Lightning-Fast Gravitational-Wave Parameter Inference Through Neural Amortization

Arnaud Delaunoy[1] Antoine Wehenkel[1] Tanja Hinderer[2,3] Samaya Nissanke[3] Christoph Weniger[3] Andrew R. Williamson[4] Gilles Louppe[1]

1 Liège université 2 Universiteit Utrecht 3 University of Amsterdam 4 University of Portsmouth

LIGO-G2100413

arXiv:2010.12931

1. Motivation

Stochastic sampling methods are widely used to carry out gravitational-wave (GW) parameter estimation, e.g., MCMC.

These analysis methods can take a significant amount of time and computational resources to analyse a single detection candidate.

This could create a significant bottleneck in the production of data products and scientific results as the rate of detection grows.

We show that neural simulation-based inference can provide a speed up factor of $\gtrsim \mathcal{O}(10^3)$.

- ϑ: GW parameters of interest
- θ: Nuisance parameters
- x: GW detector data

Given data containing a GW signal x_0, compute $p(\vartheta|x=x_0)$ based upon a model for $p(x|\vartheta,\theta)$ and a prior $p(\vartheta,\theta)$

$$p(\vartheta|x=x_0) = \frac{p(x_0|\vartheta)}{p(x_0)}p(\vartheta) = \frac{\int p(x_0|\vartheta,\theta)d\theta}{\int p(x_0|\vartheta,\theta)d\vartheta d\theta}\;p(\vartheta)$$

intractable

MCMC methods achieve this by sampling many times directly from $p(\vartheta,\theta|x=x_0)$, which can be very slow

2. Signal Model

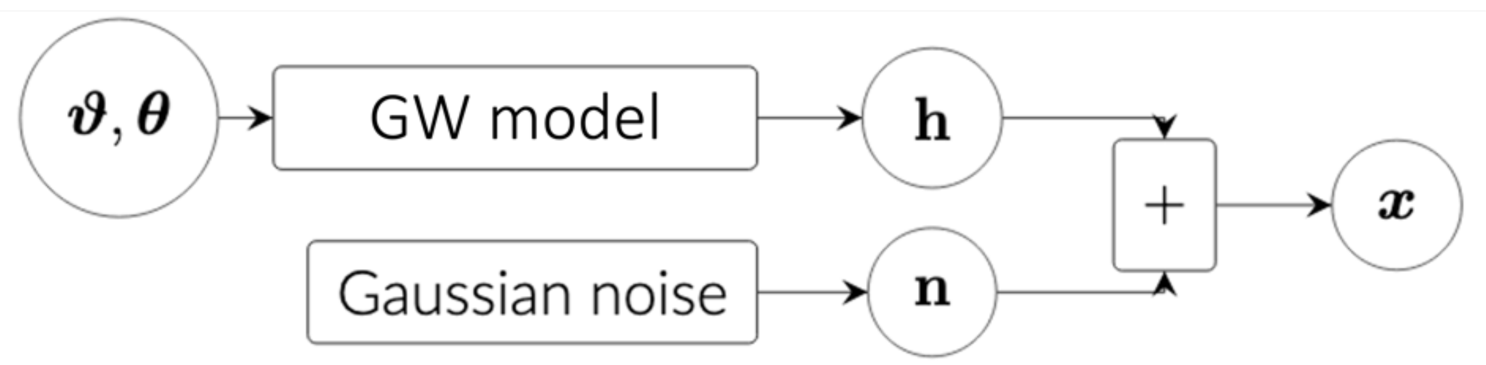

We consider precessing binary black hole signals h contained within 4 second data segments sampled at 2048 Hz. Data x are whitened and high-pass filtered at 20 Hz before being passed to the neural network.

3. Amortization

We aim to approximate the likelihood-to-evidence ratio

$$r(x|\vartheta) \equiv \frac{p(x|\vartheta)}{p(x)}$$

We train a convolutional neural network to discriminate between

- Associated data & parameter pairs:

$$(x,\vartheta) \sim p(x,\vartheta) \rightarrow y = 1$$

- Independent data & parameter pairs:

$$(x,\vartheta) \sim p(x)p(\vartheta) \rightarrow y = 0$$

We use it to compute an approximation of $r(x|\vartheta)$ [1]

$$\hat{r}(x|\vartheta) = \frac{s(x,\vartheta)}{1-s(x,\vartheta)}$$

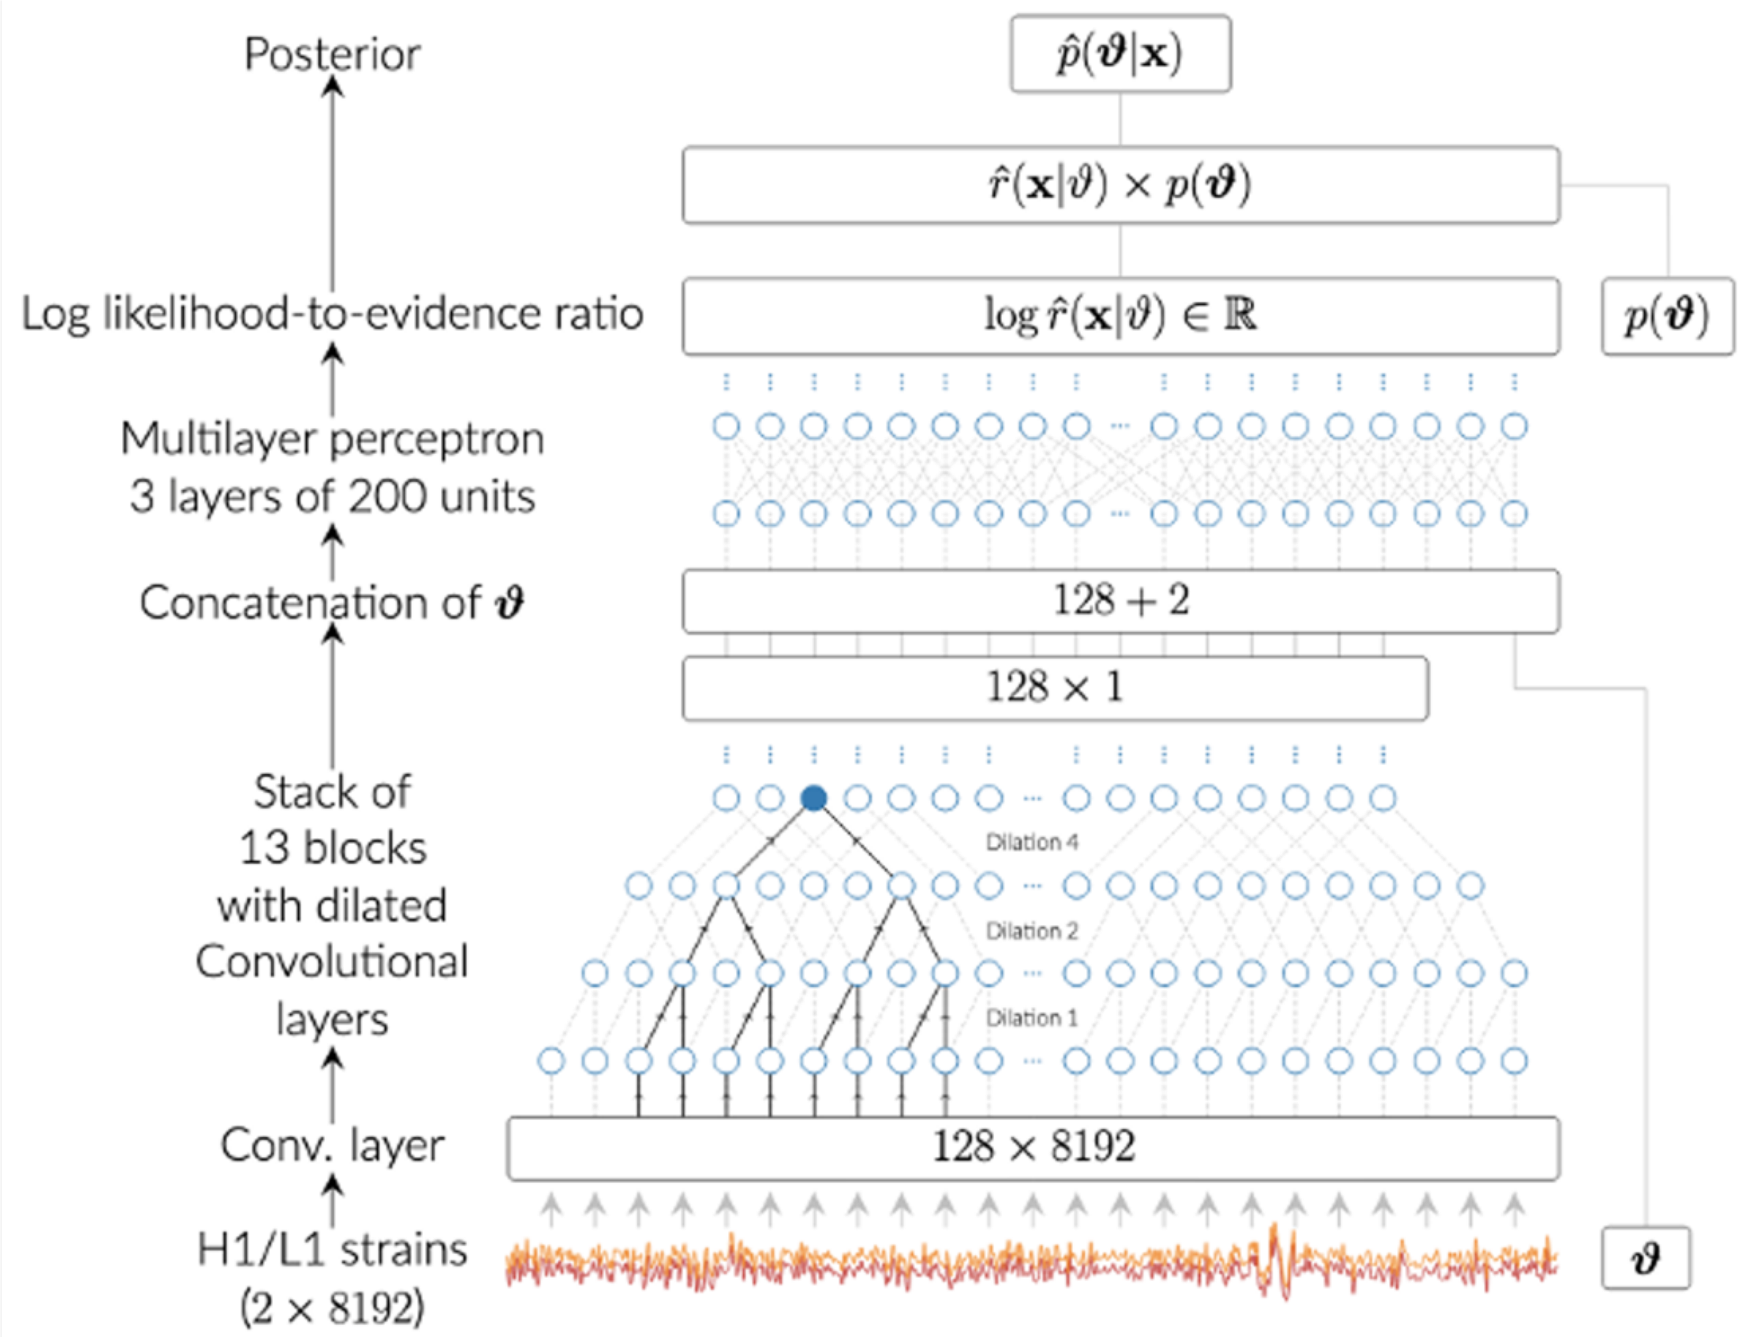

Figure: Network achitecture (adapted from [2])

Amortization

- Build a model for $p(\vartheta|x)$ beforehand (**slow**)
- Use this model to evaluate $p(\vartheta|x=x_0)$ (**fast**)

4. Results

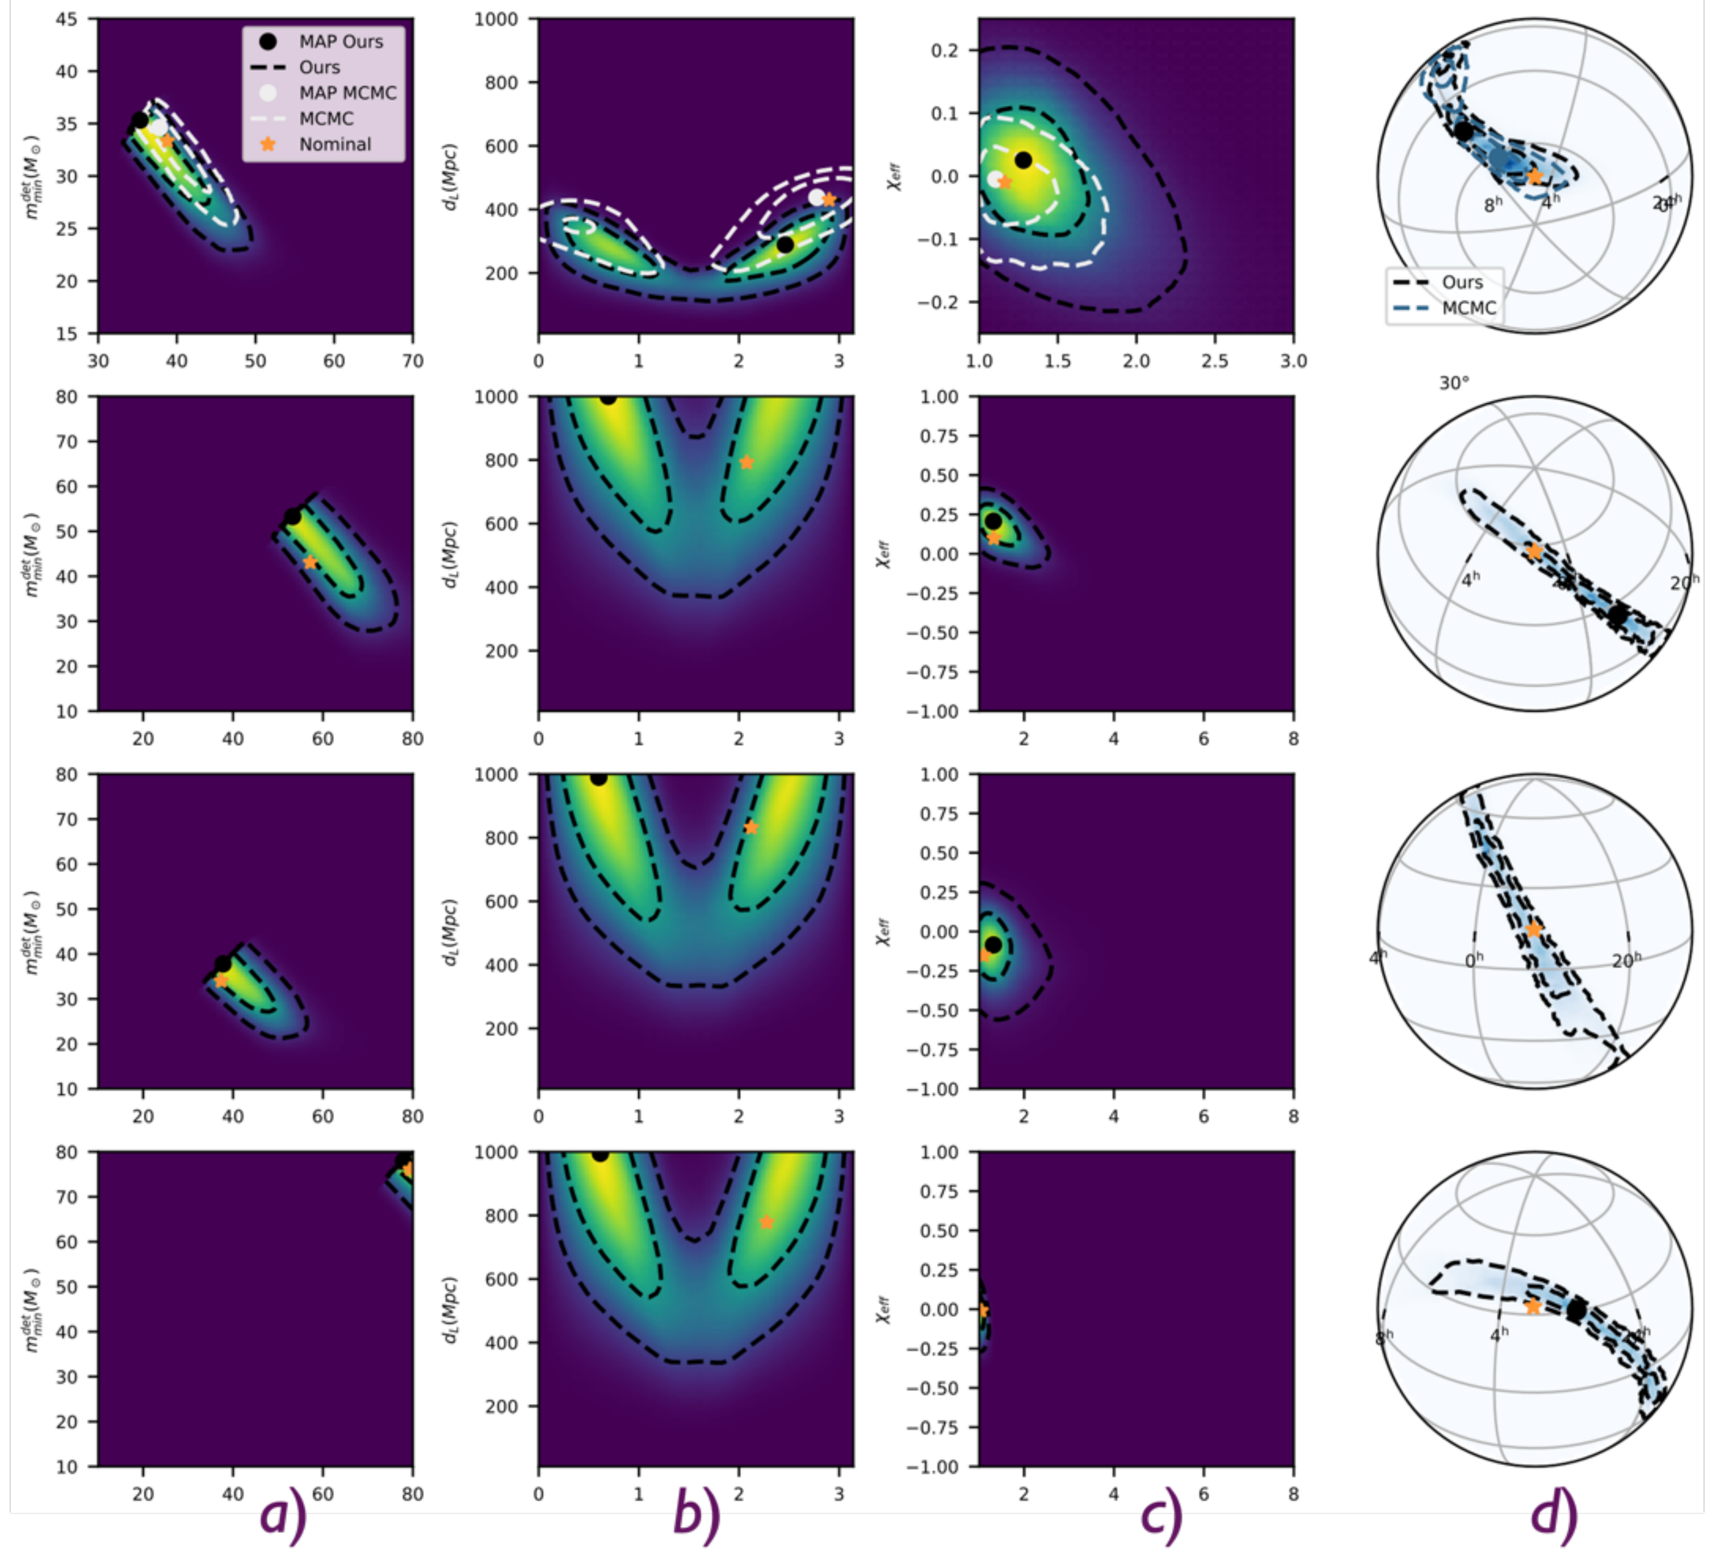

Figure: Marginal credible intervals produced by our neural networks with run times $\mathcal{O}(1\text{ minute})$. Columns: *a)* Primary v secondary masses; *b)* Orbital inclination v distance; *c)* Precessing spin param. v effective spin param.; *d)* sky location. 50% & 90% contours are shown. MCMC analysis shown in white for comparison (top row). Orange stars mark 'true' values; black/white spots mark maximum *a posteriori* values.

Take-home Message

- Neural amortization reduces inference time from days to minutes
- Our results are promising, though credible intervals tend to be slightly larger than those from MCMC
- Further assessments of the statistical validity of the estimated posteriors would be needed before making reliable scientific claims
- Work is ongoing to develop this work further via `swyft` [3]

References:

[1] Hermans et al., *arXiv:1903.04057*
[2] Gebhard et al., *Phys. Rev. D*, **100**(6):063015, (2019)
[3] Miller et al., *arXiv:2011.13951*