

FAST AND ACCURATE GRAVITATIONAL-WAVE MODELLING WITH PRINCIPAL COMPONENT REGRESSION

Cyril Cano¹ Éric Chassande-Mottin² Nicolas Le Bihan¹

¹ Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-Lab, F-38000 Grenoble, France

² Université de Paris, CNRS, Astroparticule et Cosmologie, F-75006 Paris, France

ABSTRACT

Inference from gravitational-wave observations relies on the availability of accurate theoretical waveform models to compare with the data. This paper considers the rapid generation of surrogate time-domain waveforms consistent with the gravitational-wave signature of the merger of spin-aligned binary black holes. Building on previous works, a machine-learning model is proposed that allows for highly-accurate waveform regression from a set of examples. An improvement of about an order of magnitude in accuracy with respect to the state of the art is demonstrated, along with a significant speed up in computing time with respect to the reference generation software tools.

Index Terms— Principal component regression – Chirp signals – Astrophysics – Gravitational waves – Waveforms generation

1. MOTIVATIONS

The first observation of gravitational waves (GW) by the LIGO/Virgo collaboration in 2015 [1] marked the advent of gravitational-wave astronomy. Since this date about 50 events have been detected [2] associated to the merger of compact star binaries, primarily binary black holes, the focus of this work.

Theoretical waveform models (or templates) are used to detect the gravitational-wave signals using matched filtering [2], or to infer the astrophysical parameters of their source using Bayesian samplers [3]. For both those tasks, a large number of template waveforms (about 10^5 to 10^6) are required to be computed to cover the relevant parameter space.

Waveform models are deduced from the resolution of the source dynamics which is a difficult relativistic problem (see [4] for a recent review). The evaluation of recent and accurate waveform models is computationally expensive and actually dominates the computational budget for parameter inference.

In the next decade, LIGO and Virgo detectors are expected to conduct at least two major observing runs with improved

sensitivity, leading to a large increase in the number of detected signals. The analysis of those future observations calls for numerically efficient, yet accurate waveform generators.

The waveform morphology being reasonably smooth and slowly varying with respect to its parameters (i.e., the binary component masses and their spins), it is possible to fit a signal model based on a generic amplitude and phase evolution from a set of example waveforms. We aim at reproducing *noise-free* waveform models with high fidelity thanks to a fast learning algorithm. The targeted accuracy is $\lesssim 10^{-3}$ in mismatch, a figure-of-merit defined later in Sec. 3.5. This has been successfully realized using reduced-order modelling (based on tensor spline fitting) [5, 6] and more standard machine learning approach such as mixture-of-experts regression [7] and artificial neural networks [8, 9].

In this work, we propose a model with reduced complexity based on principal component regression, that is able to improve the overall regression accuracy by about an order of magnitude. It can be implemented using off-the-shelf algorithms from `Scikit-learn` software library [10]. These performances are notably achieved thanks to a different choice of features.

2. GRAVITATIONAL-WAVE WAVEFORMS

Gravitational waves are oscillations of the space-time curvature propagating through space [1]. So far, the only sources detected are coalescences of compact star binaries, composed of either black holes and/or neutron stars. Here we focus on “spin-aligned” binary black hole (BBH) where the spins $S_i \in \mathbb{R}^3$ for $i = 1, 2$ of both objects are aligned with the orbital angular momentum \vec{L} i.e., the z axis normal to the orbital plane. In this case, the parameter space reduces to four parameters $\{m_1, m_2, S_{1z}, S_{2z}\}$ [11] with m_i for $i = 1, 2$ the mass of each black hole. The gravitational-wave signal is also determined by extrinsic parameters such as the luminosity distance r in megaparsecs (Mpc), i.e. the distance between the source and the observer, and the direction of the line of sight parametrized by the two angles (ι, φ) , namely the inclination of the source and the initial phase, see Fig. 1.

General relativity predicts the existence of two polarization modes for gravitational waves, the *plus* and *cross* polar-

This work was supported by the ANR research project RICOCHET n°ANR-21-CE48-0013.




Fig. 1. Physical parameters of a compact star binary – Intrinsic parameters $\{m_1, m_2, S_1, S_2\}$ (component masses and spins) and extrinsic parameters such as the direction of the line of sight (ι, φ) and the distance r . For spin-aligned binaries component spins are aligned with the orbital angular momentum \vec{J} so the intrinsic parameters reduce to $\{m_1, m_2, S_{1z}, S_{2z}\}$.

izations, denoted $h_+(t)$ and $h_\times(t)$. It is customary to express the two polarizations as a single complex-valued time series $h(t) = h_+(t) - ih_\times(t)$ referred to as *gravitational-wave strain*.

The computation of $h(t)$ needs to resolve the source dynamics. There are no exact closed-form solution for this relativistic problem, however accurate approximations are available. In the spin-aligned case and for binaries with low mass ratio q (defined as $q = m_1/m_2$ and with $q \geq 1$) and low total mass $M = m_1 + m_2$, the following expression [4] can be used for the strain

$$h(t) \propto \frac{M}{r} a(t) \left[\frac{1 + \cos^2 \iota}{2} \cos(\Phi(t) - 2\varphi) - i \cos \iota \sin(\Phi(t) - 2\varphi) \right] \quad (1)$$

where $a(t)$ and $\Phi(t)$ stand for the amplitude and phase of the dominant mode in the spin-weighted spherical harmonics expansion of $h(t)$ (see [4] for details).

General relativity’s scale invariance implies $h(t; m_1, m_2) = h(\lambda t; \lambda m_1, \lambda m_2)$ [7]. As a consequence, the total mass $M^* = m_1 + m_2$ can be fixed to a fiducial value¹ and in the case of spin-aligned BBH, waveforms can be described using a reduced number of three intrinsic parameters, e.g., $\{q, S_{1z}, S_{2z}\}$. The extrinsic parameters r^* , ι and φ correspond to simple scaling or phase factors that can be applied *a posteriori* in the waveform computation. We thus fix $r^* = 1$ Mpc, and $(\iota^*, \varphi^*) = (0, 0)$ in the sequel, with no loss of generality. With such a setting, the GW signal

¹Here, we used $M^* = 20 M_\odot$ with M_\odot the solar mass.

$h(t) \propto a(t)e^{-i\Phi(t)}$ appears as an amplitude and frequency modulated signal, often referred to as “chirp” for short, with $a(t)$ and $\Phi(t)$ its amplitude and phase parameters exhibiting non-oscillatory and smooth behaviours (see Fig. 2).

A range of models allows the generation of approximated yet accurate waveforms for data analysis purposes [4]. Here, we use SEOBNRv4 [11] based on the effective-one-body formalism (EOB). Fig. 2 shows an example of the expected gravitational-wave signal from a BBH merger.




Fig. 2. Top: example of gravitational-wave waveform obtained with the SEOBNRv4 model for a spin-aligned BBH. Bottom: associated amplitude and phase waveform attributes obtained from the waveform following Sec. 3.1.

3. WAVEFORM GENERATION MODEL

Following [7] we now introduce a generative model able to regress waveforms given $\theta = \{m_1, m_2, S_{1z}, S_{2z}\}$ the set of features that collects the input astrophysical parameters. The proposed regressor is a principal component regressor [12], that consists in regressing the PCA coefficients of the attributes from polynomial combining of the features. The waveform generation essentially performs the reverse process: obtain estimates for the PCA coefficients from the regressor, that are inverted to compute attributes, which in turn are mapped to waveforms.

3.1. Mapping to attributes

Over its entire duration, the chirp signal $h(t)$ goes through a sequence of three phases associated to different dynamical regimes: *inspiral* when the two black holes are far apart, *merger* when they are close and “plunge” onto each other

and *ringdown* when the merger remnant settles down to equilibrium [4]. The amplitude and phase evolve over different timescales during these phases. To capture their variations with a uniform accuracy over the entire waveform duration, amplitude and phase are discretized in time with a varying sampling resolution.

Time grid goes from $t_{\text{start}} = -20$ s to $t_{\text{end}} = 0.006$ with $t = 0$ at maximum amplitude. Concretely, 4096 sample points are computed over the non-uniform time grid $\tilde{t} = \text{sign}(t) |t|^{\frac{1}{\alpha}}$ where $\alpha = 0.35$ [7]. The sampling rate is reduced in the inspiral phase where both the amplitude and phase evolutions are slow, while it is increased when getting closer to the merger where the amplitude reaches its maximum, and where the phase evolves rapidly. This resampling has two benefits: data reduction during inspiral², and high accuracy near the merger.

3.2. Data reduction

The training set is reduced through a principal component analysis (PCA) of each waveform attribute after an alignment to 0 at $t = t_{\text{start}}$ for PCA efficiency purposes³. The minimal number of principal components (PC) required to reach the maximal capabilities of the model is determined by cross-validation (see Sec. 4.2).

3.3. Regression

Schmidt et al. [7] use a mixture of experts (MoE) regressor to infer the PCA coefficients from the set of features $\{q, \chi_{1z}, \chi_{2z}\}$ where $\vec{\chi}_i = \vec{S}_i/m_i^2$ for $i = 1, 2$ are the dimensionless spins. MoE is an ensemble learning method, based on a weighted sum of linear regressors called “experts” [12].

In this work, a single linear regressor is preferred and applied to a different set of features, expanded with polynomial combinations up to a pre-determined order. This maximum order is computed to maximize a score, see Sec. 4.2.

Choosing the initial set of features has a significant impact on the final score and regression accuracy. To leading order, the amplitude and phase evolution are known [4] to depend on the chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / M^{1/5}$, the mass ratio q and the effective spin $\chi_{\text{eff}} = (q\chi_{1z} + \chi_{2z}) / (1 + q)$. It is natural to think that those physically motivated parameters are good candidates to fit the data.

Systematic tests with subsets made of different feature combinations were performed and the final model use the feature set leading to the best score (see Sec. 4.2.2).

²This allows to easily process waveforms sampled at 16384 Hz on a standard laptop.

³The amplitude and phase offsets subtracted by the alignment procedure can be fitted and added back at the generation stage to produce the full waveform. Though we don’t detail this part here but this can be done with good accuracy with the same regressor.

3.4. Waveform generation

Once the model is trained, the usage goes as follows: the PC coefficients of the attributes are predicted for the parameters of the desired compact binary. The attribute time series are deduced from the predicted PC coefficients and the waveform amplitude and phase are interpolated from the attributes on the (uniform) time grid used for data analysis. The GW polarizations $h_+(t)$ and $h_\times(t)$ are finally computed by applying (1) with the requested distance, inclination and initial phase.

3.5. Scoring

The regression accuracy is evaluated by a specific metric called mismatch or unfaithfulness [13]. The mismatch between two waveforms $h, g \in \mathbb{C}^N$ is defined as:

$$\text{mismatch}(h, g) = \min_{\tau \in \mathbb{R}} \left[1 - \frac{|\langle h_\tau, g \rangle|}{\|h_\tau\| \|g\|} \right] \quad (2)$$

where $h_\tau(t) := h(t - \tau)$ and with the scalar product: $\langle h, g \rangle = \int \frac{h(f)g^*(f)}{S(f)} df$ defined in the Fourier domain and with f being the frequency variable. This metric is a loss function (smaller is better) and it is phase-shift and time-shift invariant. The definition in (2) allows for a frequency-dependent weighting, usually fixed to the GW detector noise power spectrum density $S(f)$. Here, we assume a flat noise curve $S(f) = 1$, which leads to a conservative constraint as the regression should then be equally accurate at all frequencies.

Errors in the waveform approximation lead to systematic errors in the astrophysical parameters estimates obtained from the observations. Those systematic errors from mis-modeling should be smaller than the statistical errors (due to the presence of noise in the observations). This principle leads to the rule of thumb (see e.g., App. G of [13]) stating that the mismatch should be $< N / (2 \text{SNR}^2)$ where N is the effective number of intrinsic parameters and SNR is the signal-to-noise ratio. Schmidt et al. [7] achieve a median mismatch value of 5×10^{-4} , with tails going to 10^{-1} in the worst case. In the spin aligned case with $N = 3$ effective parameters, this corresponds to an applicability range that goes up to $\text{SNR} = 54$ (3 in the worst case).

4. RESULTS

4.1. Training and testing datasets

A dataset of 4000 randomly distributed BBH waveforms was simulated with the SEOBv4 model from the `LALSimulation` software library [14]. The mass ratio q was randomly drawn uniformly in the range $[1, 20]$ and spins χ_{1z}, χ_{2z} uniformly over $[-0.8, 0.95]$. Such dataset is comparable to the one used in [7].

For the subsequent analysis, the dataset is splitted into a training and a testing set that corresponds to 80 % and 20 %

of the main dataset respectively. All reported results are obtained with the testing set.

4.2. Hyperparameter tuning

4.2.1. Number of PC after truncation

The overall modelling error is essentially due to the truncation of the PC expansion and the regression for each attribute. We first evaluate the former in Fig. 3 (circles and boxes). The phase approximation plays a dominant role. For instance, the truncation to only one PC leads to a mismatch of 5×10^{-5} for the amplitude and 7×10^{-1} for the phase.

Fig. 3 shows that the overall error after regression (crosses and plus signs) stabilizes at 6 PCs for the phase with a mismatch median score of about 10^{-5} . This is the retained number of PCs after truncation as this indicates the regressor fails to estimate higher order PC coefficients. For simplicity, the PCA expansion is truncated to the same number of PCs for both the amplitude and phase.




Fig. 3. Median mismatch vs number of principal components retained after PCA truncation. The mismatch is computed after fitting the amplitude only (phase is exact) for circles and plus signs, and the phase only for boxes and crosses. Circles and boxes represent mismatches solely due to the PCA truncation whereas crosses and plus signs represent mismatches due to the overall model. The dashed line indicates the selected PC truncation level at 6 PC.

4.2.2. Feature selection

Different feature combinations were tested and compared through their mismatch median value and dispersion. Promising sets of features were first pre-selected based on their r^2 score [10] obtained for the regression of the first PC of the phase, which essentially governs the overall performance (see Sec. 4.2.1). This procedure was performed for subsets that collect up to 6 parameters from the pool that includes q , \mathcal{M} , χ_{eff} and χ_{iz} , m_i and $1/m_i$ for $i = 1, 2$.

We found that about twenty feature sets lead to the best median mismatch of order 10^{-5} . Among those sets, the two following $\{\chi_{2z}, \chi_{\text{eff}}, \mathcal{M}\}$ and $\{\chi_{1z}, \chi_{2z}, q, m_2\}$ are of particular interest. The first has the nice property to possess three

features only, similarly to the number of intrinsic physical parameters. However, the mismatch distribution obtained from the testing set has much smaller variance with the second feature set. Using this feature set it is possible to produce many of the terms that appear in the so-called post-Newtonian expansion that approximates the signal phase before the final merger [4] (e.g., the component spins S_{1z}, S_{2z} before normalization, or the reduced mass $m_1 m_2 / (m_1 + m_2)$), thus explaining the good performances. Interestingly feature sets such as $\{q, \chi_{1z}, \chi_{2z}\}$ (used in [7]) or $\{q, \chi_{\text{eff}}, \mathcal{M}\}$ (a “natural” candidate composed of the physically motivated parameters) do not perform as well, with a median mismatch of 3×10^{-2} and 1×10^{-1} respectively.

4.2.3. Maximum order of polynomial features

The feature set $\{q, m_2, \chi_{2z}, \chi_{\text{eff}}\}$ is expanded by polynomial combinations of the initial features up to a maximum order [10]. The minimum order necessary to achieve best performances was determined by cross-validation. We found that a 7th order polynomial feature expansion is sufficient to minimize the median mismatch.

4.3. Accuracy and runtime benchmarking

Fig. 4 shows the mismatch distribution obtained with the testing set using $\{\chi_{1z}, \chi_{2z}, q, m_2\}$ as feature set. The median mismatch is 1.8×10^{-5} (average is 6.8×10^{-5}) and the 5% and 95% percentiles are 2.2×10^{-6} and 1.6×10^{-4} , resp. The worst case mismatch is $\sim 10^{-2}$. As shown in Fig. 4 we were not able to reproduce the performances given [7] with the Schmidt et al. model. Instead we found a median mismatch of 2.3×10^{-3} on the testing set.

No matter which reference is considered ([7] or our reproduction here), we obtain an improvement by at least an order of magnitude compared to the Schmidt et al. model. This makes the approximation error much smaller than the numerical and modelling error intrinsic to the SEOBNRv4 model (which is less than 1%) [11], uniformly over the considered parameter space. Following the rule given in Sec. 3.5, this corresponds to an applicability range that goes up to $\text{SNR} \approx 288$ (12 in the worst case) which covers the loudest SNR expected during the up-coming LIGO and Virgo science runs.

We evaluated the speed up factor defined as the ratio between the runtimes required to compute the original waveform model⁴ (using the LALSIMULATION software library [14]) to that of the proposed regression model. To perform this measurement, a set of 500 waveforms was generated by drawing the total mass uniformly between 40 and 100 M_{\odot} with the other properties as described in Sec. 4.1. The median speed up is measured to be $\sim 10^2$ (comparable to [7]). This shows that the proposed method can greatly accelerate parameter estimation with Bayesian samplers [3].

⁴The initial frequency is set to $f_{\text{min}} = 15$ Hz




Fig. 4. Mismatch histograms obtained with the proposed principal component regression model (orange) and that of Schmidt et al. (purple) [7].

The runtime of the generation with the regression model is dominated by the interpolation step from the non-uniform time grid of attributes to the uniform time grid of the user. The speed up factor could thus be further improved by dropping the non-uniform time grid (and thus the need for interpolation). Preliminary tests indicate that the same accuracy level can be obtained without the non-uniform time grid.

5. PERSPECTIVES

This work paves the way to further extensions relevant to gravitational-wave astronomy, such as the inclusion of higher/subdominant modes other than ($\ell = 2, m = \pm 2$). Another important extension is the case of precessing binaries (with arbitrary, misaligned spins) leading to arbitrarily polarized waveforms (not necessarily circularly polarized as it is the case here). In this more complex case with more input parameters and a larger waveform variability, fast and highly accurate modelling is still an open problem.

While the method presented here is evaluated in the context of gravitational-wave astronomy, it applies in principle to any area that involves chirp signals (see e.g., [15] for examples). One of the lessons learned from this study is that pragmatic regression approaches based on off-the-shelf tools may provide better results than more sophisticated approaches. A judicious choice of features appears critical to optimize the approximation accuracy.

6. REFERENCES

- [1] B. P. Abbott et al., “Observation of gravitational waves from a binary black hole merger,” *Phys. Rev. Lett.*, vol. 116, pp. 061102, 2016.
- [2] R. Abbott et al., “GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the 3rd observing run,” arXiv:2010.14527, 2020.
- [3] E. Thrane and C. Talbot, “An introduction to bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models,” *PASA*, vol. 36, pp. e010, 2019.
- [4] P. Schmidt, “Gravitational waves from binary black hole mergers: Modelling and observations,” *Front. Astron. Space Sci.*, vol. 7, no. 28, 2020.
- [5] M. Pürrer, “Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins,” *Phys. Rev. D*, vol. 93, no. 6, pp. 064041, 2016.
- [6] B. D. Lackey et al., “Surrogate model for an aligned-spin effective one body waveform model of binary neutron star inspirals using Gaussian process regression,” *Phys. Rev. D*, vol. 100, no. 2, pp. 024002, 2019.
- [7] S. Schmidt et al., “Machine learning gravitational waves from binary black hole mergers,” *Phys. Rev. D*, vol. 103, pp. 043020, 2021.
- [8] S. Khan and R. Green, “Gravitational-wave surrogate models powered by artificial neural networks,” *Phys. Rev. D*, vol. 103, pp. 064015, Mar 2021.
- [9] A. J. K. Chua et al., “Reduced-order modeling with artificial neurons for gravitational-wave inference,” *Phys. Rev. Lett.*, vol. 122, pp. 211101, May 2019.
- [10] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [11] A. Bohé et al., “Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors,” *Phys. Rev. D*, vol. 95, pp. 044028, 2017.
- [12] T. Hastie et al., *The elements of statistical learning*, Springer, 2nd edition, 2017.
- [13] K. Chatziioannou et al., “Constructing Gravitational Waves from Generic Spin-Precessing Compact Binary Inspirals,” *Phys. Rev. D*, vol. 95, no. 10, pp. 104004, 2017.
- [14] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite,” 2018, DOI: 10.7935/GT1W-FZ16.
- [15] P. Flandrin, “Time-frequency and chirps,” in *Wavelet Applications VIII*, Orlando, US, 2001, vol. 4391.