Introduction to gravitational-wave parameter estimation

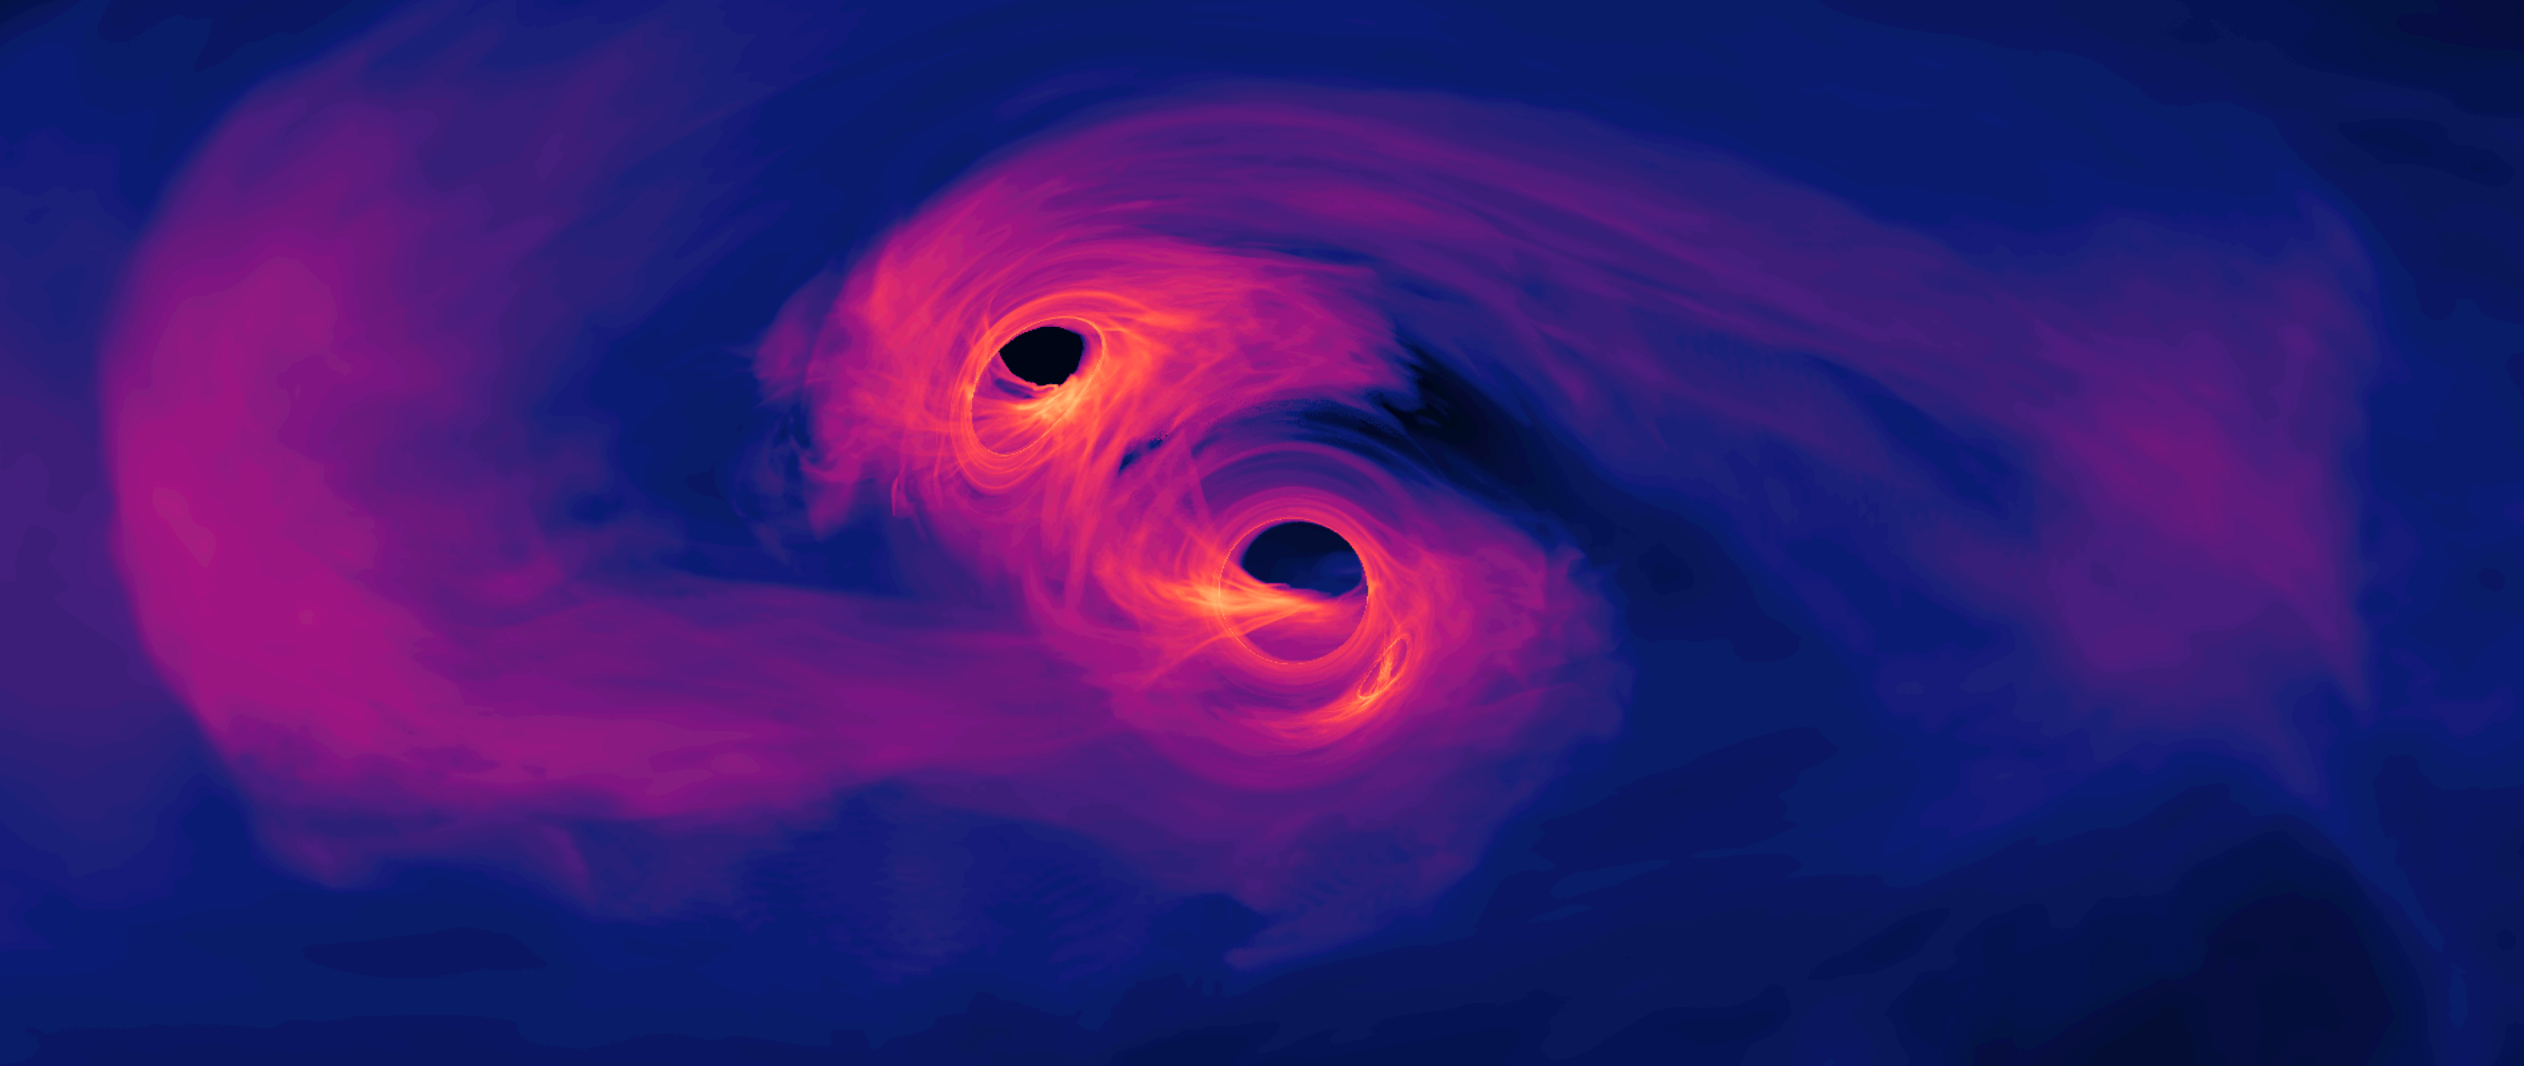

Alan M. Knee

GWANW 2023 Student Workshop

June 26 @ LHO

dcc.ligo.org/G2201051

Image credit: NASA GSFC/Scott Noble

Gravitational-wave sources

- Ground-based detectors are sensitive to gravitational waves (GWs) from several sources, including stellar-mass compact binary coalescences (CBCs)

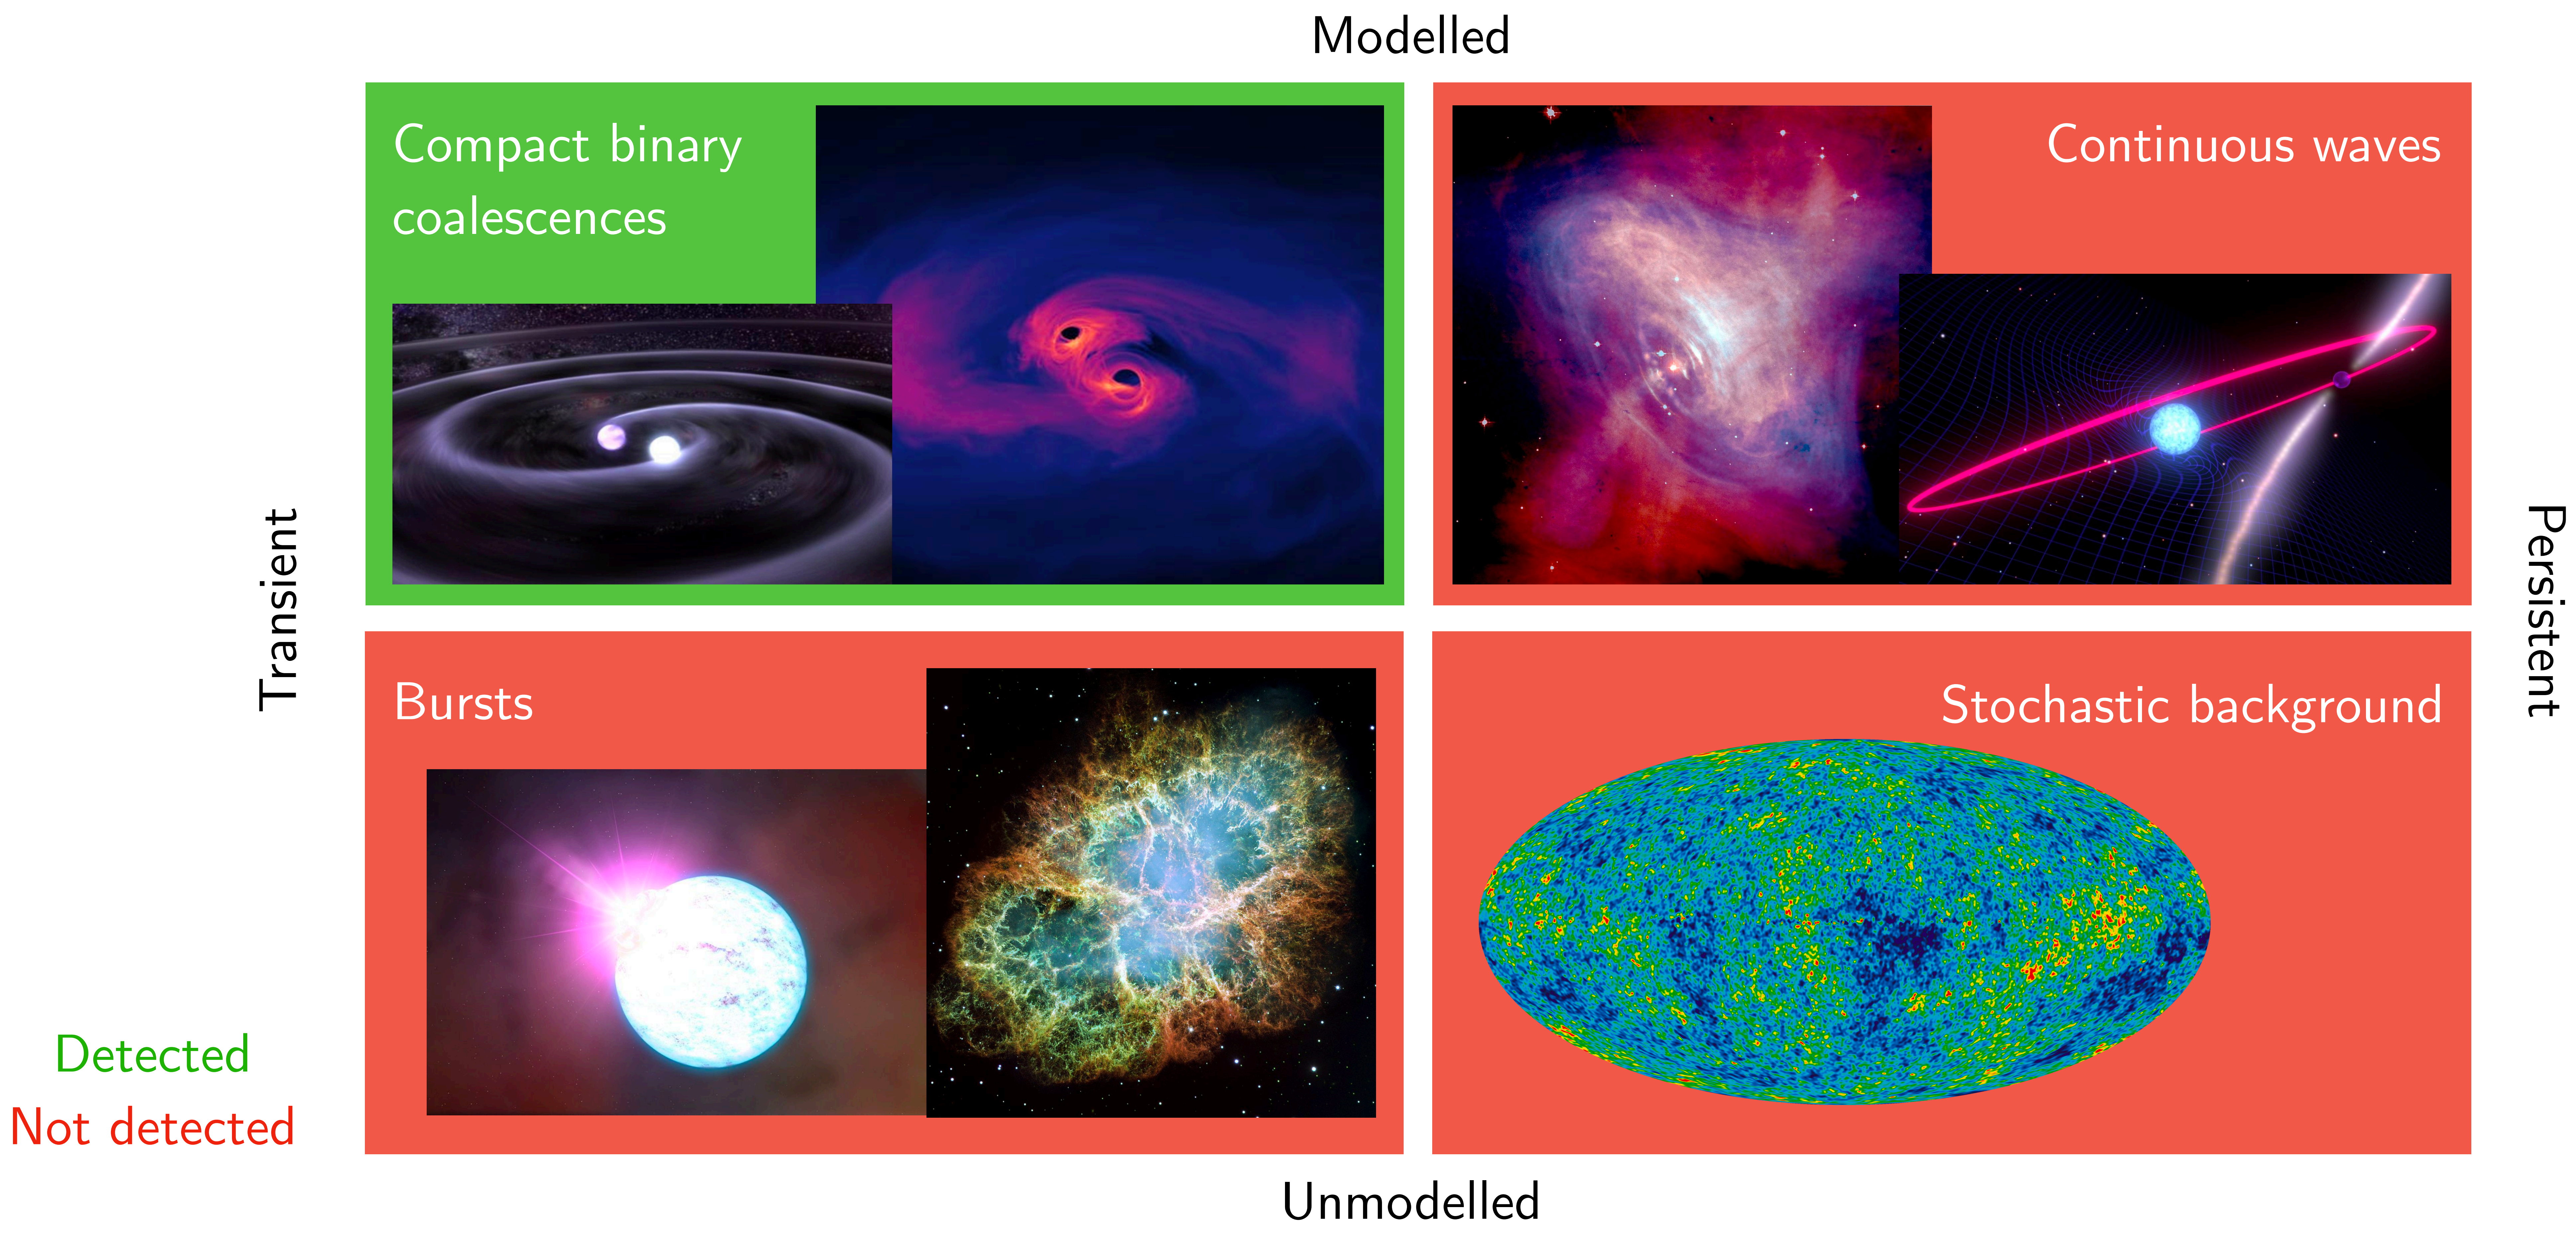

What we want to know

- GWs tell us about the physics of the binary that emitted them, such as ...

Masses

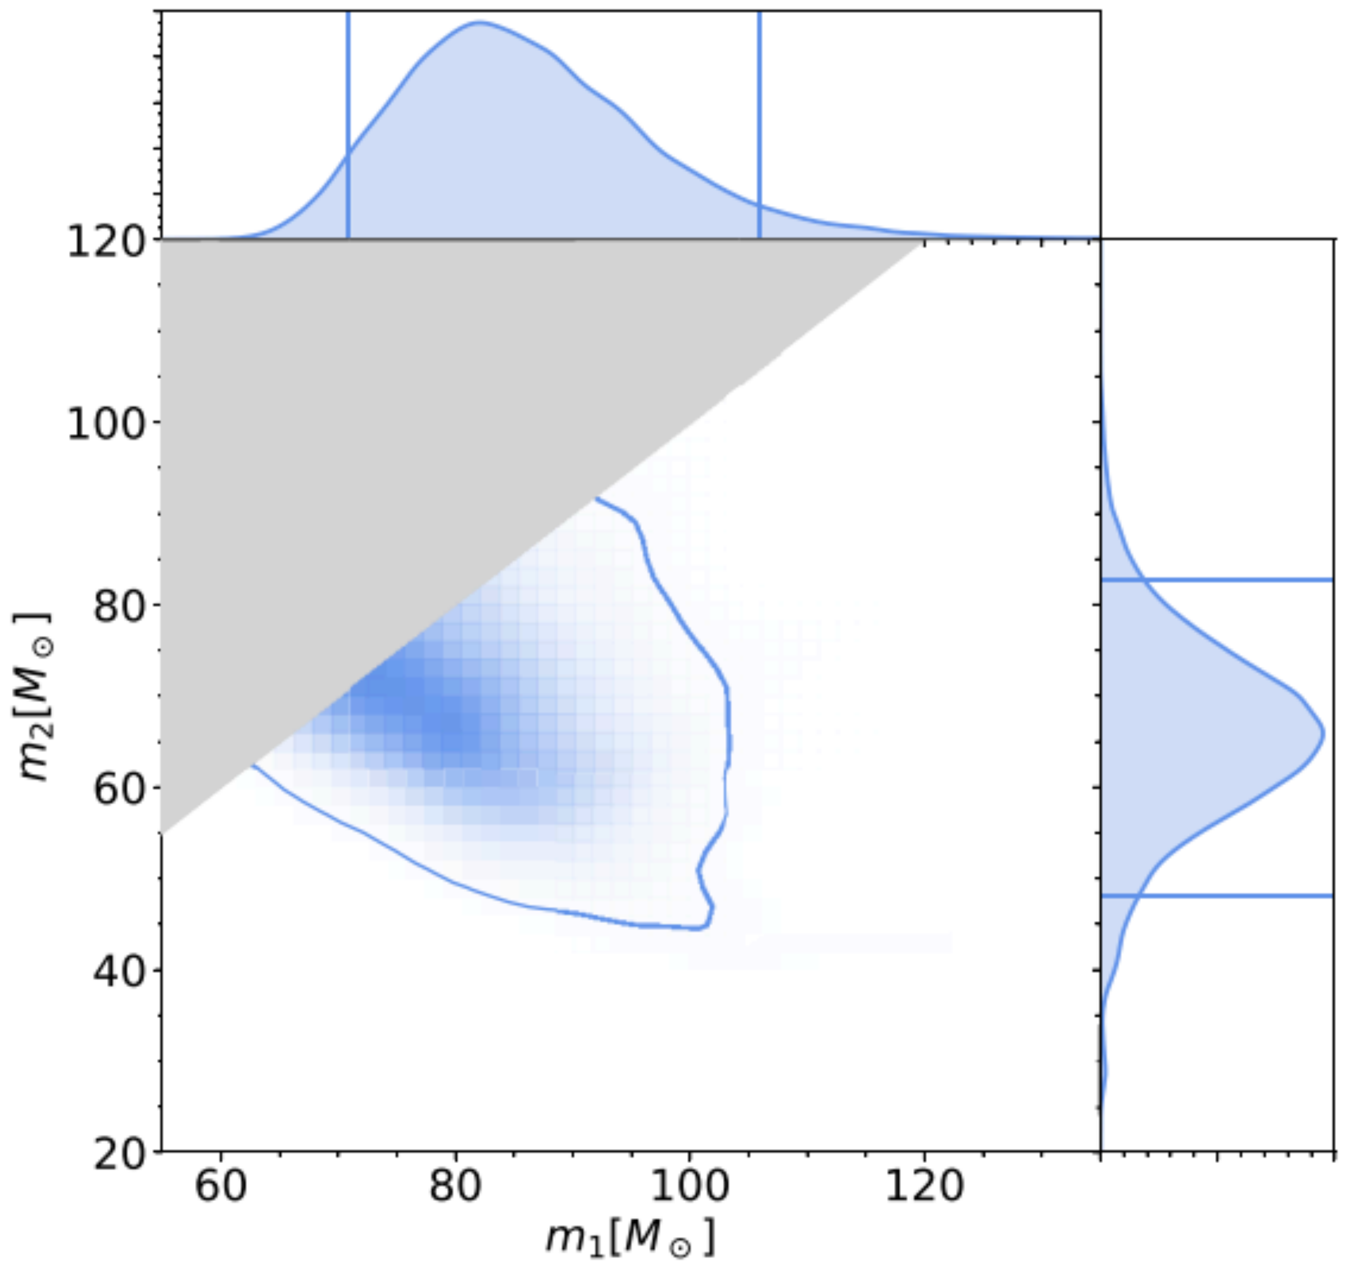

[2009.01075]

Spins

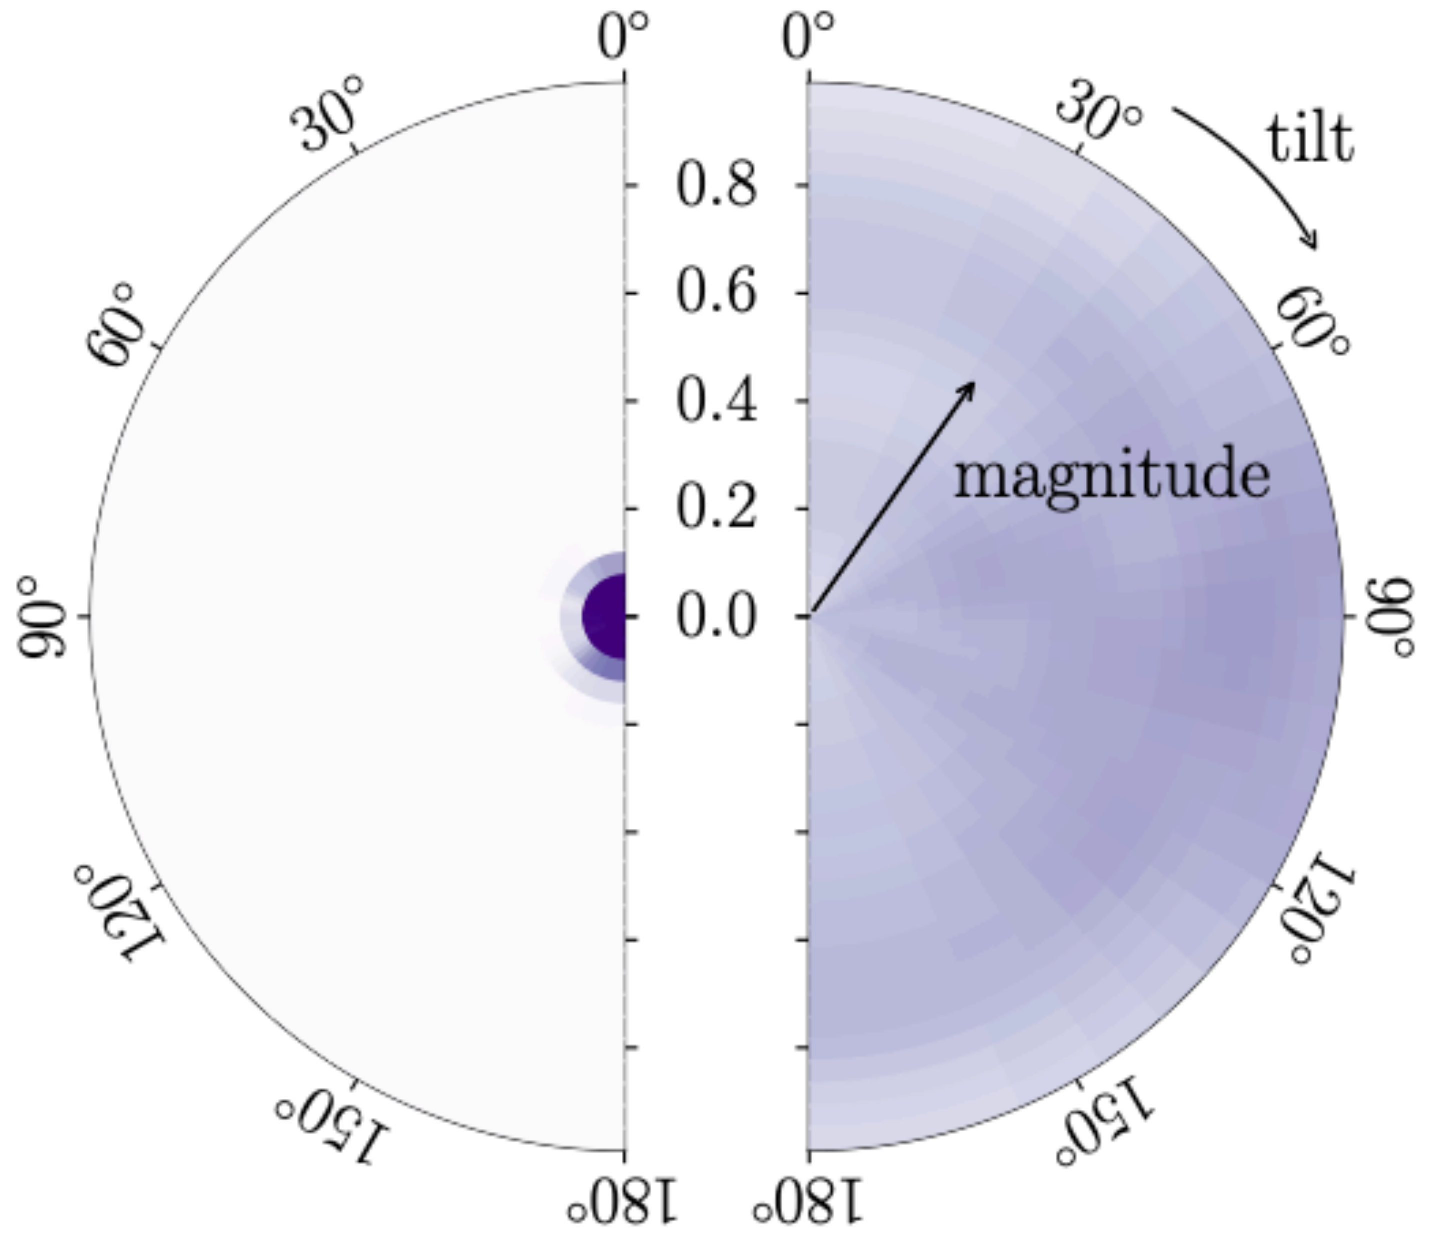

[2006.12611]

Location

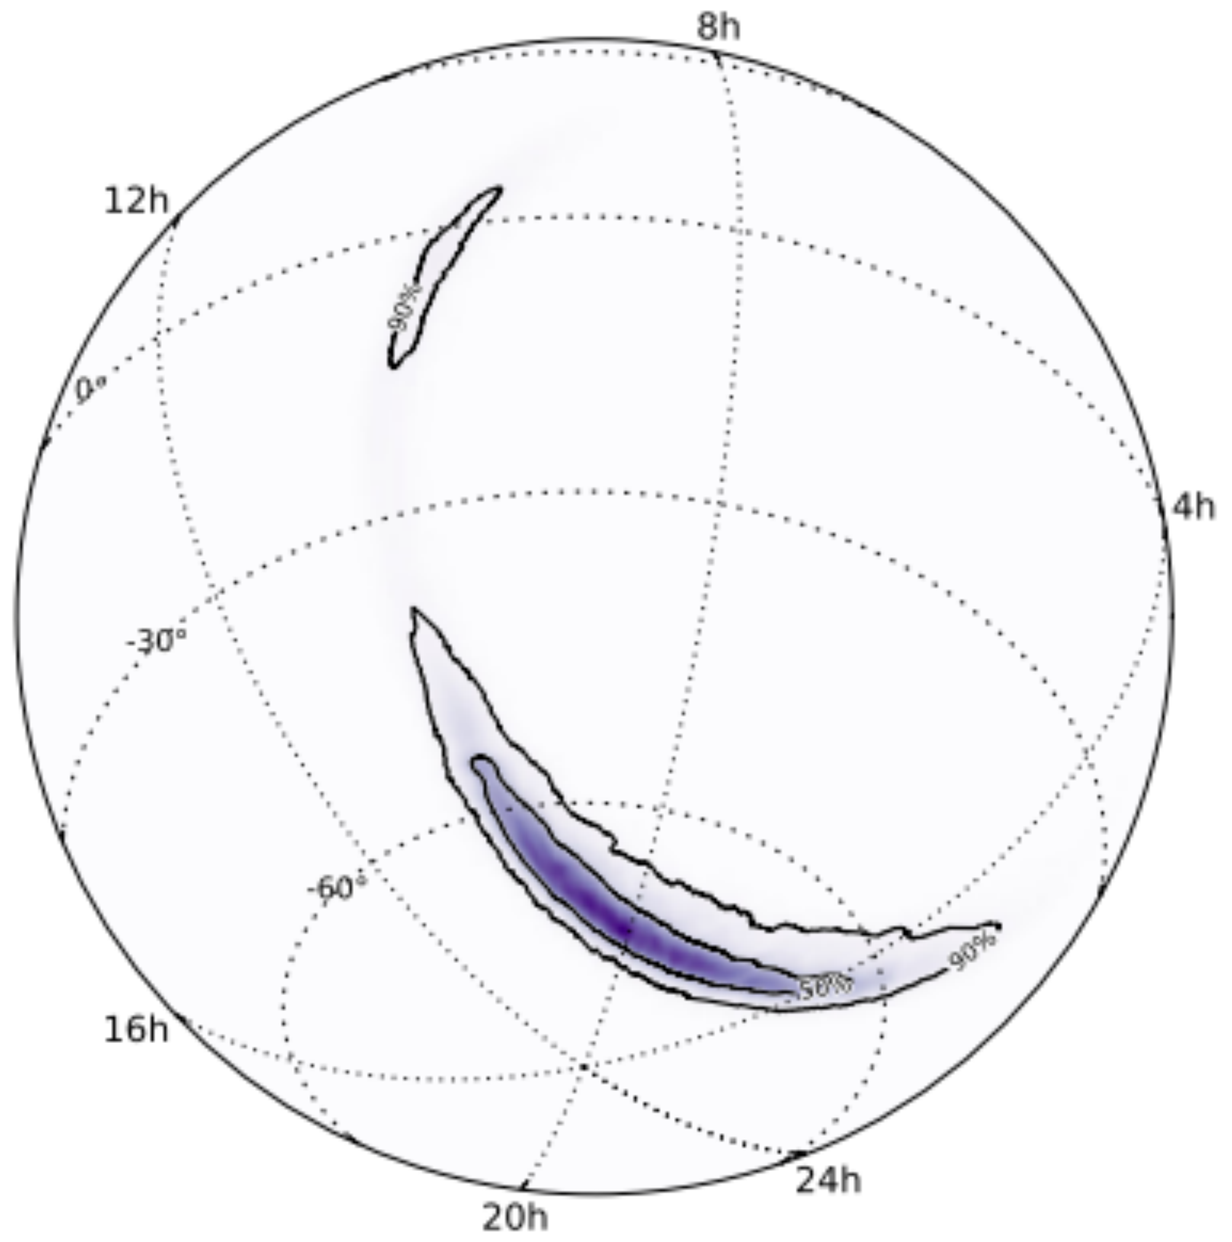

[1602.03840]

- Goal of parameter estimation is to measure these properties from the GW signal

Outline

- How to estimate parameters from the strain data? Use the framework of Bayesian inference

Strain data

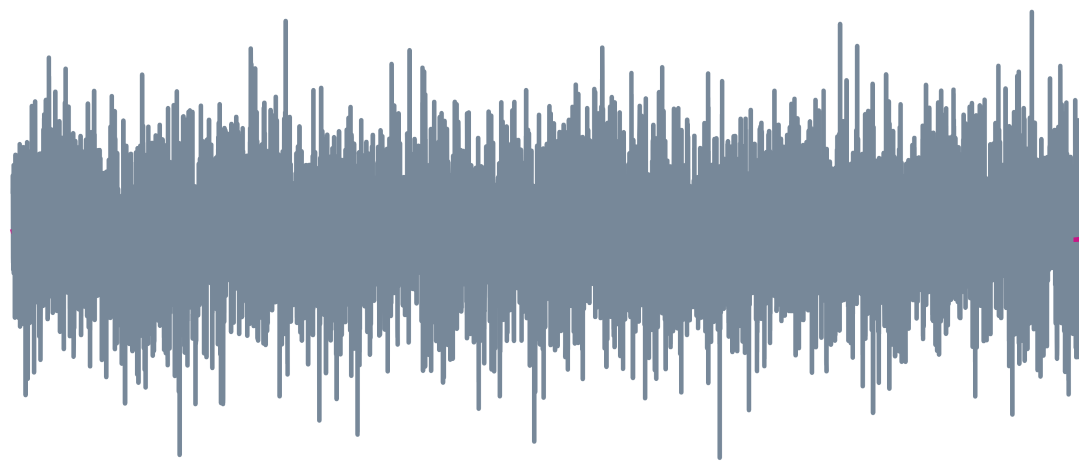

Waveform model

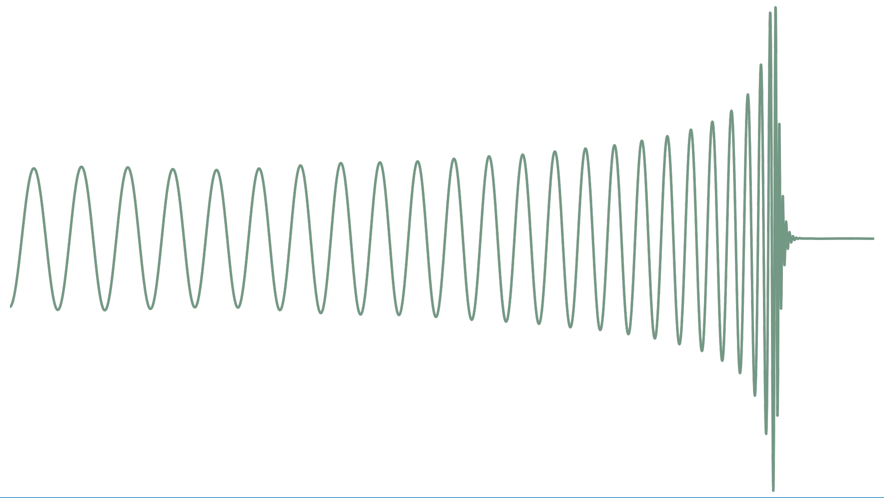

Bayesian inference

$$P(A\,|\,B) = \frac{P(B\,|\,A)P(A)}{P(B)}$$

Posterior distributions

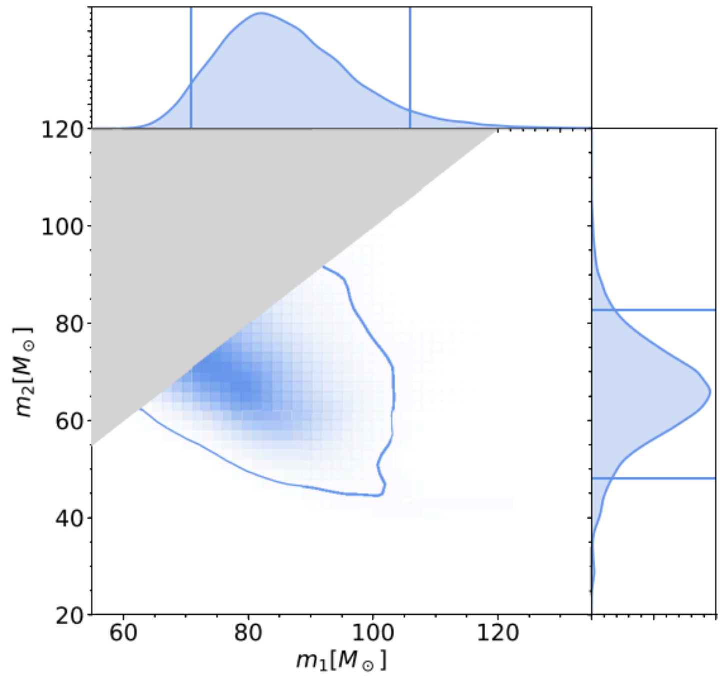

Waveform model

$$f_{\mathrm{GW}} = 2f_{\mathrm{orb}}$$

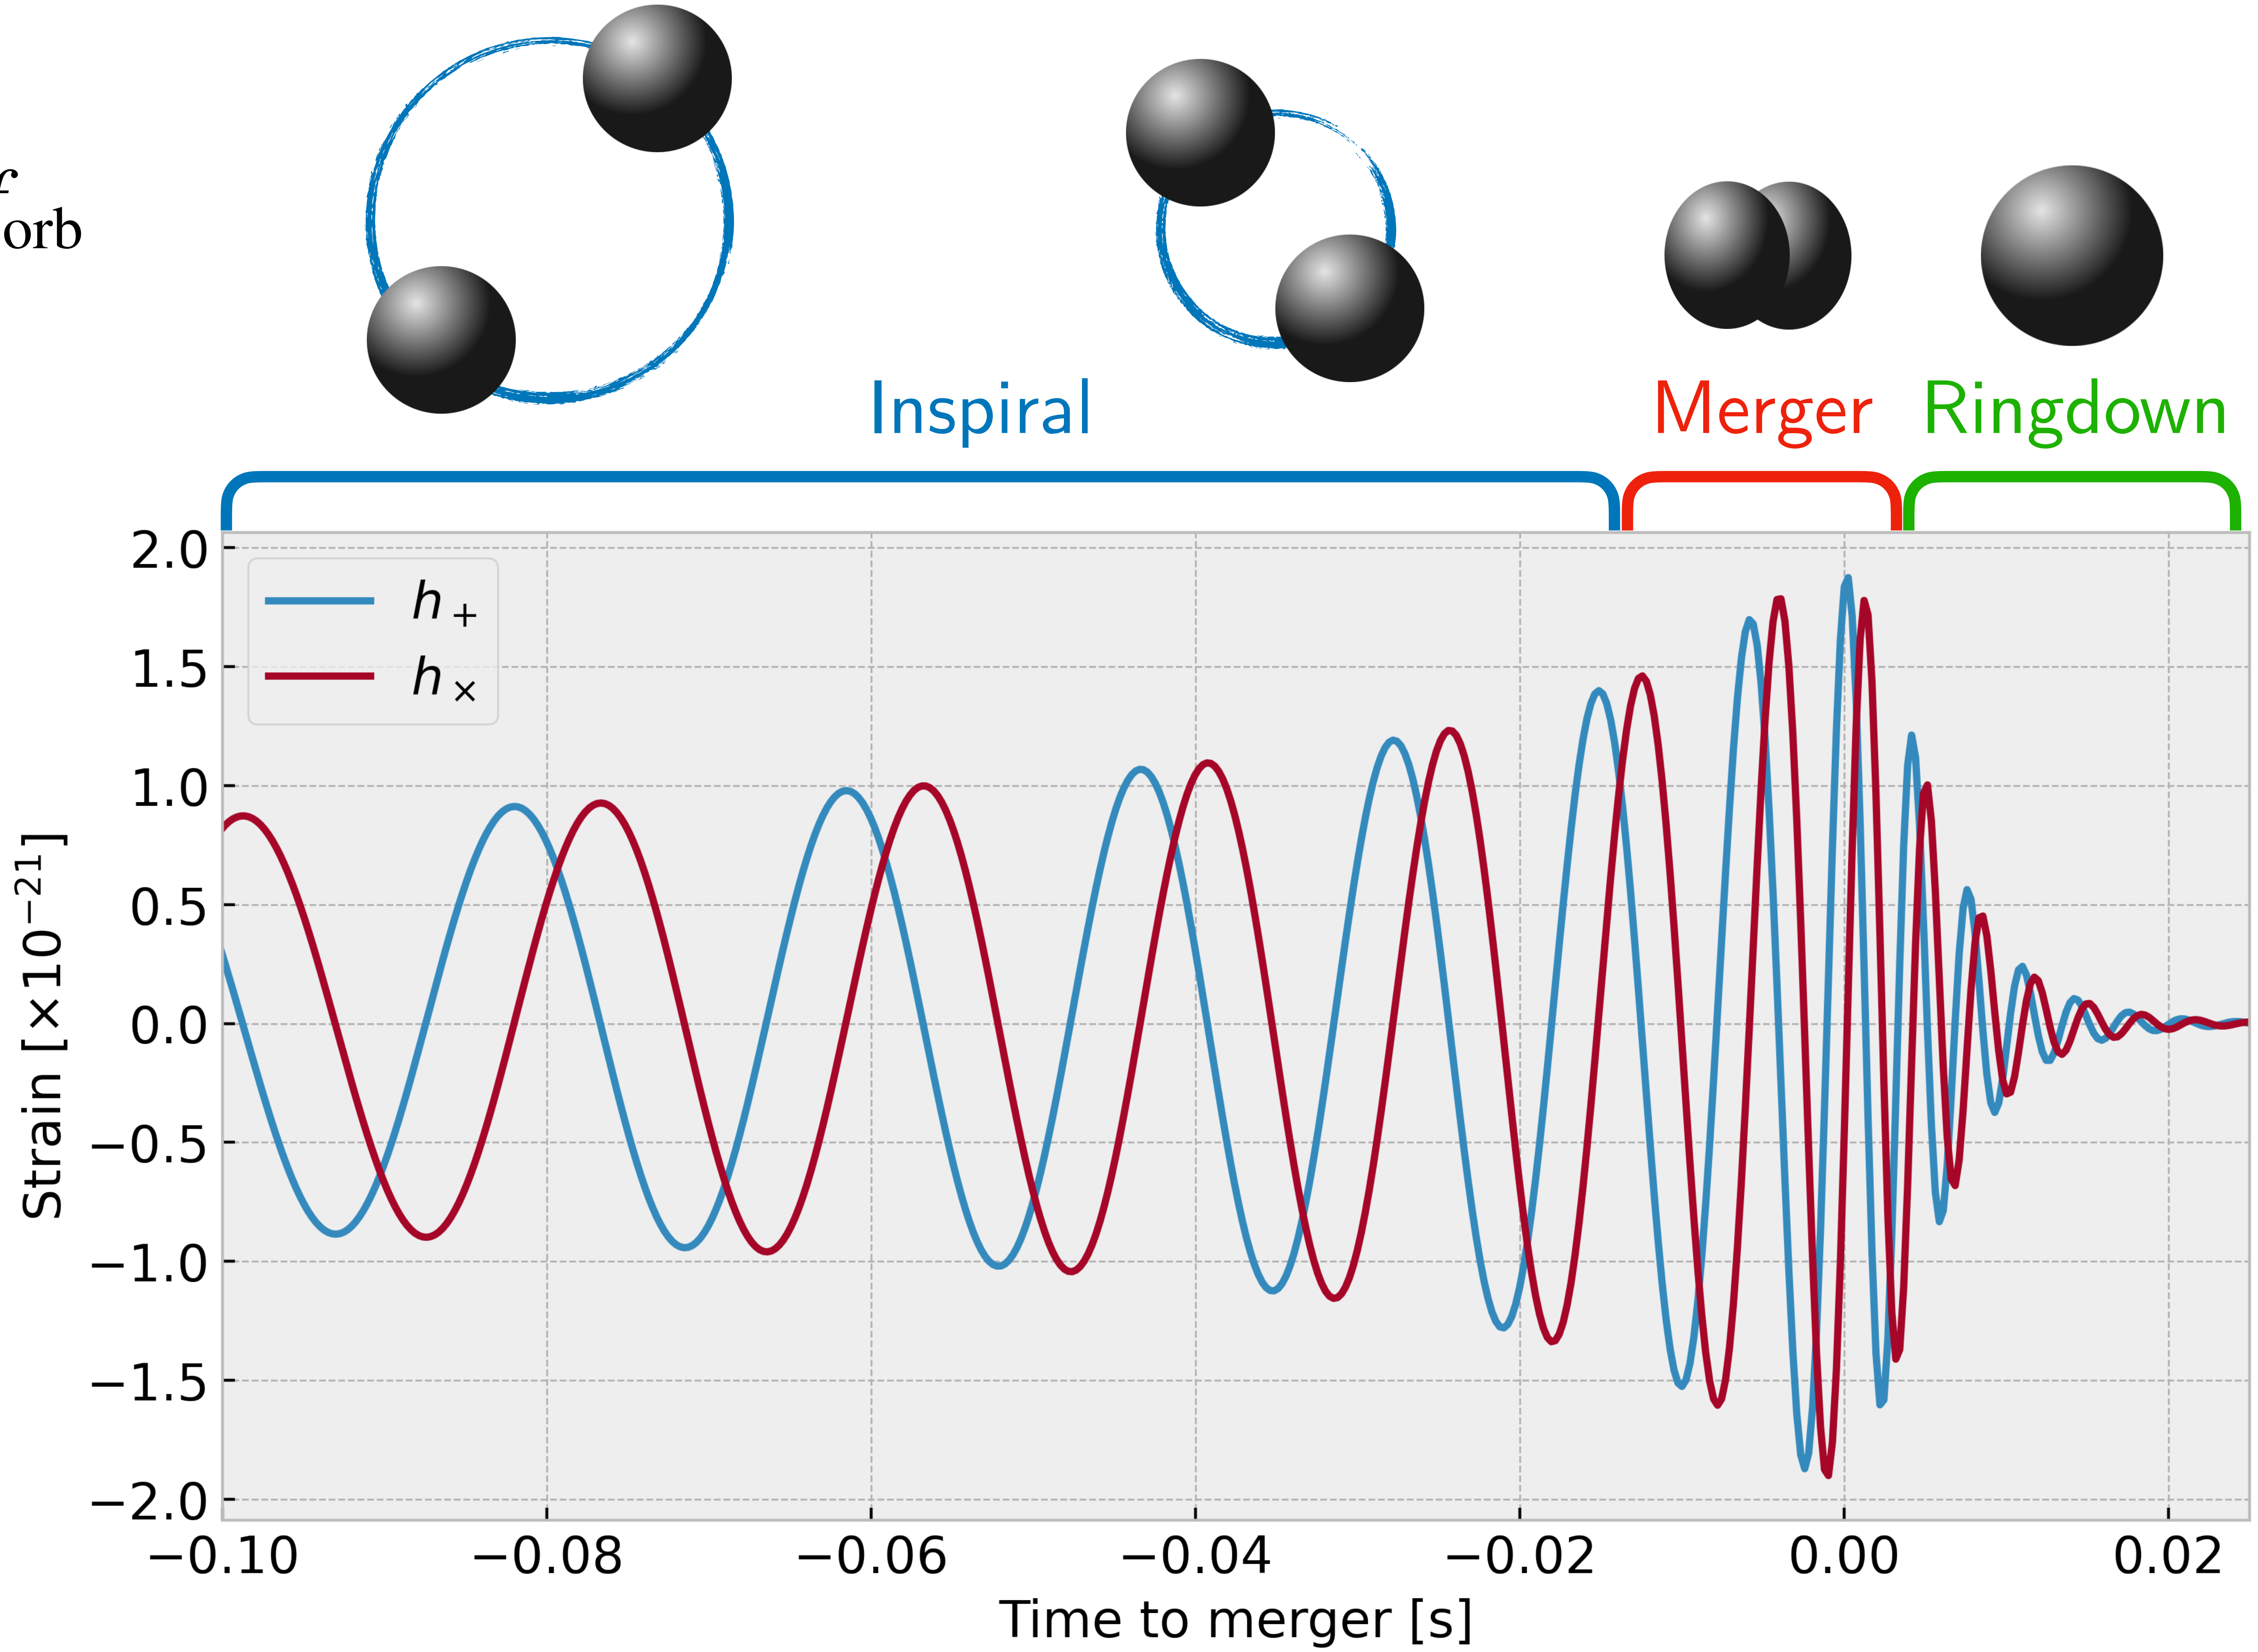

Parameters

- In general, 8 *intrinsic* parameters: two masses, six spin elements
- Neutron stars add a tidal parameter

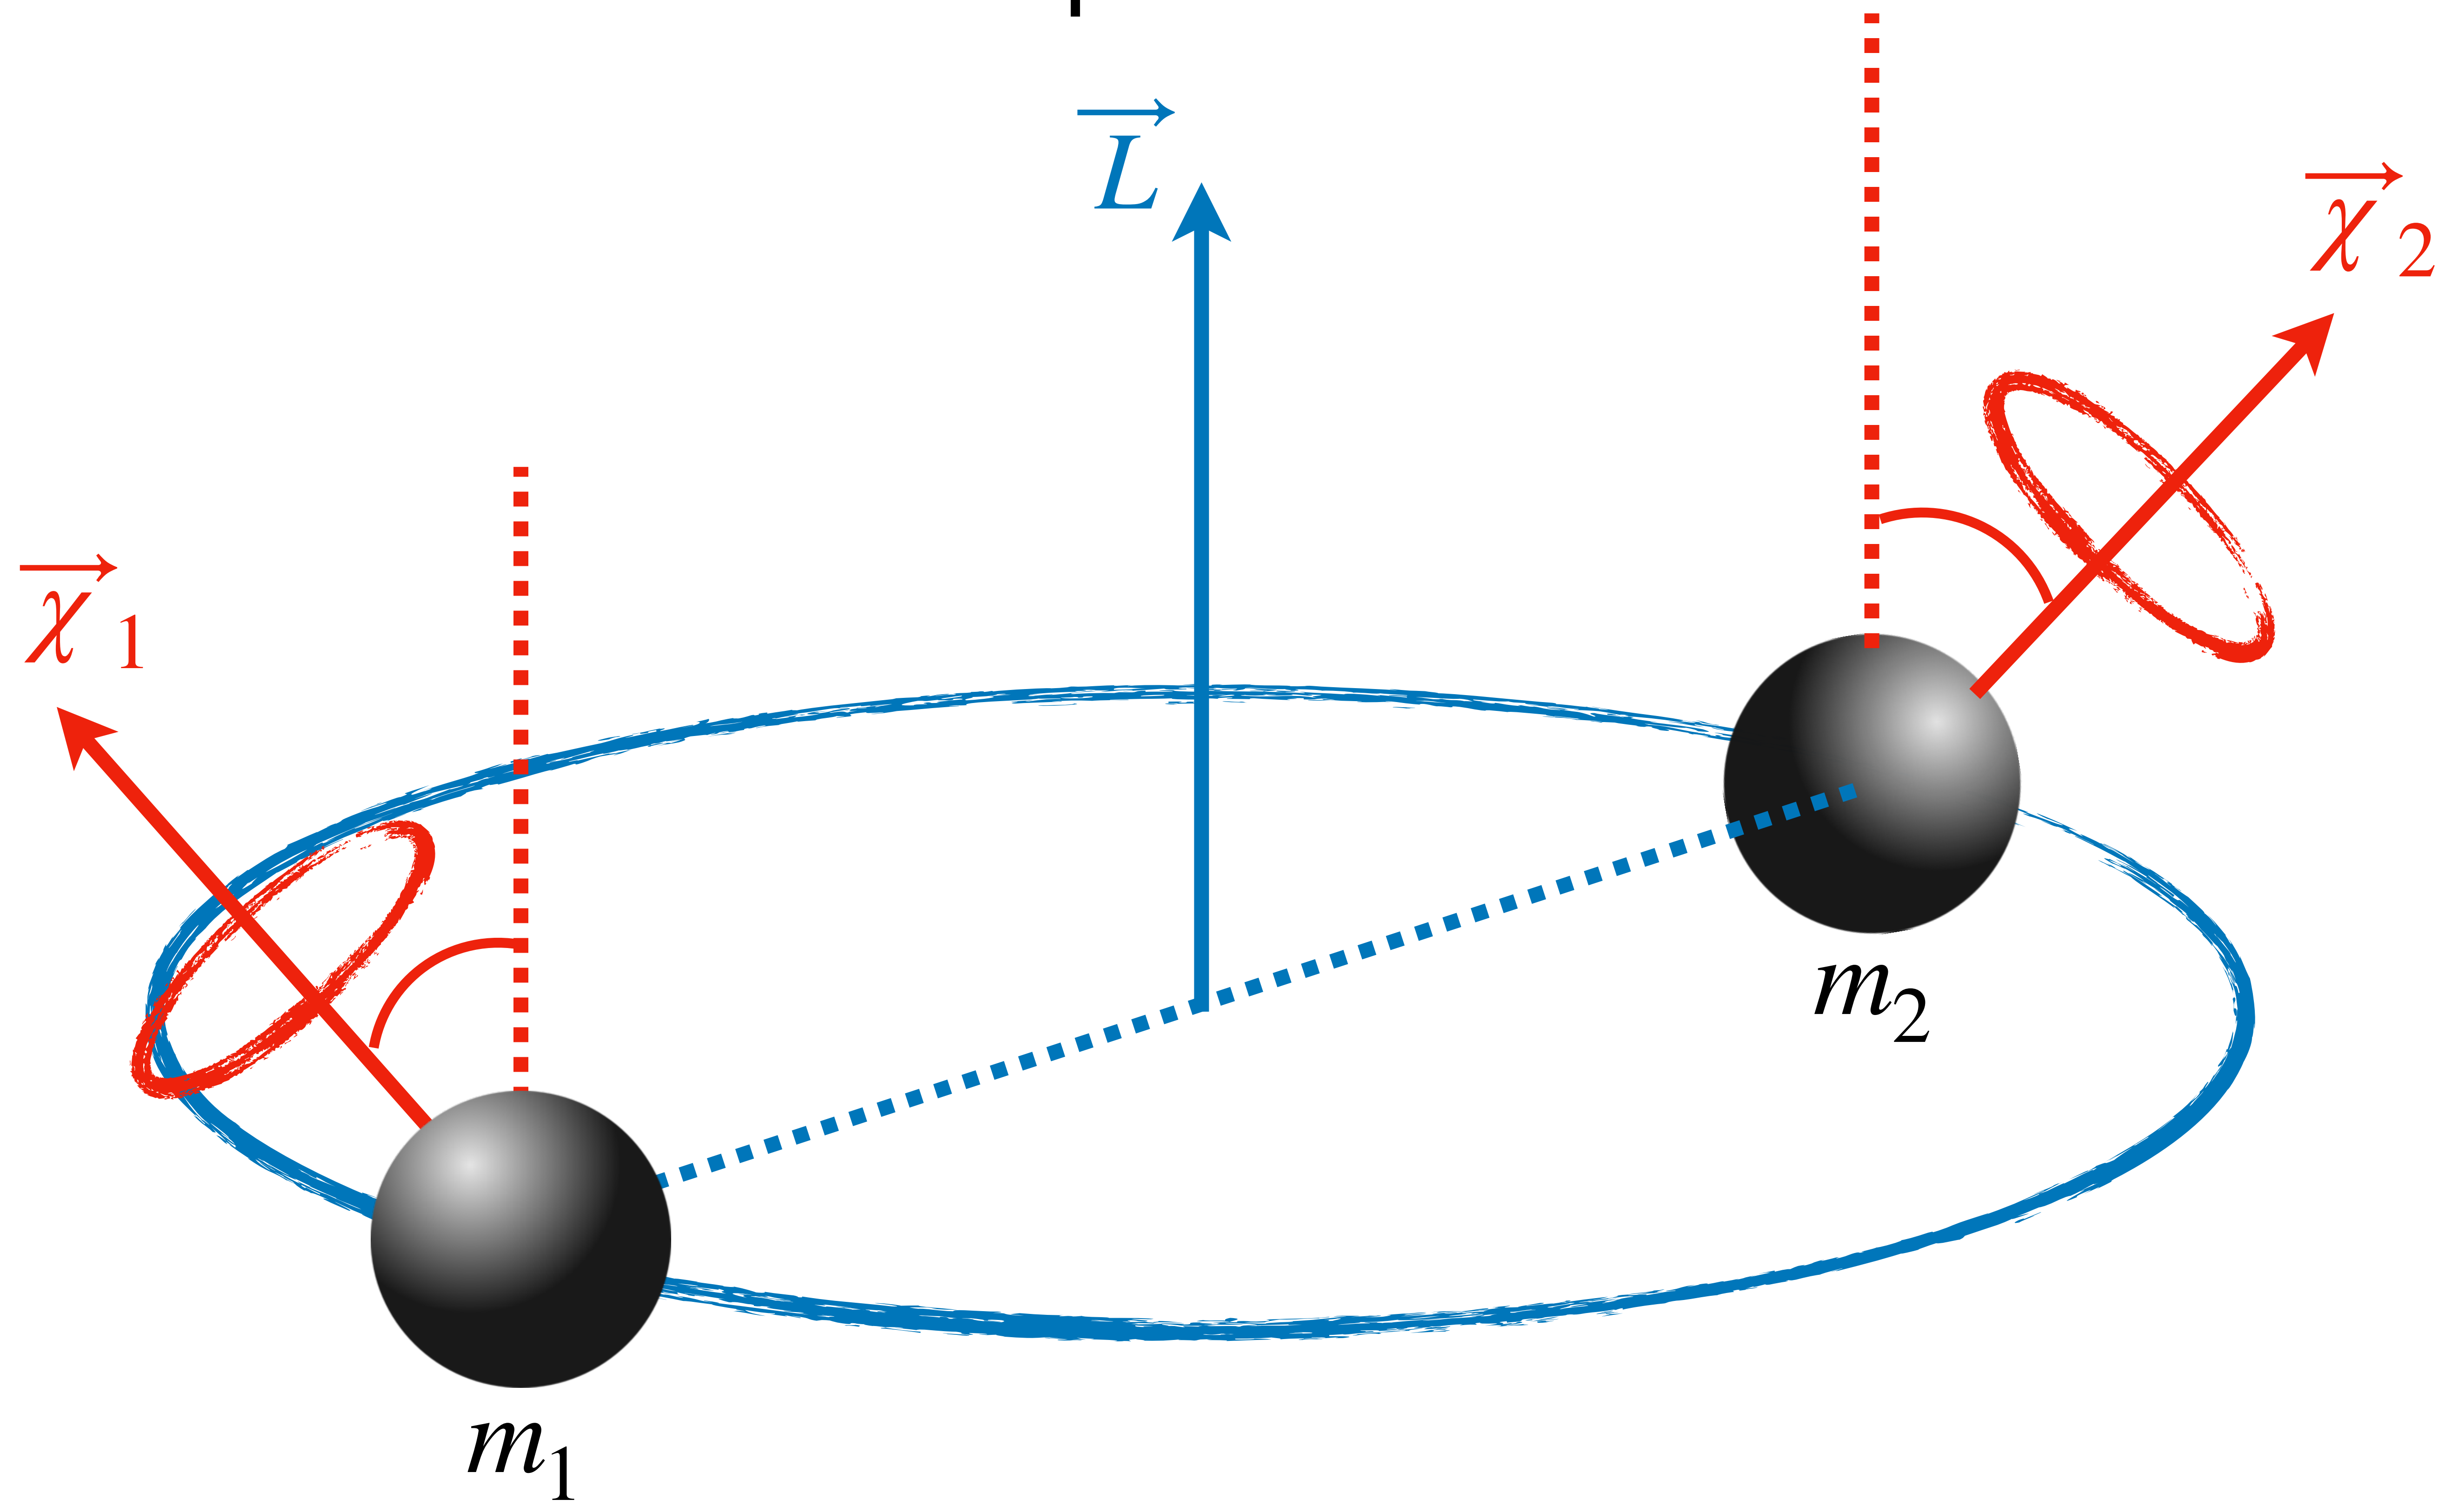

"Chirp mass" sets leading-order frequency evolution:

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}$$

Mass ratio, spins have higher-order effects

Parameters

- 7 *extrinsic* parameters, describing location/orientation of the system relative to the detector

Inclination θ_{JN}

Polarization angle ψ

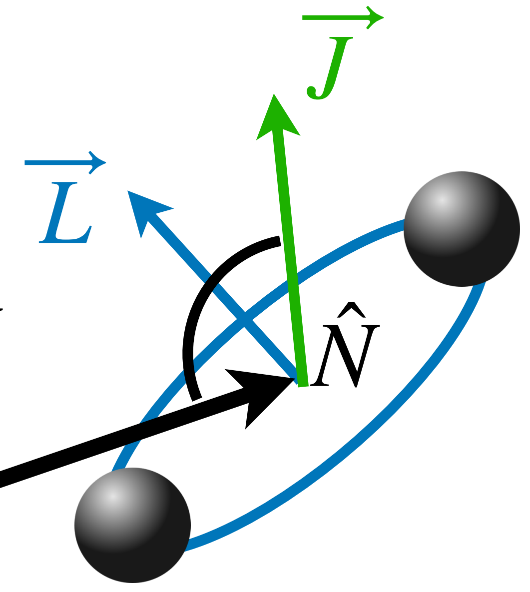

Source

Sky position (α, δ)

Time t_c and phase ϕ_c of coalescence

Detector

Distance D_L

Observed strain:

$$h(t) = F_+(\alpha, \delta, \psi)h_+(t) + F_\times(\alpha, \delta, \psi)h_\times(t)$$

Antenna patterns

Parameters

- Have a waveform model for the signal parameterized by

Intrinsic

Masses: m_1, m_2

Spins: $\vec{\chi}_1, \vec{\chi}_2$

Tidal (for NSs): Λ_1, Λ_2

Extrinsic

Luminosity distance: D_L

Inclination: θ_{JN}

Sky position: (α, δ)

Polarization angle: ψ

Coalescence time: t_{c}

Coalescence phase: ϕ_{c}

- Bayesian statistics gives us a set of tools to *infer* these parameters and their uncertainties from the data

Bayesian inference

- Bayes' theorem is a statement about *conditional probabilities*

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

- A and B are statements, e.g. "it will rain tomorrow" or "the total mass of GW190521 is 160 $M_\odot$"

Bayesian inference

- In the context of GW parameter estimation, can write Bayes' theorem like this

Goal is to determine this

$$\underbrace{P(\theta|d, h)}_{\text{Posterior}} = \frac{\overbrace{\mathcal{L}(d|\theta, h)}^{\text{Likelihood}}\overbrace{\pi(\theta|h)}^{\text{Prior}}}{\underbrace{\mathcal{Z}(d|h)}_{\text{Evidence}}}$$

Model parameters

Calibrated strain data

Model/hypothesis

Likelihood function

$$P(\theta|d,h) = \frac{\boxed{\mathcal{L}(d|\theta,h)}\pi(\theta|h)}{\mathcal{Z}(d|h)}$$

- The probability of observing the data d given a waveform model h with parameters θ
- Usually approximate the noise as being stationary and Gaussian

Tildes denote Fourier transforms

$$\log \mathcal{L}(d|\theta,h) \propto -\sum_k \frac{2|\tilde{d}_k - \tilde{h}_k(\theta)|^2}{S_\mathrm{n}(f_k)T}$$

Sum over frequency bins f_k

Noise power spectral density (PSD)

Data segment duration

Priors

$$P(\theta|d,h) = \frac{\mathcal{L}(d|\theta,h)\boxed{\pi(\theta|h)}}{\mathcal{Z}(d|h)}$$

- Prior distributions represent our *a priori* (initial) assumptions about the parameter values
- Often pick uniform/isotropic distributions
  - e.g., uniform in component masses, isotropic in spin angles, etc.
- ... or incorporate previous measurements

Evidence

$$P(\theta|d,h) = \frac{\mathcal{L}(d|\theta,h)\pi(\theta|h)}{\boxed{\mathcal{Z}(d|h)}}$$

- Normalizes the posterior distribution

$$\mathcal{Z}(d|h) = \int \mathcal{L}(d|\theta,h)\pi(\theta|h)\,\mathrm{d}\theta$$

- Can construct Bayes' factors to compare evidence for competing models

$$\mathcal{B} = \frac{\mathcal{Z}(d|h_1)}{\mathcal{Z}(d|h_2)}$$

- Larger Bayes' factor means hypothesis h_1 is favoured by the data over h_2

Marginalized posterior

$$\boxed{P(\theta|d,h)} = \frac{\mathcal{L}(d|\theta,h)\pi(\theta|h)}{\mathcal{Z}(d|h)}$$

- The posterior distribution is multi-dimensional, but can recover a 1D posterior for a single parameter, θ_1, by marginalizing over every *other* parameter

$$P(\theta_1|d,h) = \int P(\theta|d,h)\,\mathrm{d}\theta_2 \ldots \mathrm{d}\theta_n$$

Obtaining the posterior

$$\boxed{P(\theta|d,h)} = \frac{\mathcal{L}(d|\theta,h)\pi(\theta|h)}{\mathcal{Z}(d|h)}$$

- Posterior distribution is given by Bayes' theorem, so we can just calculate it directly, right?
- Not so fast — remember that we have ~15-17 free parameters
- Imagine we are evaluating the likelihood over a coarse grid in parameter space, using just 10 values for each parameter, 1ms waveform generation

$$10^{15} \text{ points} \times 10^{-3} \text{ seconds/point} \approx 30{,}000 \text{ years!}$$

- High dimensionality renders brute-force calculation impractical, need to try something else

Stochastic sampling

- Instead we use a stochastic sampler to infer the posterior distribution, e.g. Markov chain Monte Carlo (MCMC) or nested sampling
- Several sampling algorithms/implementations publicly available, e.g. dynesty
- Get results typically on scale of hours to days, depending on setup/processing power

[2205.15570]

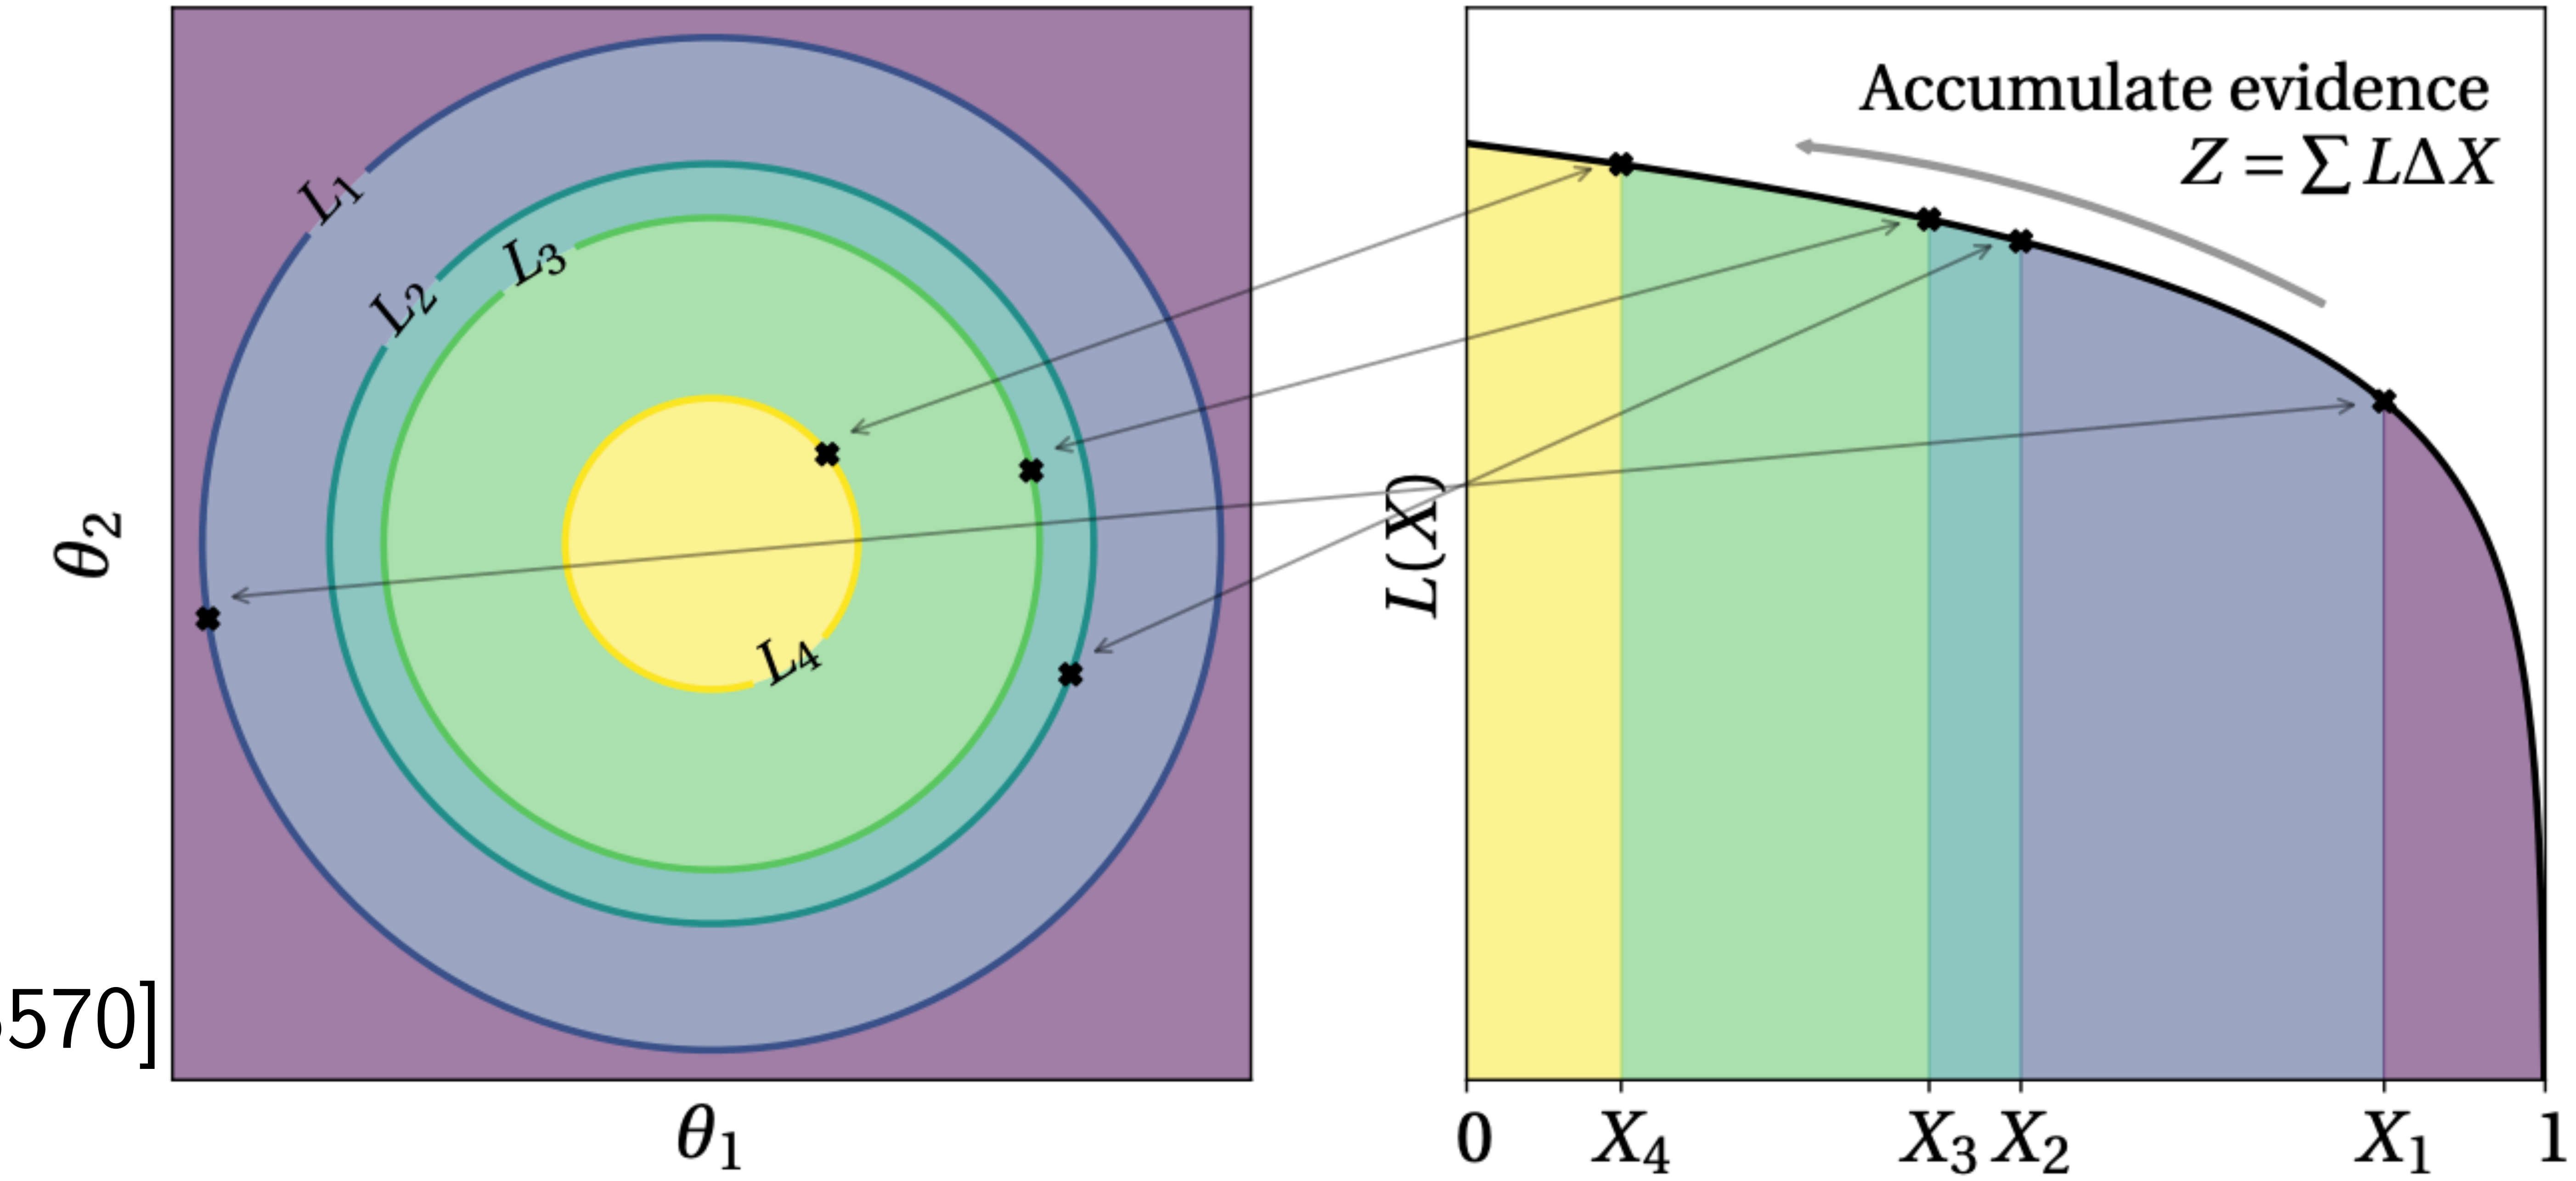

Bilby

- Bilby* is a publicly available Python package for Bayesian parameter estimation
- Designed for GW analyses but can also tackle more general applications
- Analyze strain data, inject simulated signals, generate random noise
- Main workhorse for PE analyses within the LVK
- See papers: [1811.02042, 2006.00714]
- Next: using Bilby to analyze a real event

*A "bilby" is a small marsupial found in Australia, where Bilby was originally developed

Visualizing the output

- This kind of diagram is called a "corner plot"
- Histograms show the 1D marginalized distributions for chirp mass and mass ratio
- Contour plot is the joint posterior distribution for both parameters
- Obtaining credible intervals a matter of computing percentiles of the marginal posteriors

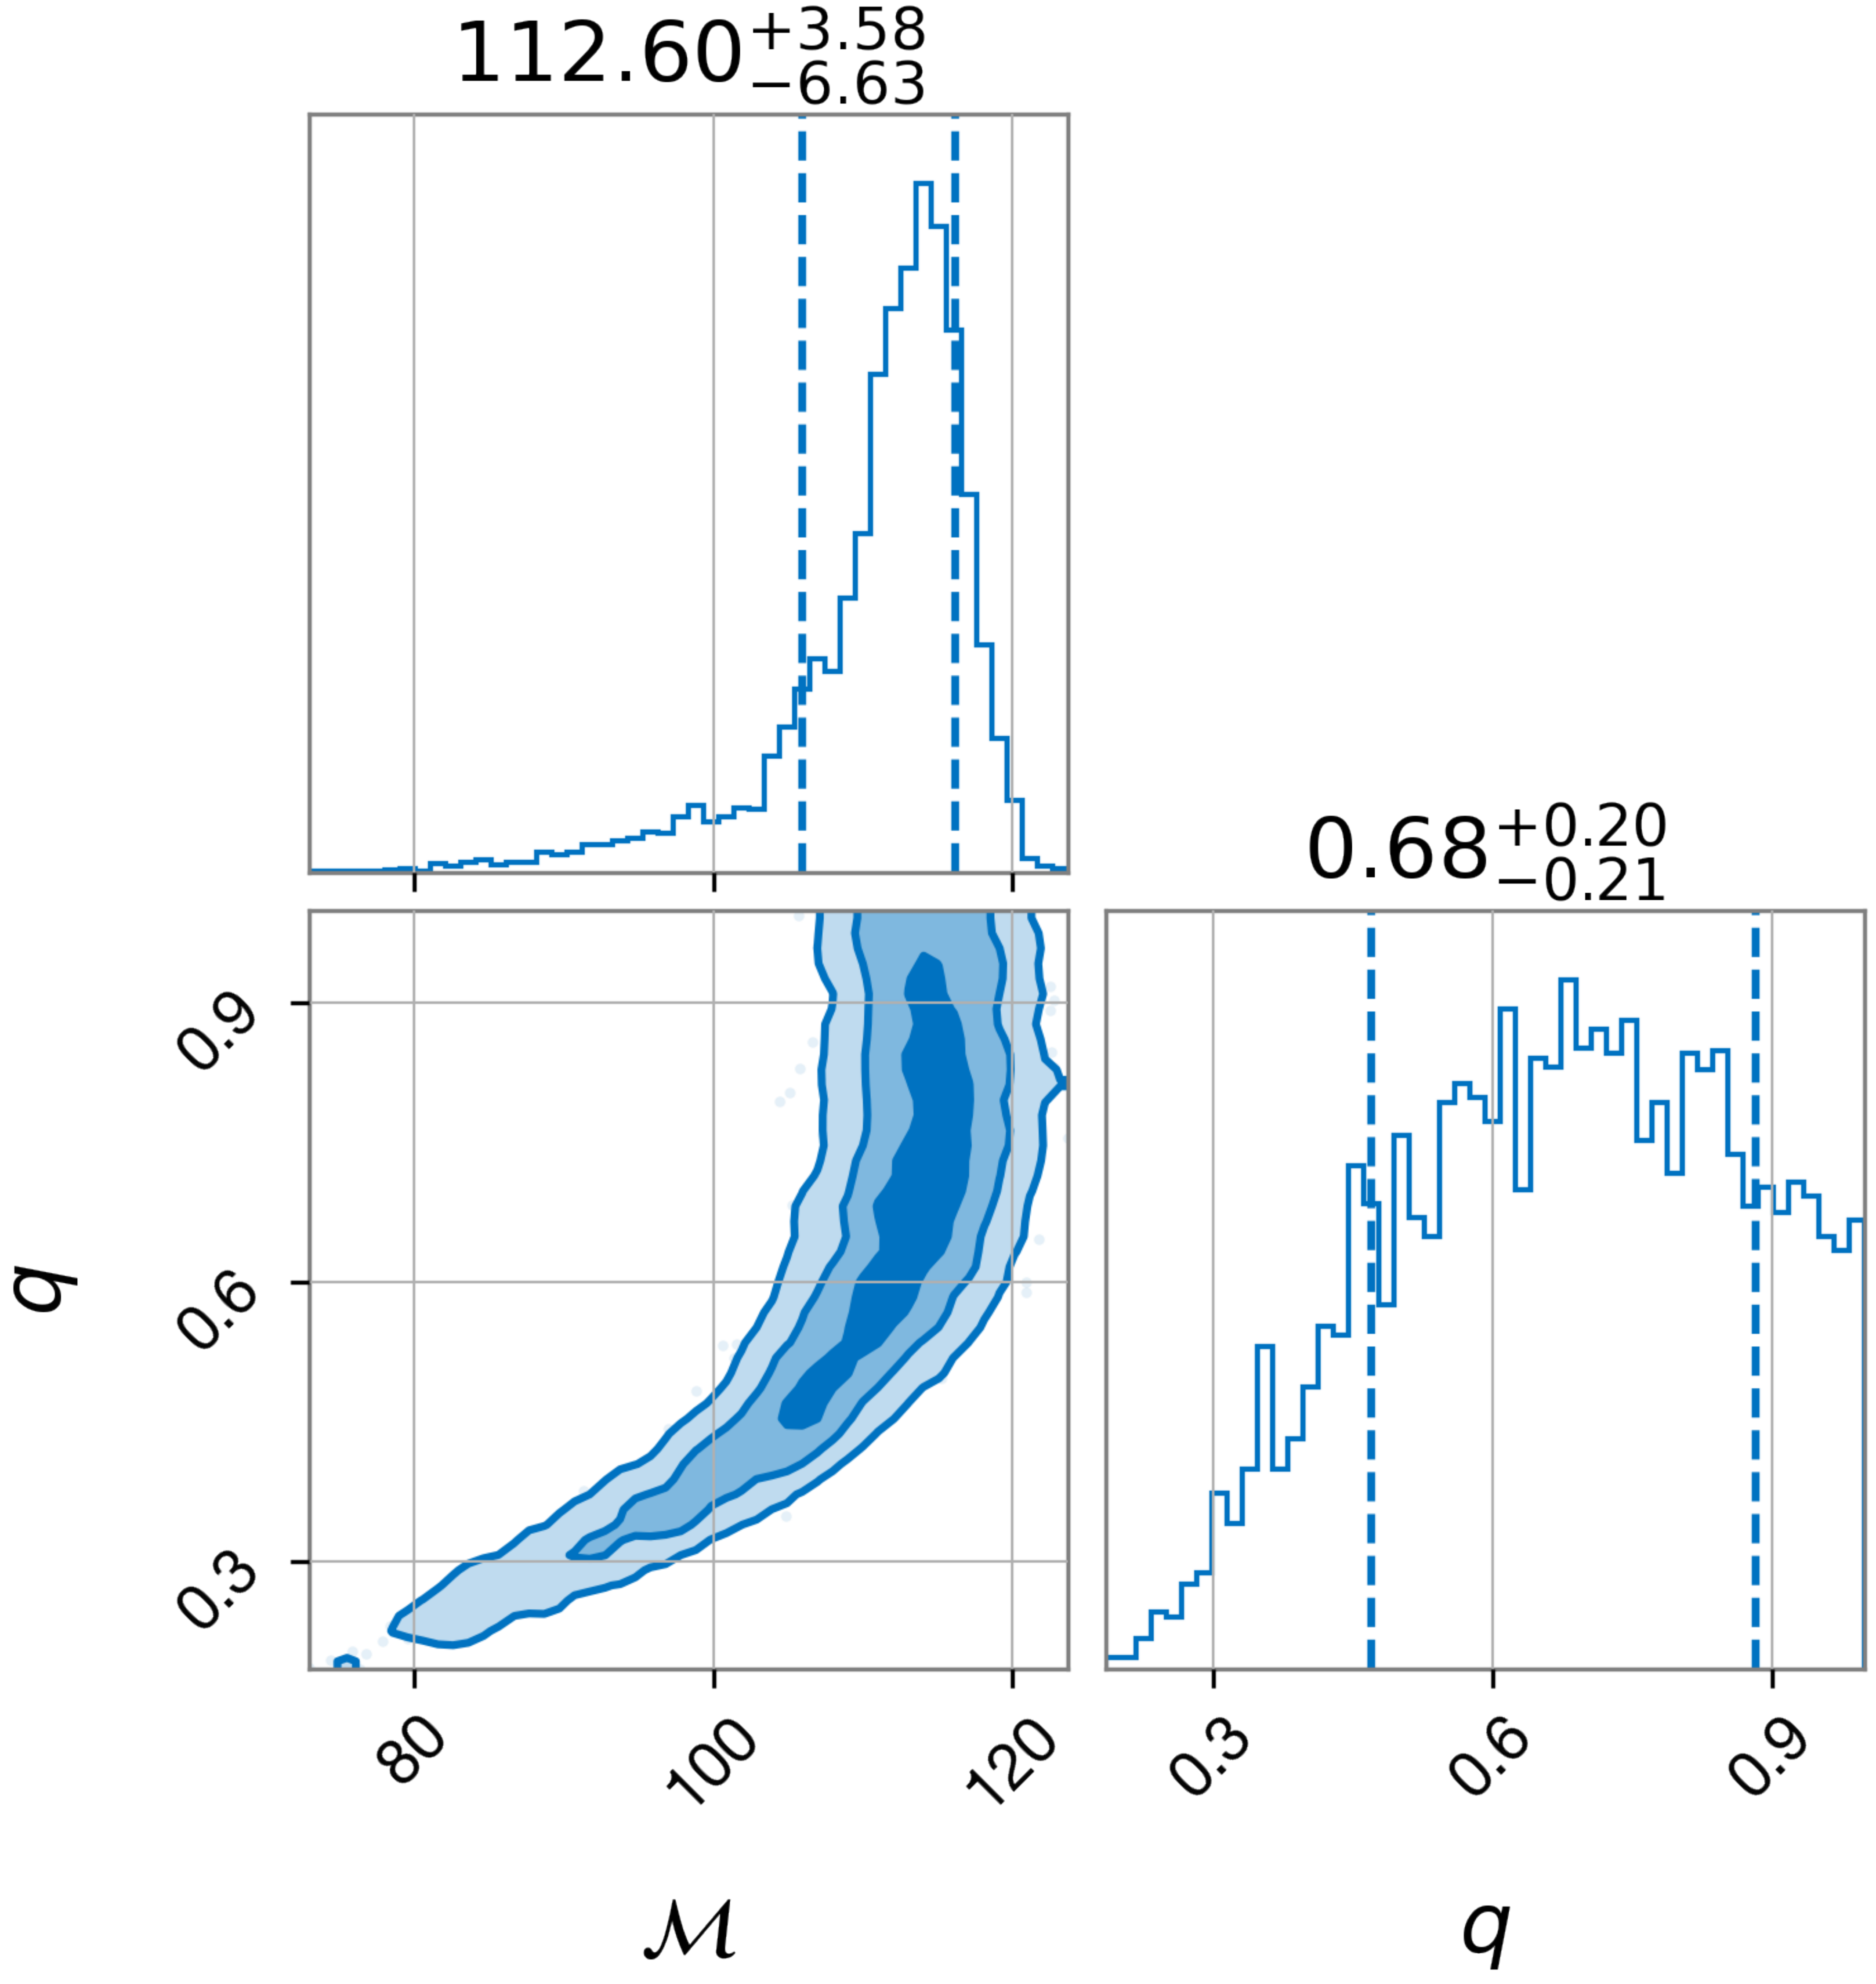

Visualizing the output

- Reconstructed waveform plotted on top of whitened and bandpassed strain in H1, L1

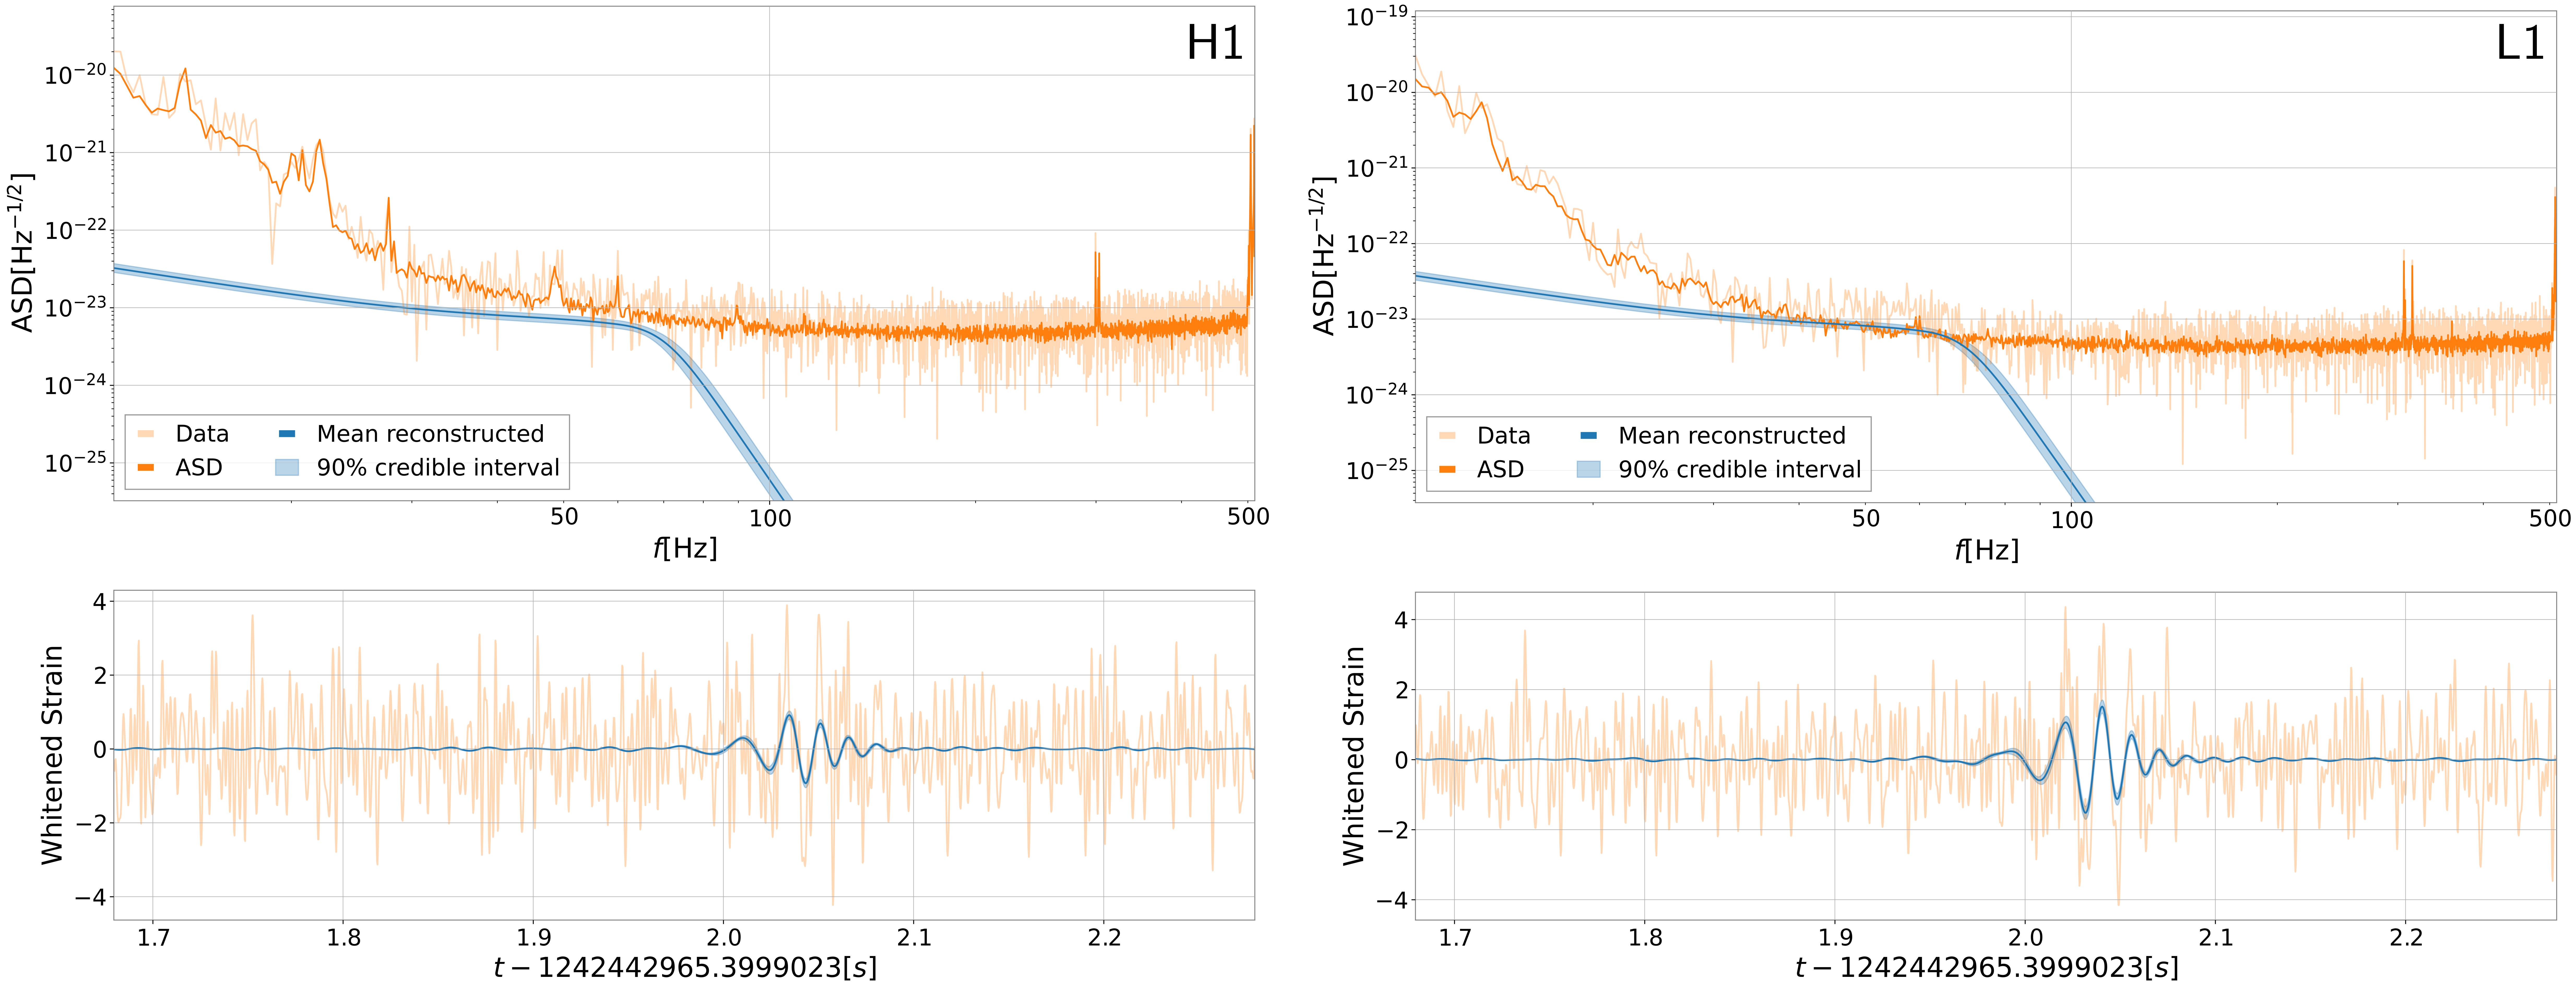

Resources

- Link to tutorial notebook: https://colab.research.google.com/github/alanknee/GWANW_PE_Tutorial/blob/main/gwanw22_bilby_tutorial.ipynb
- Bilby documentation: https://lscsoft.docs.ligo.org/bilby/
- More examples: https://lscsoft.docs.ligo.org/bilby/examples.html