J. Mushkin

*jonathan.mushkin@weizmann.ac.il

Warp & Weft: Bringing the Industrial Revolution to Parameter Estimation

Jonathan Mushkin*[1], Javier Roulet[2], Barak Zackay[1], Tejaswi Venumadhav[3], Seth Olsen[4], Digvijay Wadekar[5] and Matias Zaldarriaga[5]

[1]Weizmann Institute of Science; [2]Caltech; [3]UC Santa Barbara; [4]Princeton; [5]Institute for Advanced Study

J. Roulet

Summary

The process of posterior evaluation is often the most demanding part of data analysis pipelines, both in human and CPU hours. We propose a new approach to the problem.

1. The likelihood can be factorized by dependence on either **intrinsic** or **extrinsic** parameters.
2. Using information from the detection-pipeline trigger, intrinsic and extrinsic samples are drawn **separately** and **at once**.
3. The respective components (e.g. waveforms, antennae response) are then computed.
4. The likelihood is evaluated at **all combinations** of intrinsic and extrinsic parameters, by using algebraic operations (i.e. "dot" product), as in the warp and weft of a fabric.

This has implications for **parameter estimation**; finding **best-fit** waveforms (consistency tests); and allowing an **optimal search test statistics** that includes the effects of **precession and higher emission modes**.

The Factorization

- We separate the 15 parameters of binary mergers [1] to *intrinsic* and *extrinsic*:

$$\theta_{\text{int.}} = m_1, m_2, \vec{s}_1, \vec{s}_2, \iota, \quad \theta_{\text{ext.}} = \alpha, \delta, \psi, t, \phi, D$$

- The waveform can then be decomposed by modes m, polarizations p, and functions of either intrinsic or extrinsic parameters:

$$h(f;\theta_{\text{int.}},\theta_{\text{ext.}}) = \sum_{\substack{m=1,\dots,4\\ p=+,\times}} h_{m,p}(f;\theta_{\text{int.}}) \times \frac{D_0}{D} \times F_p(\alpha,\delta,\psi) \times e^{im(\phi-\phi_0)} \times e^{-2\pi i(t-t_0)f}$$

- where F_p is the antenna response to the p polarization, and D_0, t_0 and ϕ_0 are some arbitrary values.
- Employ relative binning [2]: we can pre-compute weights $W^{\langle d|h\rangle}_{m,p,f_b}$, $W^{\langle h|h\rangle}_{m,m',p,p',f_b}$ (using a single reference $\theta_{\text{int.}}$), then use frequencies f_b of length $\mathcal{O}(10^2)$ instead of f of length $\mathcal{O}(10^5)$.

$$\begin{aligned}\log\mathcal{L}(\theta_{\text{int.}},\theta_{\text{ext.}}) \propto\ & \mathrm{Re}\langle d|h\rangle - \frac{1}{2}\langle h|h\rangle = \\ & \mathrm{Re}\sum_{m,p,f_b}\left[W^{\langle d|h\rangle}h\right]_{m,p,f_b,\theta_{\text{int.}}} \times \left[\frac{D_0}{D}F_p e^{im(\Delta\phi)}e^{-2i\pi f_b\Delta t}\right]_{m,p,f_b,\theta_{\text{ext.}}} \\ & -\frac{1}{2}\sum_{\substack{m,m',\\ p,p'}}\left[\sum_{f_b}W^{\langle h|h\rangle}h\,h^*\right]_{\substack{m,m',\\ p,p',\theta_{\text{int.}}}} \times \left[\left(\frac{D_0}{D}\right)^2 F_p F_{p'} e^{i(m-m')\Delta\phi}\right]_{\substack{m,m',\\ p,p',\theta_{\text{ext.}}}}\end{aligned}$$

Sampling Procedure

- Samples drawn using importance / rejection sampling, given the detection trigger. Employ quasi Monte Carlo for better convergence.
- D and ϕ can be marginalized over.
- No MCMC / nested sampling required.
- Current benchmarks: $N_{\text{int.}} = 10^4$, $N_{\text{ext.}} = 10^3$ (see Figure 1). Include 4 modes, 2 polarization, ~ 400 frequencies for $\sim 10^{10}$ operations. After optimization, we aim for a run-time of $\sim 1 - 10$ seconds.

Implication for Search Pipelines

The optimal test statistic (the evidence ratio [3]) can be computed with samples from a single-detector posterior, along with mock-data produced cheaply by timeslides of the other detectors, for the background (null hypothesis H_0) and foreground (signal hypothesis H_1). For example, when three detectors are available:

$$\mathcal{Z} = \int \Pi\mathcal{L}d\theta = \int \left(\Pi\mathcal{L}^{(1)}\right)\mathcal{L}^{(2)}\mathcal{L}^{(3)}d\theta \approx \frac{\mathcal{Z}^{(1)}}{N_{\text{samples}}}\sum_{\theta\sim\Pi\mathcal{L}^{(1)}}\mathcal{L}^{(2)}\mathcal{L}^{(3)}$$

$$\log\mathcal{L}^{(1,2)} \propto \mathrm{Re}\,\langle d|h(\theta)\rangle - \frac{1}{2}\langle h(\theta)|h(\theta)\rangle$$

Timeslide detectors (2,3) by FFT for mock data

$$\langle d|h(\theta)\rangle \rightarrow \langle d|h(\theta)\rangle + \langle h(\theta^{\text{injection}})|h(\theta)\rangle$$

Foreground distribution by injections of waveforms into mock data

- The test statistic includes higher modes and precession, and is coherent between detectors.
- From a single event, the probability ratio $\Pr(\mathcal{Z}|H_1)/\Pr(\mathcal{Z}|H_0)$ is obtained.
- The method demands a fast and reliable way to produce samples from a posterior.
- Comes as a refined search, after a coarser search pipeline.

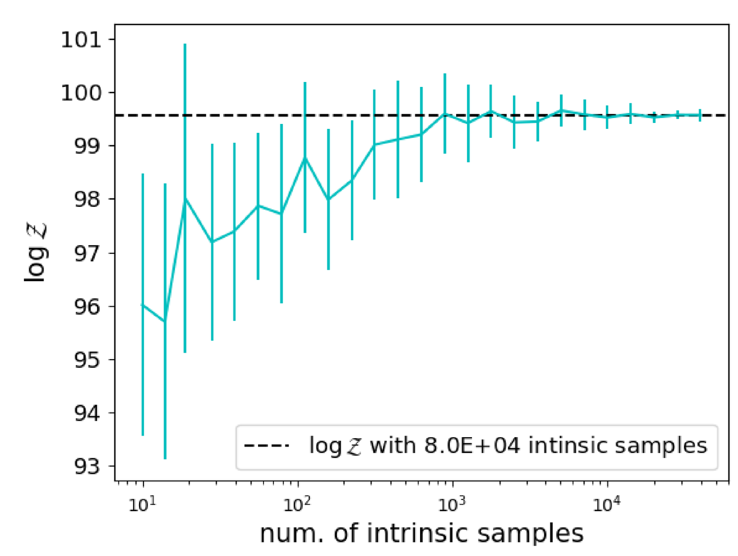

Figure 1: Convergence test of a signal injected in Livingston, for fixed $N_{\text{ext.}} \approx 10^3$ and different $N_{\text{int.}}$

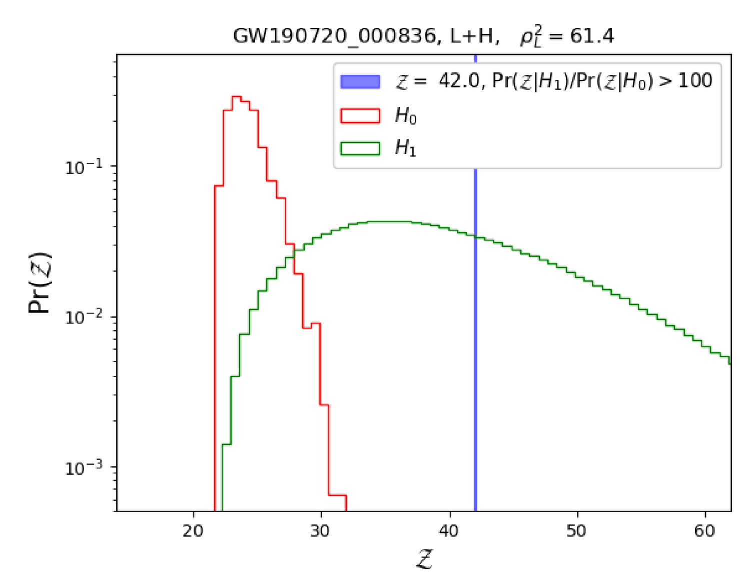

Figure 2: Declared event GW190623_183652. High $\Pr(\mathcal{Z}|H_1)/\Pr(\mathcal{Z}|H_0)$ is in agreement with reported IFAR [4, 5]. Sampling was performed on Livingston, timeslides performed on Hanford.

Astrophysical Implication

The optimal test statistic will improve the sensitivity to binary mergers exhibiting precession or higher modes. The discovery of precessing binaries would suggest dynamical formation, while the lack thereof would hint towards isolated co-evolution [6].

References

[1] J. Roulet *et al.* (2022), arXiv:2207.03508, accepted to PRD

[2] N. Leslie, L. Dai, G. Pratten (2021), 10.1103/PhysRevD.104.123030.

[3] J. Neyman and E. Pearson (1933), 10.1098/rsta.1933.0009.

[4] R. Abbott *et al.* (2021) 10.1103/PhysRevX.11.021053

[5] A. Nitz *et al.* (2021), arxiv:2112.06878

[6] M. Mapelli (2021), in doi.org/10.1007/978-981-15-4702-7.