Statistically principled learning for gravitational-wave inverse problems

Alvin Chua
NUS

AISSAI workshop
Toulouse, France
1 October 2024

National University of Singapore

Inverse problems in GW science

- Forward problem: Physical parameters ↦ expected data
  - Waveform models: Detector strain response to source GW signal
  - Population models: Source distribution, aggregate GW signal
  - Noise models: PSD (stochastic), phenomenology (transients)

$$\theta \mapsto h(\theta)$$

Inverse problems in GW science

- Forward problem: Physical parameters ↦ expected data
  - Waveform models: Detector strain response to source GW signal
  - Population models: Source distribution, aggregate GW signal
  - Noise models: PSD (stochastic), phenomenology (transients)
- Inverse problem: Detector data ↦ inferred parameters
  - Significance quantification: Presence of signal
  - Uncertainty quantification: Probability of parameters

$$\theta \mapsto h(\theta)$$

$$x \mapsto p(\theta|x)$$

Inverse problems in GW science

- Forward problem: Physical parameters ↦ expected data
  - Waveform models: Detector strain response to source GW signal
  - Population models: Source distribution, aggregate GW signal
  - Noise models: PSD (stochastic), phenomenology (transients)
- Inverse problem: Detector data ↦ inferred parameters
  - Significance quantification: Presence of signal
  - Uncertainty quantification: Probability of parameters
- Standard theoretical/computational framework
  - Matched filtering
  - Bayesian inference

$$\theta \mapsto h(\theta)$$

$$x \mapsto p(\theta|x)$$

$$\langle \cdot | \cdot \rangle \qquad \ln p(x|\theta) = -\frac{1}{2}\langle x - h(\theta) | x - h(\theta) \rangle + c$$

Where does ML fit in?

- Everywhere! But...
  - Should clarify how any proposed method relates to standard filtering-based Bayesian framework, which is already rigorously founded on well-understood principles
  - Easiest way to do this is by using ML to augment or emulate standard framework
  - More defensible on rigour and principle, thus more scientifically viable

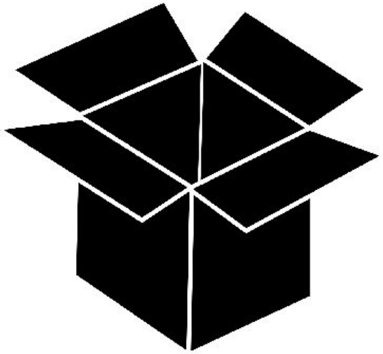

Where does ML fit in?

- Everywhere! But...
  - Should clarify how any proposed method relates to standard filtering-based Bayesian framework, which is already rigorously founded on well-understood principles
  - Easiest way to do this is by using ML to augment or emulate standard framework
  - More defensible on rigour and principle, thus more scientifically viable
- The GW problem poses distinctive challenges for ML too
  - Sensitivity: Data dominated by “noise” (instrumental and astrophysical)
  - Generalisability: Data space can have large information volume
  - Precision: Precise models used to make precise inferences

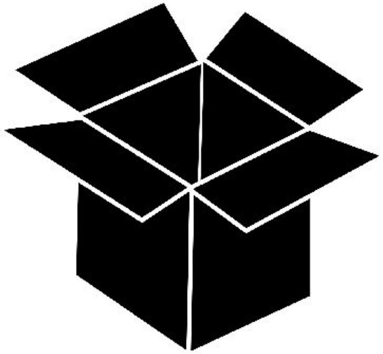

Where does ML fit in?

- Everywhere! But...
  - Should clarify how any proposed method relates to standard filtering-based Bayesian framework, which is already rigorously founded on well-understood principles
  - Easiest way to do this is by using ML to augment or emulate standard framework
  - More defensible on rigour and principle, thus more scientifically viable
- The GW problem poses distinctive challenges for ML too
  - Sensitivity: Data dominated by “noise” (instrumental and astrophysical)
  - Generalisability: Data space can have large information volume
  - Precision: Precise models used to make precise inferences
- I’ll discuss two classes of task in the inverse problem
  - Forward-model fitting
  - Posterior estimation

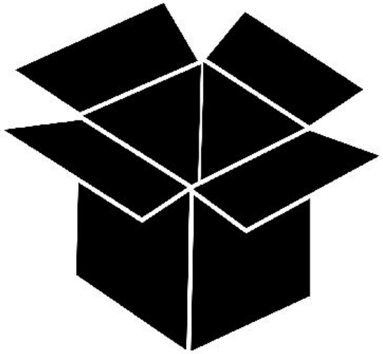

Forward-model fitting

- Let's focus on generic waveform models for a single class of (binary) source
  - Population distribution fitting [Natalia's talk]
  - Stochastic signals from populations [Riccardo's talk]
- Waveforms are often represented as long time/frequency series
  - For LISA, the dimensionality of this representation can be millions

$$\theta \mapsto h(\theta)$$

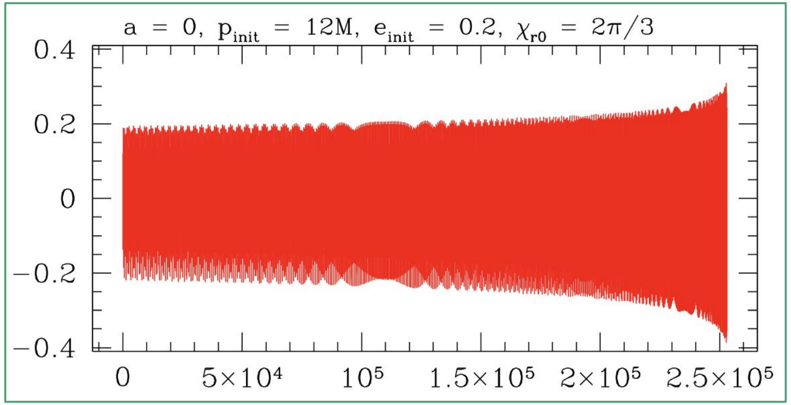

Hughes et al. (+AC) 2021

Forward-model fitting

- Let's focus on generic waveform models for a single class of (binary) source
  - Population distribution fitting [Natalia's talk]
  - Stochastic signals from populations [Riccardo's talk]

$$\theta \mapsto h(\theta)$$

- Waveforms are often represented as long time/frequency series
  - For LISA, the dimensionality of this representation can be millions
- Can we learn this directly?
  - More precisely: We want to fit model output M to waveforms by minimising L2 distance

$$\text{loss} = \sum_i |h_i - M(\theta_i)|^2$$

- Sure, maybe with the help of dimensionality reduction

Forward-model fitting

- First attempted using reduced-order modelling and neural networks
  - AC, Galley & Vallisneri 2019; Khan & Green 2021
  - Reasonably fast and accurate
  - Another nice perk of using neural networks: Analytical waveform derivatives for free

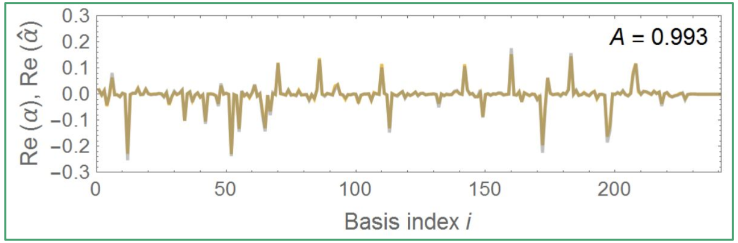

AC, Galley & Vallisneri 2019

Forward-model fitting

- First attempted using reduced-order modelling and neural networks
    - AC, Galley & Vallisneri 2019; Khan & Green 2021
    - Reasonably fast and accurate
    - Another nice perk of using neural networks: Analytical waveform derivatives for free

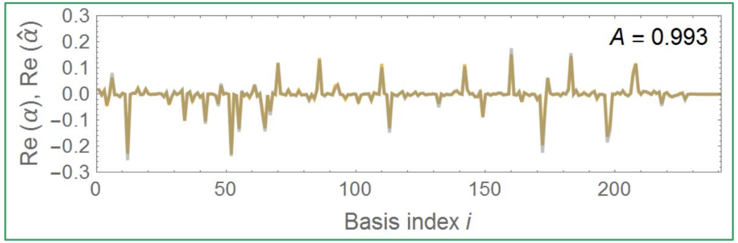

AC, Galley & Vallisneri 2019

- On the flip side: Difficult to achieve greater precision in practice
    - Limits from network architecture/capacity, or machine-precision constraints
- One thing that doesn't help: Output features are completely arbitrary
    - Dimensionality reduction is data-driven, does not leverage knowledge of underlying physics

Forward-model fitting

- Feature engineering (on both input and output representations)
  - Essentially, physically motivated alternative parametrisations or intermediate quantities
- Input example: Chirp-time coordinates
  - Preserve the information metric, can facilitate training-set placement
- Output example: Mode decompositions
  - Fit mode amplitudes instead of full waveform [AC et al. 2021; Katz et al. (+AC) 2021]

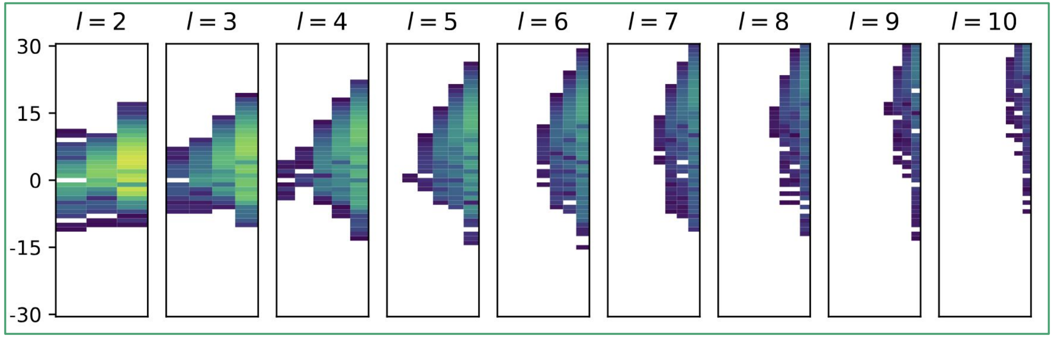

Katz et al. (+AC) 2021

Forward-model fitting

- Feature engineering (on both input and output representations)
  - Essentially, physically motivated alternative parametrisations or intermediate quantities
- Input example: Chirp-time coordinates
  - Preserve the information metric, can facilitate training-set placement
- Output example: Mode decompositions
  - Fit mode amplitudes instead of full waveform [AC et al. 2021; Katz et al. (+AC) 2021]
- Feature learning
  - Leverage emergent structure after dimensionality reduction
- Example: Spiral representations
  - Improves accuracy by orders of magnitude [Nousi et al. 2022]

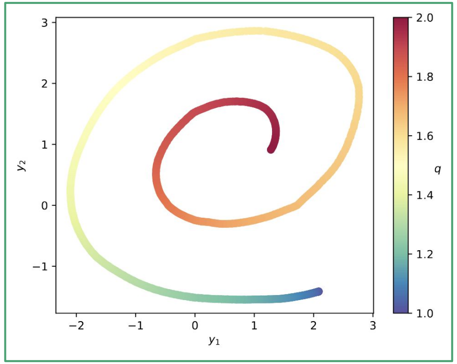

Nousi et al. 2022

Posterior estimation

- Early attempts at using neural networks for GW detection were binary classifiers with no or external (frequentist) significance quantification
  - Gebhard et al. 2017; George & Huerta 2018; Gabbard et al. 2018
  - ML might be used for significance quantification, e.g., conformal prediction [Ashton et al. 2024]

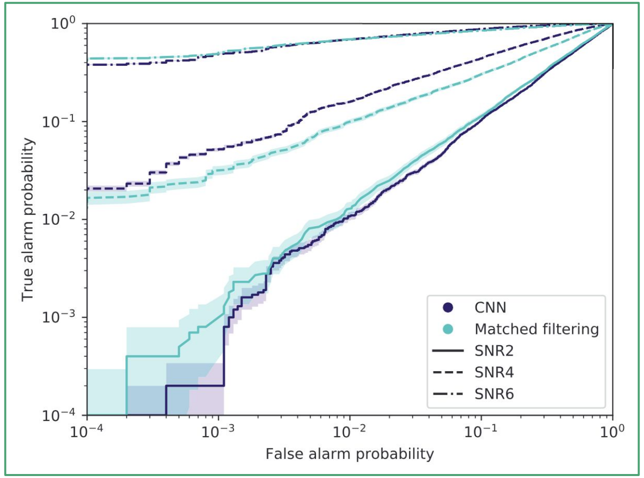

Gabbard et al. 2018

Posterior estimation

- Early attempts at using neural networks for GW detection were binary classifiers with no or external (frequentist) significance quantification
  - Gebhard et al. 2017; George & Huerta 2018; Gabbard et al. 2018
  - ML might be used for significance quantification, e.g., conformal prediction [Ashton et al. 2024]
- Let's focus on ML for GW inference instead
  - Posterior estimation for uncertainty quantification: Probabilities, expectations, credible sets
- Learning proposal distributions for MCMC
  - E.g., by fitting physical priors or posteriors with normalising flows [Natalia's talk]
  - Or more general principles, e.g., maximising entropy of proposal [Li, Chen & Sommer 2021]
- Variational inference
  - E.g., SGD variational inference with normalising flows [Vallisneri et al. 2024]

Posterior estimation

- Simulation-based (likelihood-free) inference
  - Essentially: Fit model output M to posterior by minimising conditional cross-entropy
- Completely front-load cost of inference
  - AC & Vallisneri 2020; Gabbard et al. 2022; Green, Simpson & Gair 2020

$$\text{loss} = -\sum_i \ln M(\theta_i|x_i)$$

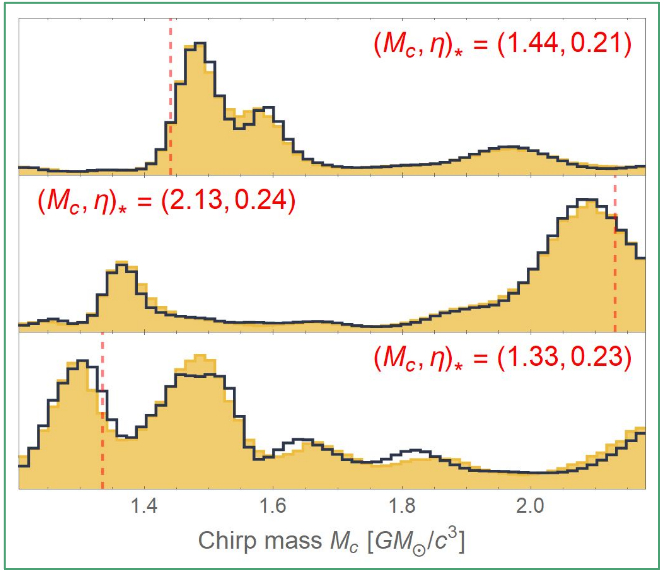

AC & Vallisneri 2020

Posterior estimation

- Simulation-based (likelihood-free) inference
  - Essentially: Fit model output M to posterior by minimising conditional cross-entropy
- Completely front-load cost of inference
  - AC & Vallisneri 2020; Gabbard et al. 2022; Green, Simpson & Gair 2020
- Potential applications
  - Fast sky localisation
  - MCMC proposal distributions
  - Large-scale exploratory studies
- But may be limited in generalisability and reliability, so not a proxy for the posterior

$$\text{loss} = -\sum_i \ln M(\theta_i | x_i)$$

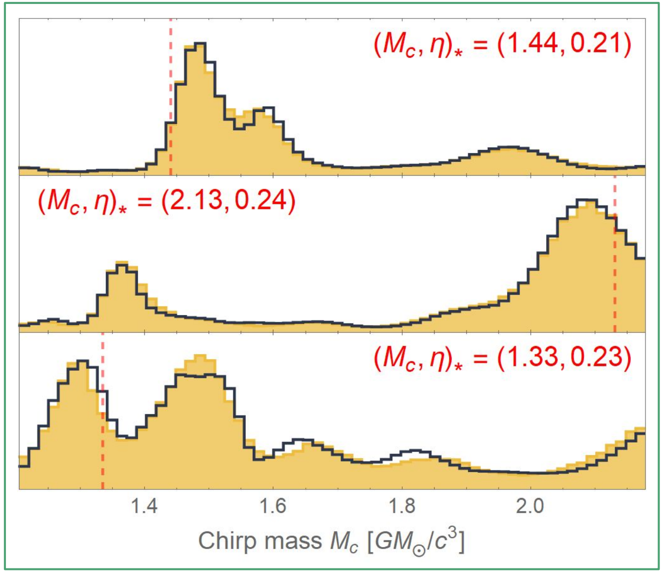

AC & Vallisneri 2020

Posterior estimation

- ML can be used to produce fast posteriors, but what about reliable ones?
  - Specifically, posteriors that are robust against model misspecification (known and unknown)
  - Dealing with known errors is an easier supervised task
- Posterior correction
  - Essentially fitting and correcting for model/likelihood error
  - If fit is a GP, can marginalise out [Moore et al. (+AC) 2016; AC et al. 2020; Liu, Li & AC 2023]

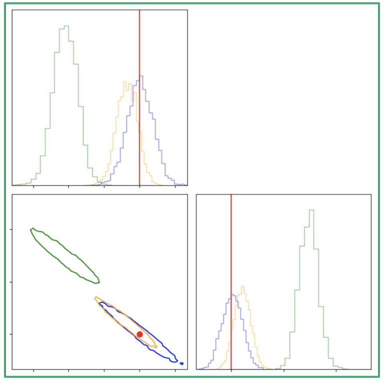

Liu, Li & AC 2023

Posterior estimation

- ML can be used to produce fast posteriors, but what about reliable ones?
  - Specifically, posteriors that are robust against model misspecification (known and unknown)
  - Dealing with known errors is an easier supervised task
- Posterior correction
  - Essentially fitting and correcting for model/likelihood error
  - If fit is a GP, can marginalise out [Moore et al. (+AC) 2016; AC et al. 2020; Liu, Li & AC 2023]
- Posterior calibration
  - Learn credible level of erroneous posterior that gives specified coverage [Mao et al. (+AC) 2024]
- Generalisability is still an issue with such methods though

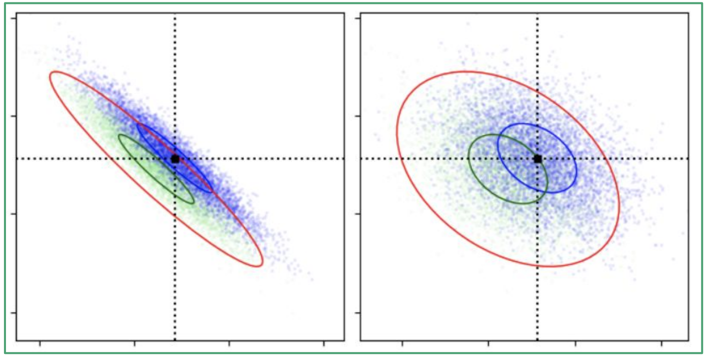

Mao et al. (+AC) 2024

Summary

- ML is potentially useful for GW inverse problems, but should always have a clear relationship with the standard filtering-based Bayesian framework
- ML can be used for various fitting tasks in forward modelling, and there is plenty of room to improve its efficacy using feature engineering/learning
- ML can be used in various ways to augment or emulate the task of Bayesian posterior estimation, and to make posteriors robust against model error

References (incomplete, sorry)

- R. Mao et al., Calibrating approximate Bayesian credible intervals of gravitational-wave parameters, Phys. Rev. D 109, 083002 (2024).
- M. Liu, X.-D. Li & A. J. K. Chua, Improving the scalability of Gaussian-process error marginalization in gravitational-wave inference, Phys. Rev. D 108, 103027 (2023).
- M. L. Katz et al., Fast extreme-mass-ratio-inspiral waveforms: New tools for millihertz gravitational-wave data analysis, Phys. Rev. D 104, 064047 (2021).
- S. A. Hughes et al., Adiabatic waveforms for extreme mass-ratio inspirals via multivoice decomposition in time and frequency, Phys. Rev. D 103, 104014 (2021).
- A. J. K. Chua et al., Rapid generation of fully relativistic extreme-mass-ratio-inspiral waveform templates for LISA data analysis, Phys. Rev. Lett. 126, 051102 (2021).
- A. J. K. Chua et al., Gaussian processes for the interpolation and marginalization of waveform error in extreme-mass-ratio-inspiral parameter estimation, Phys. Rev. D 101, 044027 (2020).
- A. J. K. Chua & M. Vallisneri, Learning Bayesian posteriors with neural networks for gravitational-wave inference, Phys. Rev. Lett. 124, 041102 (2020).
- A. J. K. Chua, C. R. Galley & M. Vallisneri, Reduced-order modeling with artificial neurons for gravitational-wave inference, Phys. Rev. Lett. 122, 211101 (2019).
- C. J. Moore et al., Improving gravitational-wave parameter estimation using Gaussian process regression, Phys. Rev. D 93, 064001 (2016).