HARMONIC: Bayesian model comparison for simulation-based inference

Matthew Docherty
mdochertyastro.com

UCL
Mullard Space Science Laboratory

with

Alessio Spurio Mancini

Matthew Price

Jason McEwen

Talk Outline

1. Learnt harmonic mean estimator for model comparison

2. Simulation-based inference in cosmology

3. Learnt harmonic mean estimator for simulation-based model comparison

4. Numerical examples

Talk Progress

1. Learnt harmonic mean estimator for model comparison

2. Simulation-based inference in cosmology

3. Learnt harmonic mean estimator for simulation-based model comparison

4. Numerical examples

Model Comparison- A Bayesian inference problem

I. Parameter Estimation

- Give a model, what parameters created the data
- Only requires unnormalised posterior

$$P(\theta|d) \propto \mathcal{L}(\theta)\pi(\theta)$$

II. Model Comparison

- Which model best describes observed data
- Relies on model evidence

$$P(\theta|d) = \frac{1}{z}\mathcal{L}(\theta)\pi(\theta)$$

$$z = P(d|\mathcal{M}) = \int \mathcal{L}(\theta)\pi(\theta)d\theta$$

- Ensures posterior is a true probability **density**
- Critical but computationally demanding - look for an estimator

Matthew Docherty

Original Harmonic Mean Estimator

$$\frac{1}{z} \equiv \rho = \mathbb{E}_{P(\theta|d)}[\frac{1}{\mathcal{L}(\theta)}]$$

Harmonic mean of likelihood given posterior samples

$$= \int \frac{1}{\mathcal{L}(\theta)} P(\theta|d) d\theta$$

Integral form of expectation

$$= \int \frac{1}{z} \frac{\pi(\theta)}{P(\theta|d)} P(\theta|d) d\theta$$

Substitute likelihood using **Bayes Theorem**

- Counter-intuitive Importance Sampling

If prior has wider tails than posterior - Problems...

Newton & Raftery (1994)

Original Harmonic Mean Estimator

Re-targeted Harmonic Mean Estimator

$$\rho = \mathbb{E}_{P(\theta|d)}[\frac{\psi(\theta)}{\mathcal{L}(\theta)\pi(\theta)}]$$

Re-targeted Criteria: (pretty general)

- Normalised
- Narrower tails than posterior

What is a good candidate for this target density?

Learnt Harmonic Mean Estimator

Optimal when target density exactly equals the normalised posterior……but don't have z to normalise!

$$\varphi(\theta) \overset{\text{ML}}{\simeq} \varphi^{\text{optimal}}(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{z}$$

Only depends on samples of the unnormalised posterior - sampler agnostic and extends easily to SBI setting.

Software Implementation - HARMONIC

GitHub harmonic | Tests passing | docs passing | codecov 96% | pypi package 1.1.0 | License GPL | arXiv 2111.12720

Harmonic

`harmonic` is an open source, well tested and documented Python implementation of the *learnt harmonic mean estimator* (McEwen et al. 2021) to compute the marginal likelihood (Bayesian evidence), required for Bayesian model selection.

Code: https://github.com/astro-informatics/harmonic
Docs: https://astro-informatics.github.io/harmonic/index.html

- Software best-practises - reviews/testing/docs
- Seamless integration with **emcee**

```
$ pip install harmonic
```

Talk Progress

1. Learnt harmonic mean estimator for model comparison
2. Simulation-based inference in cosmology
3. Learnt harmonic mean estimator for simulation-based model comparison
4. Numerical examples

Simulation-based inference (SBI)

Cosmological **motivation** - Intractable likelihoods for complex physics but can forward run model and simulate data

Many approaches for SBI but can largely be split into 2 groups:

1. ABC (Approx. Bayesian Computation)
2. Surrogate modelling

ML methods session - ML implementation of method 2 - Neural Density Estimation (NDE)

- Generate simulations then use to train NN for density estimation

- **Amortised -** Fast, - only condition on data post hoc - Generalises well
- **Sequential -** More efficient, - conditioned on data ad hoc - can mitigate bias

(Sequential) Neural Density Estimate - (S)NDE

(Sequential) Neural Posterior Estimate - (S)NPE *Papamakarios & Murray (2016)*

- Learn to approximate the normalised posterior directly with negative log likelihood loss

(Sequential) Neural Likelihood Estimate - (S)NLE *Papamakarios+ (2019)*

- Learn to approximate the likelihood using KL divergence
- Obtain posterior samples with MCMC

(Sequential) Neural Ratio Estimate - (S)NRE *Hermans+ (2019)*

- Indirectly learn the normalised posterior
- Practically done by training binary classifier to learn the likelihood ratio between joint and marginalised distributions
- Obtain normalised posterior samples with MCMC

Talk Progress

1. Learnt Harmonic mean estimator for model comparison
2. Simulation-based inference in Cosmology
3. Learnt harmonic mean estimator for simulation-based model comparison
4. Numerical examples

(S)NDE + Harmonic Methodology

Spurio Mancini, Docherty+ (in prep.)

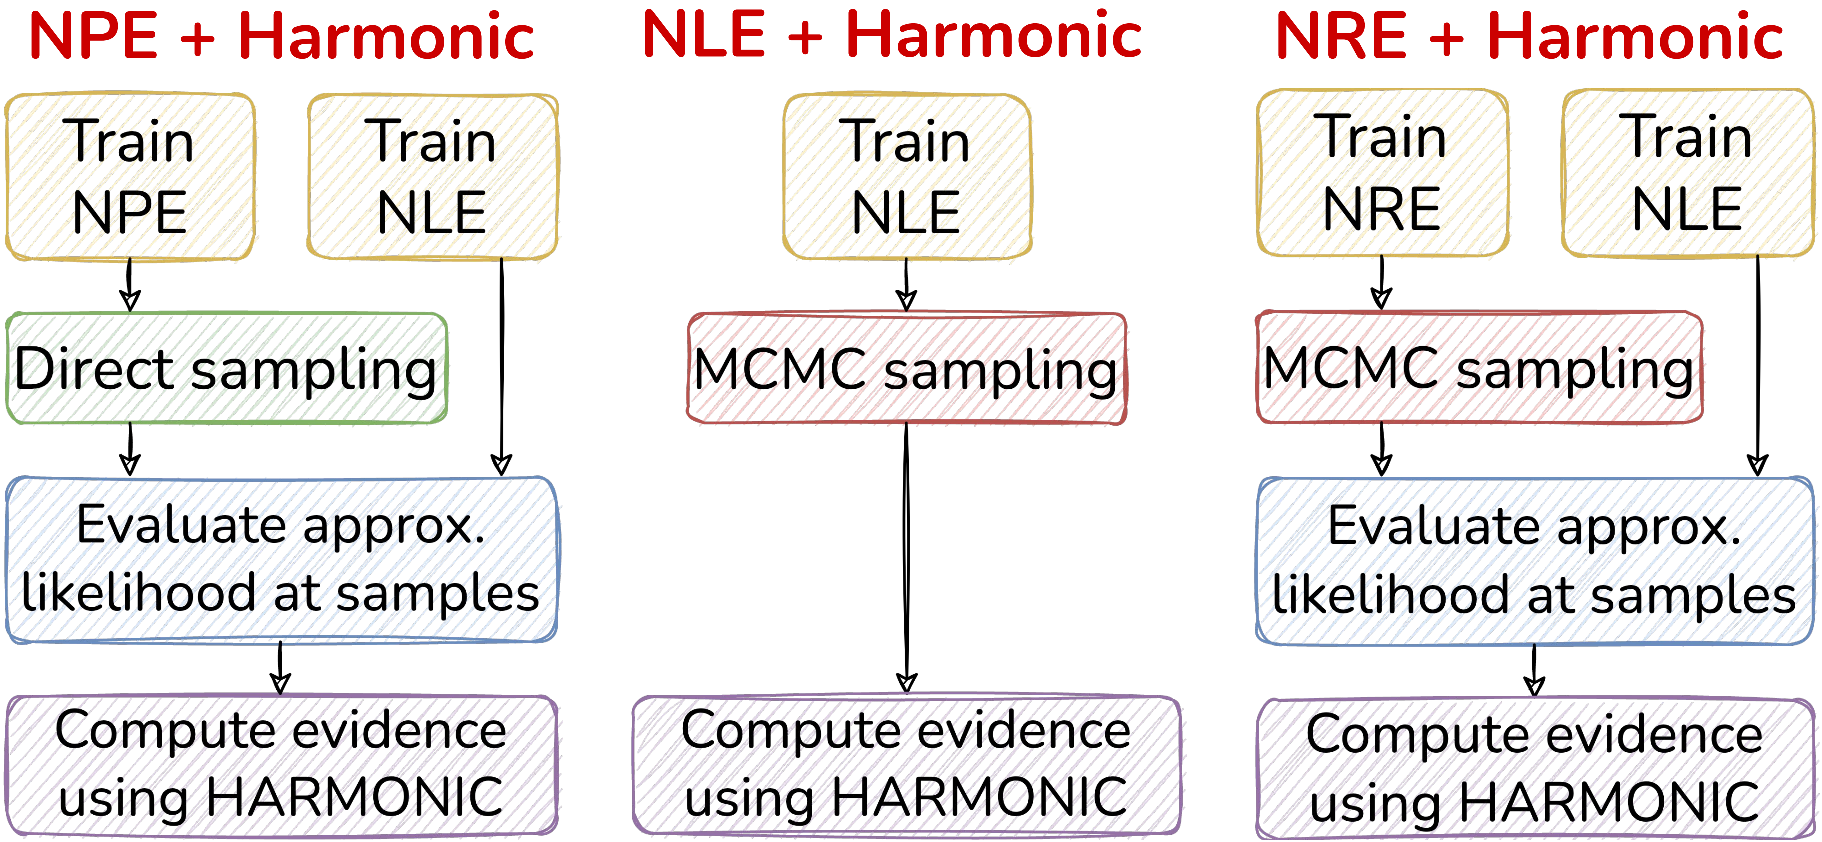

Methodological Properties

	SNPE	SNLE	SNRE
No external MCMC sampling	✅	❌	❌
No need for 2-stage composite inference	❌	✅	❌

Talk Progress

1. Learnt harmonic mean estimator for model selection

2. Simulation-based inference in cosmology

3. Learnt harmonic mean estimator for simulation-based model comparison

4. Numerical examples

Spurio Mancini, Docherty, Price & McEwen (in prep.)

$$d_i = \theta_i + \mathcal{N}(0,1)\,,\; i = 1,2,3$$

Linear Gaussian

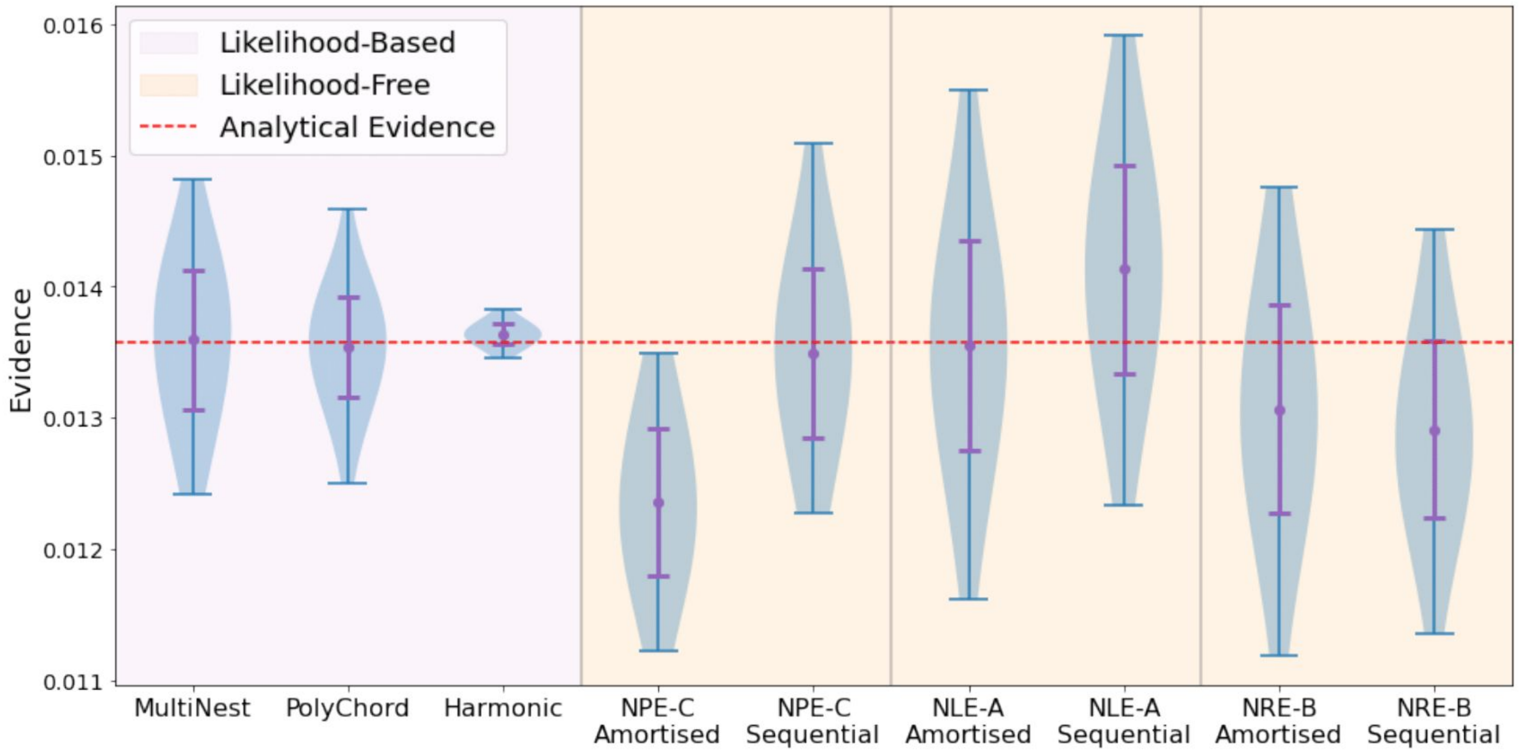

Matthew Docherty

Gravitational Waves

Background/Motivation:
GW model comparison - benefit of using higher-order model waveforms for inference?

Setup:
Simulate BBH merger using **pyCBC** to infer individual masses

	logZ (Source)	logZ (Aux)
Likelihood-based for validation (nested sampling)	-57.5	-59.3
Simulation-based (NLE + Harmonic)	-57.4	-59.6

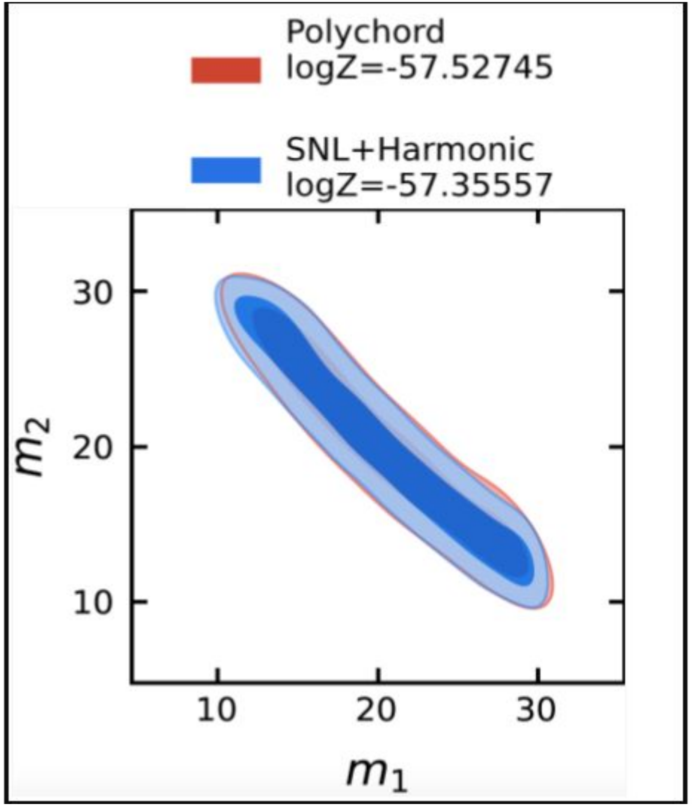

Summary Slide

- Harmonic **agnostic** to sampling strategy → **Essential for SBI evidence pipeline (no MCMC)**

- Introduced 3 novel methods of **simulation-based model comparison** with promising preliminary results

- Future work: extend to **higher dimensions** (both data space and parameter space)

- Come chat to me if you have posterior samples and would like to work Harmonic into your pipeline, or just run: **`pip install harmonic`**

https://github.com/astro-informatics/harmonic