


Accurate and efficient simulation-based inference for massive black-hole binaries with LISA

Dixième assemblée générale du GdR Ondes Gravitationnelles

Lyon, France - 7th October 2026


Alice Spadaro, J. Gair, D. Gerosa, S. Green, R. Buscicchio,
N. Gupte, R. Tenorio, S. Clyne, M. Püerrer, N. Korsakova

The millihertz gravitational-wave sky




LISA Definition Study Report (2024)

Challenges of LISA data analysis



- This work: massive black-hole binaries
 - The dual challenge:
 - Rapid low latency (EM follow-up)
 - High accuracy (astrophysical catalogs)
- ↓
- Tradeoff:
Fast (approximate) ↔ Accurate (expensive)
- Non-stationarities: data gaps & glitches

Challenges of LISA data analysis




- This work: massive black-hole binaries
- The dual challenge:
 - Rapid low latency (EM follow-up)
 - High accuracy (astrophysical catalogs)



Tradeoff:
Fast (approximate) ↔ Accurate (expensive)

- Non-stationarities: data gaps & glitches

The gravitational-wave inverse problem




$$\theta = \{m_1, m_2, \chi_1, \chi_2, d_L, \dots\}$$




Forward problem

Inverse problem


$$d = h_{GW}(\theta) + n$$



The gravitational-wave inverse problem




- Conventional stochastic samplers




- Likelihood evaluation: $p(d|\theta)$?
- Explore parameter space via sequential stochastic sampling

- Neural-posterior estimation



- Likelihood simulation: $d \sim p(d|\theta)$
- Train neural-network density estimator: $q(\theta|d) \approx p(\theta|d)$


Importance sampling



Sample $\theta_i \sim q(\theta | d)$ Corresponding weight $w_i = \frac{p(\theta_i | d)}{q(\theta_i | d)} = \frac{p(d | \theta_i) p(\theta_i)}{q(\theta_i | d)}$

Effective sample size $n_{\text{eff}} = \frac{\left(\sum_i w_i\right)^2}{\sum_i (w_i)^2}$ Sampling efficiency $\epsilon = \frac{n_{\text{eff}}}{n} \in (0, 1]$




Importance sampling



Sample $\theta_i \sim q(\theta | d)$ Corresponding weight $w_i = \frac{p(\theta_i | d)}{q(\theta_i | d)} = \frac{p(d | \theta_i) p(\theta_i)}{q(\theta_i | d)}$

Effective sample size $n_{\text{eff}} = \frac{\left(\sum_i w_i\right)^2}{\sum_i (w_i)^2}$ Sampling efficiency $\epsilon = \frac{n_{\text{eff}}}{n} \in (0, 1]$


Unleashing DINGO on LISA



Spadaro A. et al. (2026) – Phys. Rev. D 114, 023038

Deep INference for Gravitational-wave Observations (DINGO)

DINGO network





1. Data pipeline

- Waveform dataset: SVD-compressed polarisations (h_+ , $h_\times$)
- Data augmentation: on-the-fly extrinsic parameters & noise

2. Neural network



- Embedding (ResNet): low-dimensional context representation of strain data
- Conditional normalising flow: spline-based density estimator ($u \leftrightarrow \theta$)

Deep INference for Gravitational-wave Observations (DINGO)




M. Dax et al. (2021)


DINGO-LISA



Model validation




Stable convergence, no overfitting




Unbiased and calibrated posteriors

DINGO-LISA: low SNR




SNR ~ 90

Sampling efficiency: 50%




DINGO inference time: < 2 seconds for 20k samples

DINGO-LISA: moderate SNR




SNR ~ 500
Sampling efficiency: 12%




DINGO inference time: < 2 seconds for 20k samples

DINGO-LISA: high SNR



SNR ~ 1000
Sampling efficiency: 0.05%



DINGO inference time: < 2 seconds for 20k samples

Conclusions

- * Accurate SBI framework for high-mass LISA MBHBs
- * 20k posterior samples in < 1 minute
- * Agrees with conventional samplers up to $\text{SNR} \sim 500$; still unbiased at $\text{SNR} \sim 1000$, though less efficient


Next steps

- * Lower masses, longer and precessing signals
- * Posteriors updated as data accumulate, with gaps, glitches and overlaps in time/time-frequency domain
- * Validation on LISA Data Challenges

Thanks for the attention!

Normalising flows and neural posterior estimation

Posterior parametrisation:



$$\theta = f_d(u),$$

$$q(\theta|d) = \mathcal{N}_{[0,1]}(f_d^{-1}(\theta)) \left| \det J_{f_d}^{-1} \right|$$

Training objective:


$$D_{KL}(p|q) = \int dd p(d) \int d\theta p(\theta|d) \log \left(\frac{p(\theta|d)}{q(\theta|d)} \right)$$

$$D_{KL}(p|q) = -\mathbb{E}_{\theta \sim p(\theta)} \mathbb{E}_{d \sim p(d|\theta)} [\log q(\theta|d)]$$



Loss function

Forward KL

Reverse KL




LISA response



Short-lived signals: frozen, low-frequency approximation

LISA response: two LIGO-type detectors rotated from each other by $\pi/4$

Prior distribution of the training set



Focus on high-SNR sources

