Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy

Hunter Gabbard[1], Chris Messenger[1], Ik Siong Heng[1], Francesco Tonolini[2], Roderick Murray-Smith[2]

SUPA School of Physics and Astronomy, University of Glasgow, United Kingdom [1]
School of Computing Science, University of Glasgow, United Kingdom [2]

Introduction

Gravitational wave (GW) detection is now commonplace [1] and as the sensitivity of GW detectors improves, we will observe $\mathcal{O}(100)$s of transient GW events per year [2]. The current methods used to estimate their source parameters employ optimally sensitive [3] but computationally costly Bayesian inference approaches [4] where typical analyses have taken between 6 hours and 5 days [5].
We show for the first time that a conditional variational autoencoder (CVAE) [6, 7] pre-trained on binary black hole (BBH) signals can return Bayesian posterior probability estimates ~ 6 orders of magnitude faster than existing techniques.

Methods Overview

1. Generate 10^7 training BBH waveforms in Gaussian noise.
2. Generate 256 test samples with posteriors produced by 4 benchmark Bayesian sampler approaches (`Dynesty`, `CPNest`, `Ptemcee` and `Emcee`).
3. Train model (`VItamin`) to infer posteriors for 7 parameters describing a BBH merger.
4. Hyperparameters chosen through Bayesian optimisation using Gaussian Processes.
5. Produce 3 figures of merit:
   - Corner plots for each test sample superimposing results from `VItamin`, `Dynesty` and `Ptemcee`.
   - P-P plot showing how consistent each approach is with the truth.
   - KL divergence plot showing how similar posteriors are between different approaches (values closer to zero indicate high degree of overlap).

Our Model

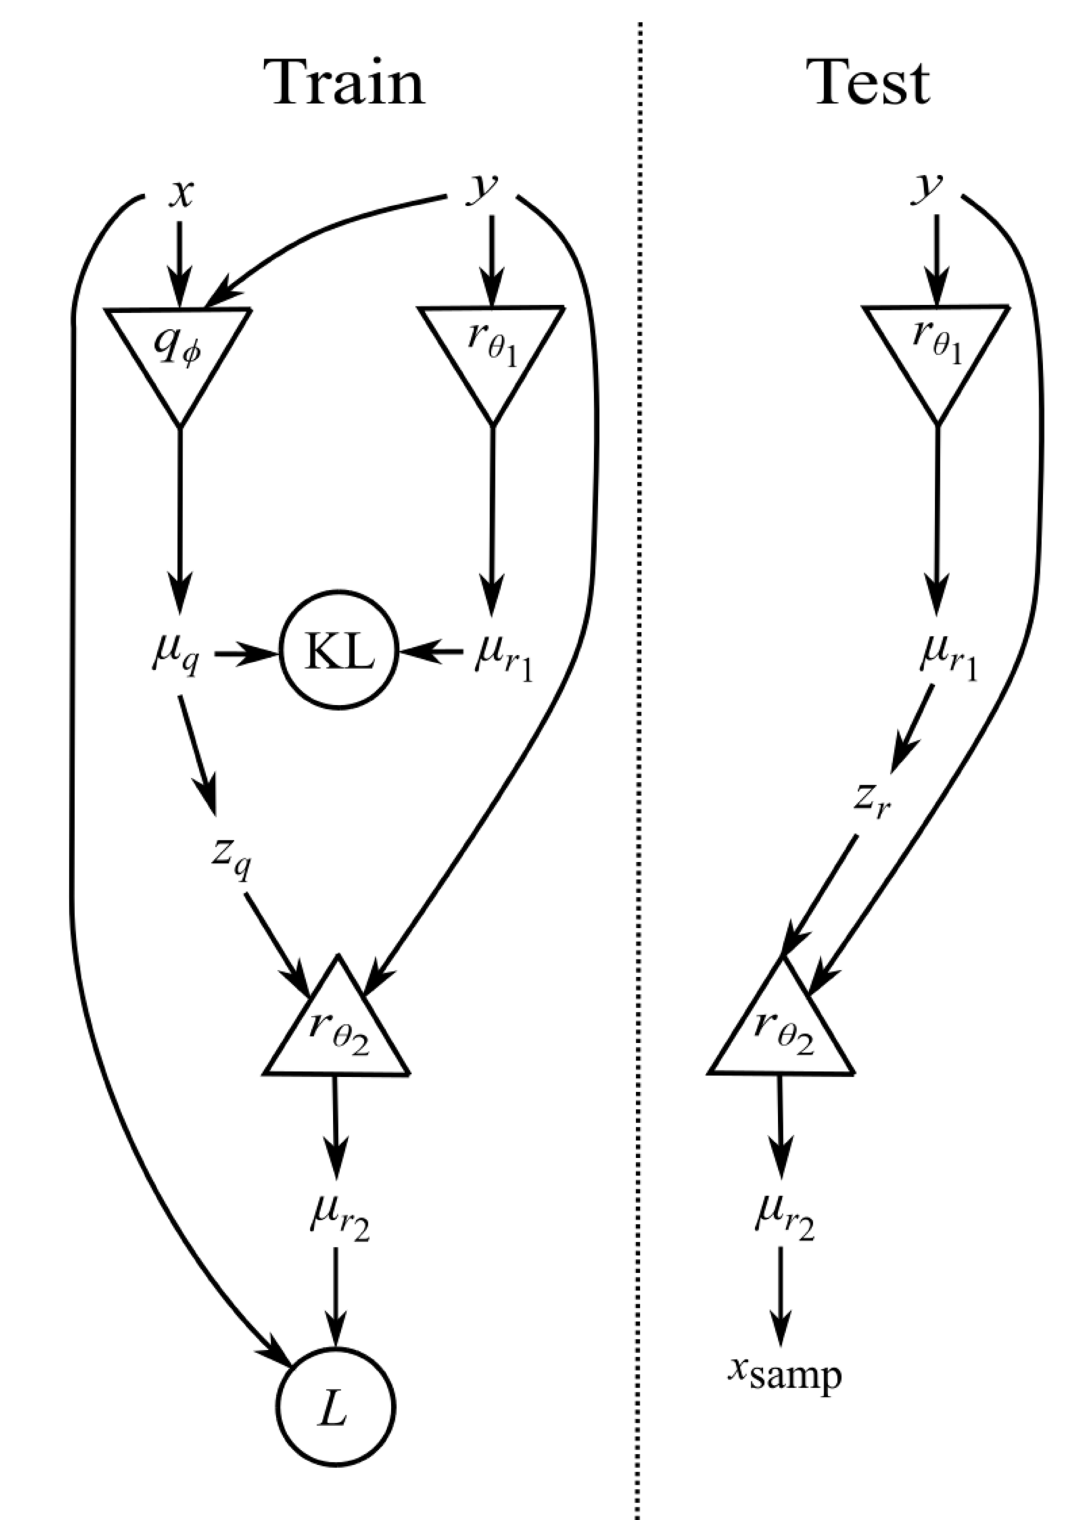

Figure 1: The configuration of the CVAE neural network. During training (left-hand side), a training set of noisy GW signals (y) and their corresponding true parameters (x) are given as input to encoder network q_ϕ, while only y is given to encoder network r_{θ_1}. The KL-divergence is computed between the encoder output latent space representations (μ_q and μ_r) forming one component of the total cost function. Samples (z_q) from the q_ϕ latent space representation are generated and passed to the decoder network r_{θ_2} together with the original input data y. The output of the decoder (μ_x) describes a distribution in the physical parameter space and the cost component L is computed by evaluating that distribution at the location of the original input x.

Results

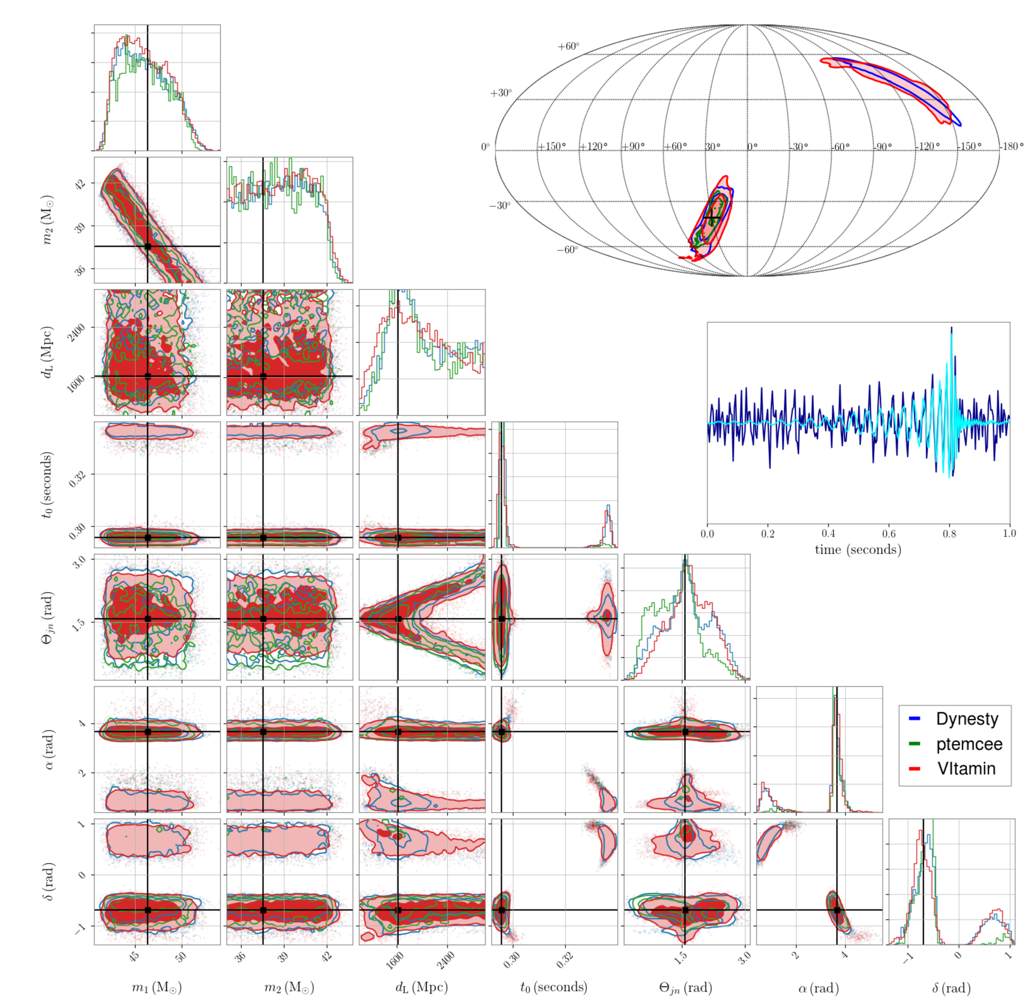

Figure 2: Corner plot showing one and two-dimensional marginalised posterior distributions on the GW parameters for one example test dataset using `VItamin` (red), `Dynesty` (blue) and `Ptemcee` (green).

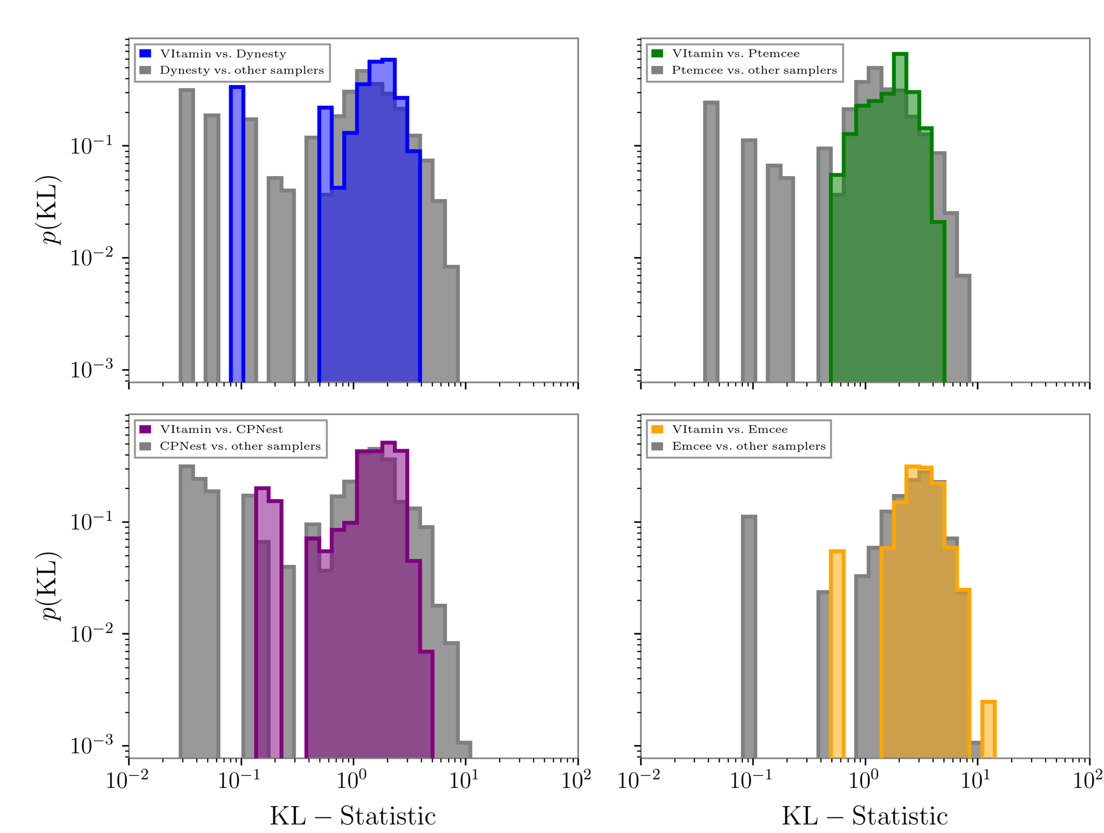

Figure 3: Distributions of KL-divergence values between posteriors produced by different samplers. We show the distribution of KL-divergences computed between a single benchmark sampler and every other benchmark sampler over all 256 GW test cases (grey) and KL-divergence distributions between the single benchmark sampler and the `VItamin` outputs (blue, green, purple, yellow).

Discussion

We have demonstrated that we are able to reproduce, to a high degree of accuracy, Bayesian posterior probability distributions generated through machine learning with the same quality of results as trusted benchmark analyses used within the LIGO-Virgo Collaboration.
The significance of our results is most evident in the orders of magnitude increase in speed over existing algorithms. Given that the predicted number of future detections of binary neutron star (BNS) mergers (~ 180 [2]) will severely strain the GW community's current computational resources using existing Bayesian methods, it is imperative that some form of fast posterior generation be employed.
Given the abundant benefits of this method, we hope that a variant of this of approach will form the basis for future GW parameter estimation.

Contact Information

- Email: h.gabbard.1@research.gla.ac.uk

References

[1] B. P. Abbott *et al.*, "Gw170817: Observation of gravitational waves from a binary neutron star inspiral," *Phys. Rev. Lett.*, vol. 119, p. 161101, Oct 2017.

[2] B. P. Abbott *et al.*, "Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA," *Living Reviews in Relativity*, vol. 21, p. 3, Apr 2018.

[3] A. C. Searle, P. J. Sutton, and M. Tinto, "Bayesian detection of unmodeled bursts of gravitational waves," *Classical and Quantum Gravity*, vol. 26, p. 155017, Aug 2009.

[4] J. Veitch, V. Raymond, B. Farr, W. M. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, W. D. Pozzo, F. Feroz, J. Gair, C.-J. Haster, V. Kalogera, T. Littenberg, I. Mandel, R. O'Shaughnessy, M. Pitkin, C. Rodriguez, C. Röver, T. Sidery, R. Smith, M. V. D. Sluys, A. Vecchio, W. Vousden, and L. Wade, "Robust parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library," *Physical Review D*, 2014.

[5] "Gracedb — gravitational-wave candidate event database (ligo/virgo o3 public alerts)." `https://gracedb.ligo.org/superevents/public/O3/`. Accessed: 2019-09-16.

[6] F. Tonolini, A. Lyons, P. Caramazza, D. Faccio, and R. Murray-Smith, "Variational inference for computational imaging inverse problems," 2019. To appear in JMLR.

[7] A. Pagnoni, K. Liu, and S. Li, "Conditional variational autoencoder for neural machine translation," 2018.

University of Glasgow
LIGO
IGR