Machine learning assisted marginal likelihood estimation: learnt harmonic mean estimator, arXiv:2111.12720
[4] Spurio Mancini et al. 2022, Bayesian model comparison for simulation-based inference, arXiv:2207.04037
[5] Polanska et al. 2023, Learned harmonic mean estimation of the marginal likelihood with normalizing flows
[6] Dinh et al. 2016, Density estimation using real NVP
[7] Friel and Wyse 2012, Estimating the evidence – a review

Concentrating the probability density with normalizing flows

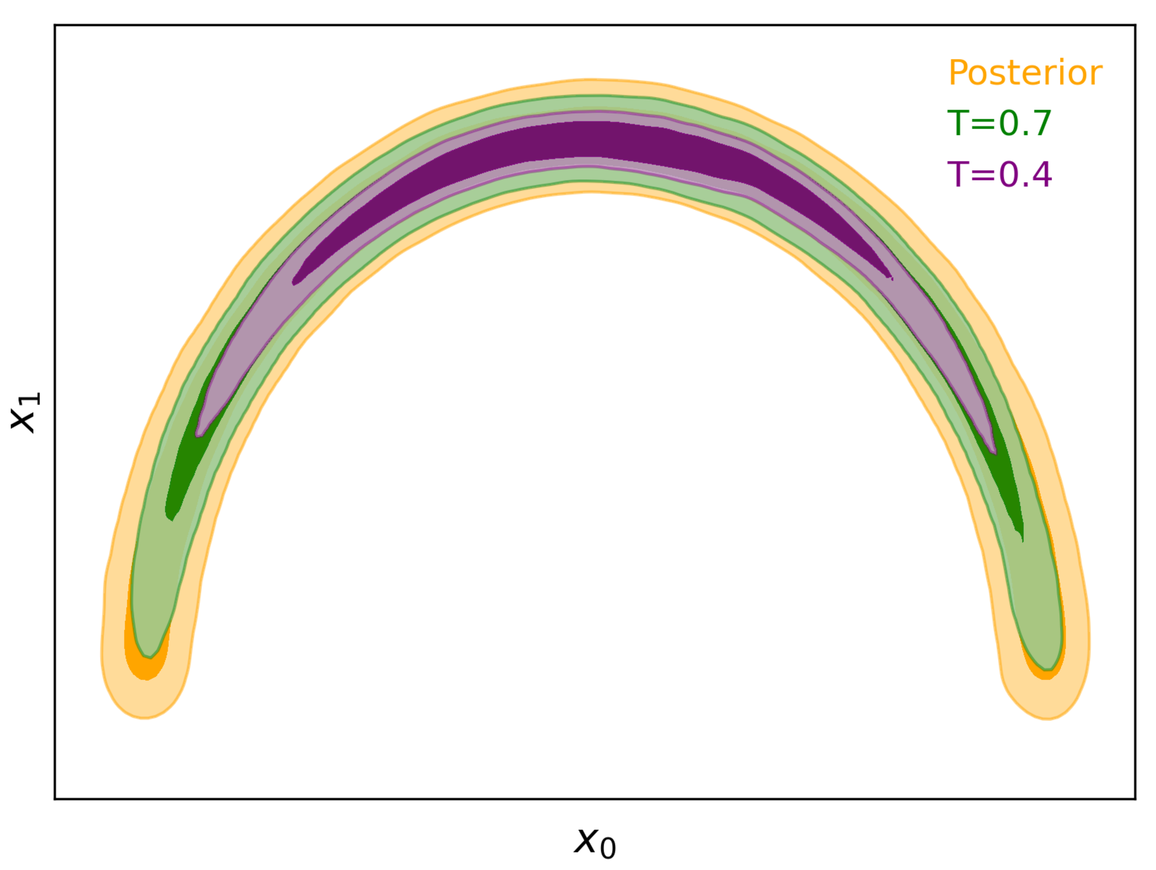

▶ **This work [5]**: Use normalizing flows in learned harmonic mean estimator for robustness and potential for scalability.

▶ Train real NVP flow [6] on samples from posterior distribution → Normalized approximation of posterior.

▶ Introduce **flow temperature parameter** T. Scale base distribution's variance to concentrate probability density, ensuring target is contained within the posterior.

▶ Eliminate the need for bespoke training approach in [3]. Robust algorithm with no need for fine-tuning.

Rosenbrock

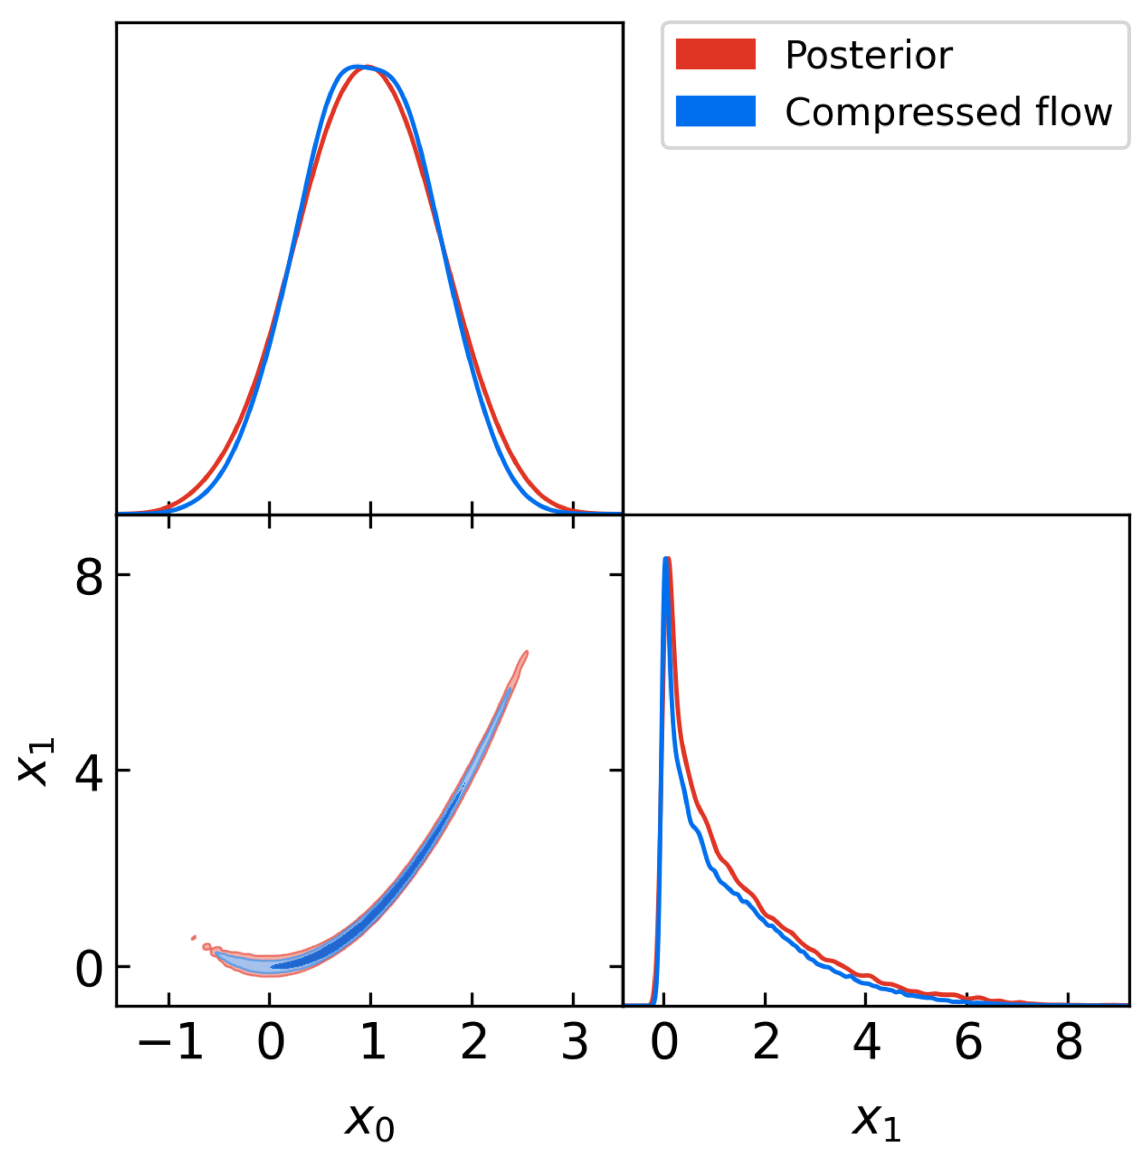

▶ Rosenbrock function exhibits a narrow curving degeneracy. Challenging to explore the posterior sufficiently to evaluate the marginal likelihood.

▶ Learned harmonic mean estimator *highly accurate*.

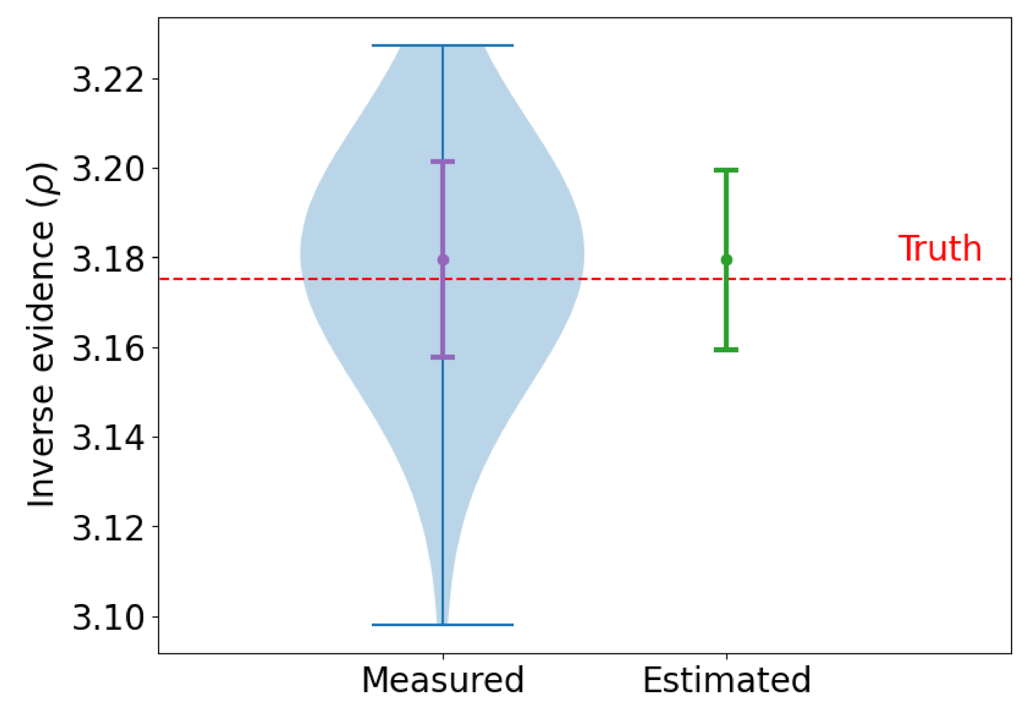

Pathological prior insensitivity: Normal-Gamma model

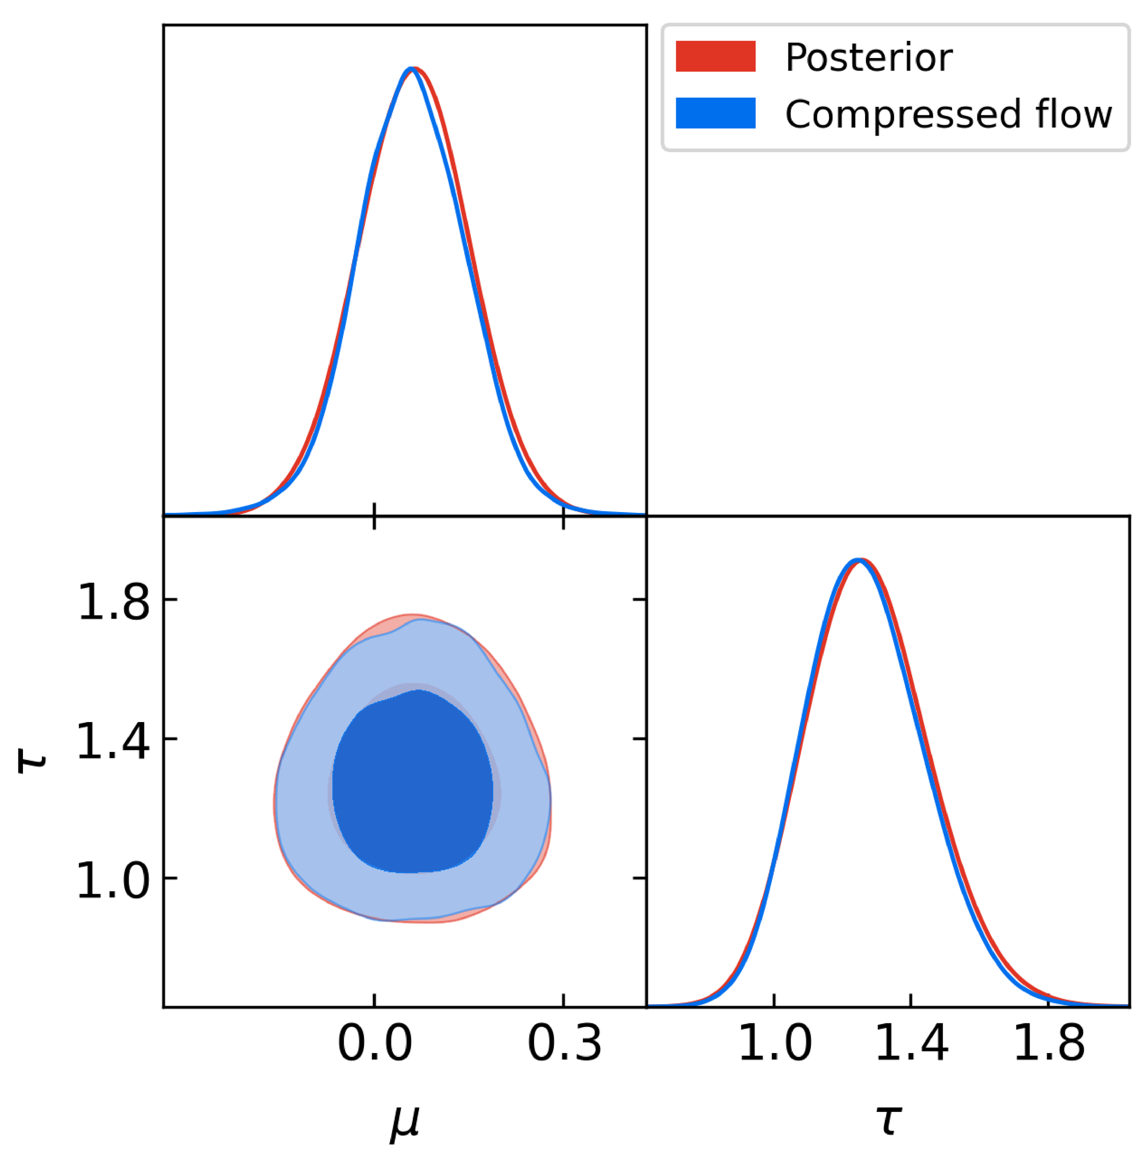

▶ Original harmonic mean estimator *insensitive* to prior τ_0 selection, *i.e.* a pathological failure [7].

▶ Learned harmonic mean estimator *accurate* and *sensitive* to prior.