


Compact binary systems waveform generation with a generative pretrained transformer

Ruijun Shi^{1,*}, Yue Zhou^{2,*}, Tianyu Zhao¹, Zhoujian Cao¹, and Zhixiang Ren^{2,†}

¹*Department of Astronomy, Beijing Normal University, Beijing 100875, China*

²*Peng Cheng Laboratory, Shenzhen 518055, China*

 (Received 31 October 2023; accepted 11 March 2024; published 10 April 2024)

Space-based gravitational wave (GW) detection is one of the most anticipated GW detection projects in the next decade, which promises to detect abundant compact binary systems. At present, deep learning methods have not been widely explored for GW waveform generation and extrapolation. To solve the data processing difficulty and the increasing waveform complexity caused by the detector's response and second-generation time-delay interferometry (TDI 2.0), an interpretable pretrained large model named Compact Binary Systems Waveform Generation with Generative Pretrained Transformer (CBS-GPT) is proposed. For compact binary system waveforms, three models were trained to predict the waveforms of massive black hole binaries, extreme mass-ratio inspirals, and galactic binaries, achieving prediction accuracies of at most 99%, 91%, and 99%, respectively. The CBS-GPT model exhibits notable generalization and interpretability, with its hidden parameters effectively capturing the intricate information of waveforms, even with the complex instrument response and a wide parameter range. Our research demonstrates the potential of large models in the GW realm, opening up new opportunities and guidance for future researches such as complex waveforms generation, gap completion, and deep learning model design for GW science.

DOI: [10.1103/PhysRevD.109.084017](https://doi.org/10.1103/PhysRevD.109.084017)

I. INTRODUCTION

The first direct detection of a binary black hole merger (GW150914) [1,2] by the Laser Interferometer Gravitational-Wave Observatory (LIGO) has opened an innovative window to understand the Universe, which provides direct evidence for the validity of Einstein's general relativity. Gravitational wave (GW) observations will clarify many questions in astrophysics, cosmology, and fundamental physics [3–10]. So far, the ground-based GW detectors have reported over a hundred compact binary coalesce events [11], and recently pulsar timing array has also successfully detected sound evidence of the existence of stochastic GW background [12–15]. To gain a deeper understanding and an overall picture of GW cosmology [16], the field of low-frequency GWs needs to be widely covered. The space-based GW detection avoids terrestrial noise [17] and makes the detection of low-frequency (10^{-4} –0.1 Hz) GW signals more promising. Spaced-based GW detectors like Laser Interferometer Space Antenna (LISA) [18], Taiji [19,20], and Tianqin [21] have been planned and are scheduled for the 2030s. In particular, future space-based GW detection is expected to detect a richer variety of GW sources including massive black hole binaries (MBHBs),

extreme mass-ratio inspirals (EMRIs), and galactic binaries (GBs) [18].

GW signals are extremely weak and usually buried in instrumental noise. With the improvement of detector sensitivity and the increasing amount of data, the computational complexity and timeliness demands for detection and parameter estimation are growing, which are challenging problems for traditional methods that based on computing power of the central processing unit. With the rapid developing of graphics processing unit (GPU) computing power, artificial intelligence (AI) methods have shed some new light on this issue. Specifically, AI techniques have been successfully applied in various subjects such as GW signal detection [22–27], parameter estimation [28–30], signal extraction and noise reduction [31–38] with promising results. Additionally, the target of space GW detectors is also one type of complex and multiscale waveforms (such as MBHB, EMRIs, and GB). Some previous studies focused on generating binary black hole (BBH) waveforms. Lee *et al.* [39] employed a recurrent neural network that is capable of generating BBH waveforms during the merging and ringdown phases of nonspinning binary black hole coalescence. Khan *et al.* [40] demonstrated that a vanilla transformer can learn quasicircular and nonprecessing BBH waveforms. Similarly, Chua *et al.* [41] used a greedy algorithm to build a reduced basis, enabling the rapid generation of BBH waveforms. Recently, large-scale language models based on attention mechanism have shown

*These authors contributed equally to this work.

†Corresponding author: renzhx@pcl.ac.cn

their tremendous power in computer vision (CV) and natural language processing (NLP) [42–44]. Some studies indicate that similar architectures can be applied to the GW data analysis [35,36]. Space-based GW detectors will observe more signals along with complex difficulties such as source confusion, gaps, and glitches [45]. It is critical to provide a set of data processing tools to address these issues. Deep learning holds promise for meeting these challenges.

In contrast to previous studies on AI waveform generation, which had limitations in considering the second-generation time-delay interferometry (TDI 2.0) responses, our paper takes a step further. Moreover, the parameter range of waveforms in prior investigations was relatively narrow. In our paper, we are committed to further investigation on more complex waveforms and train a model to facilitate solving downstream problems. We introduce the (CBS-GPT) model, which is an interpretable, transformer-based, and self-supervised large model for prediction of compact binary sources (MBHB, EMRIs, and GB). In CBS-GPT (Fig. 1), patching and hybrid embedding mechanisms are proposed for full extraction of waveform features. By utilizing the self-attention mechanism and mean square error loss, CBS-GPT is trained for each GW waveform source. The experiment results illustrate that CBS-GPT can accurately predict the subsequent waveform based on the input waveform. In this study, two models were trained to achieve extrapolation with different input-to-prediction length ratios. In the 20:1 extrapolation, the average overlap between the predicted and target waveforms of MBHB, EMRIs, and GB reaches 0.981, 0.912, and 0.991, respectively. In the 1:1 extrapolation, the average overlaps reached 0.990, 0.807, and 0.992 for MBHB, EMRIs, and GB, respectively. We have also

discovered that waveform complexity can significantly influence the model’s prediction performance, and CBS-GPT can match the key frequencies effectively. Finally, through attention map visualization and correlation calculation, we discover that the attention map and its corresponding waveform present similar periodic distribution, which illustrates that CBS-GPT is able to learn waveform features even under the complex instrument response and a wide parameter range.

The rest of this paper is organized as follows. Section II describes data generation and the CBS-GPT model architecture. In Sec. III, we present our overlap and attention map results, and discuss interpretability outcomes as well as potential applications. Finally, Sec. IV highlights our findings based on the results.

II. METHODOLOGY

A. Data

Space-based GW detectors’ targets are GW signals at frequencies of $[10^{-4}, 0.1]$ Hz. We focus on three compact binary sources that are of major interest for LISA: MBHB, EMRIs, and GB. Figure 2 displays data examples. Detailed information of the data generation process is given below.

1. MBHB

MBHBs are one of the space-based GW detector’s main detection targets [18]. In this paper, SEOBNRV4_OPT [46] ($l = m = 2$ mode) is used to generate the MBHB waveforms. The parameter space of the MBHB dataset is shown in Table I. In Fig. 2(a), the time-delay interferometry (TDI) 2.0 transfer function significantly affects high-frequency transmissions due to the lower total mass of MBHB. First,

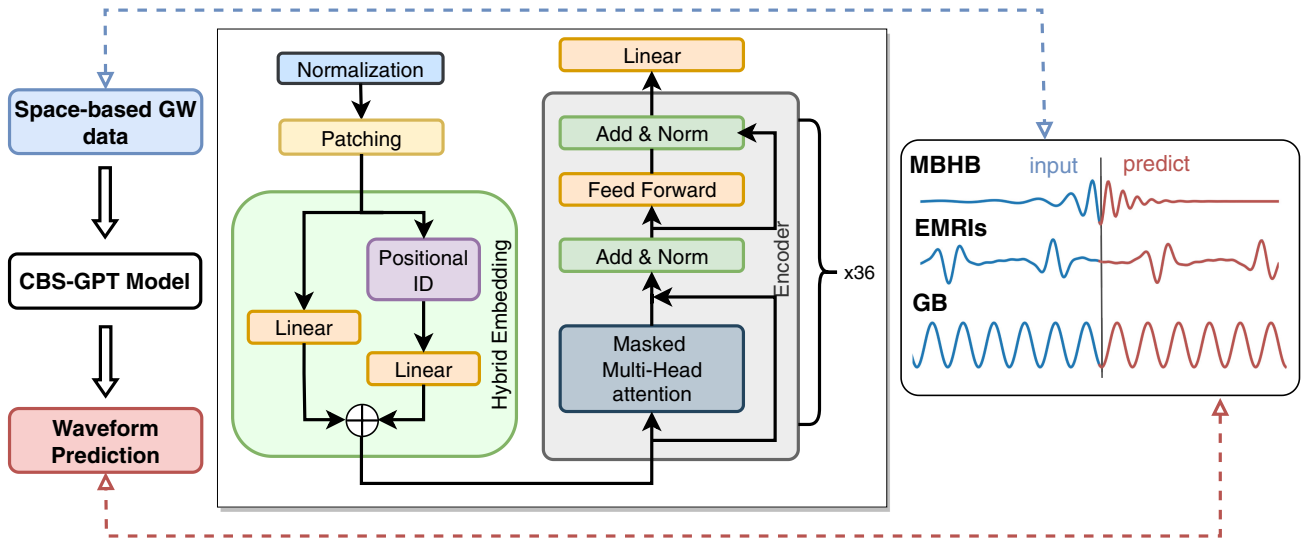


FIG. 1. Overview of CBS-GPT. The CBS-GPT model was trained separately for three kinds of GW sources (MBHB, EMRIs, and GB). The subsequent waveform can be extrapolated after feeding its corresponding preceding waveform into CBS-GPT. Details of data and model description are in Sec. II.

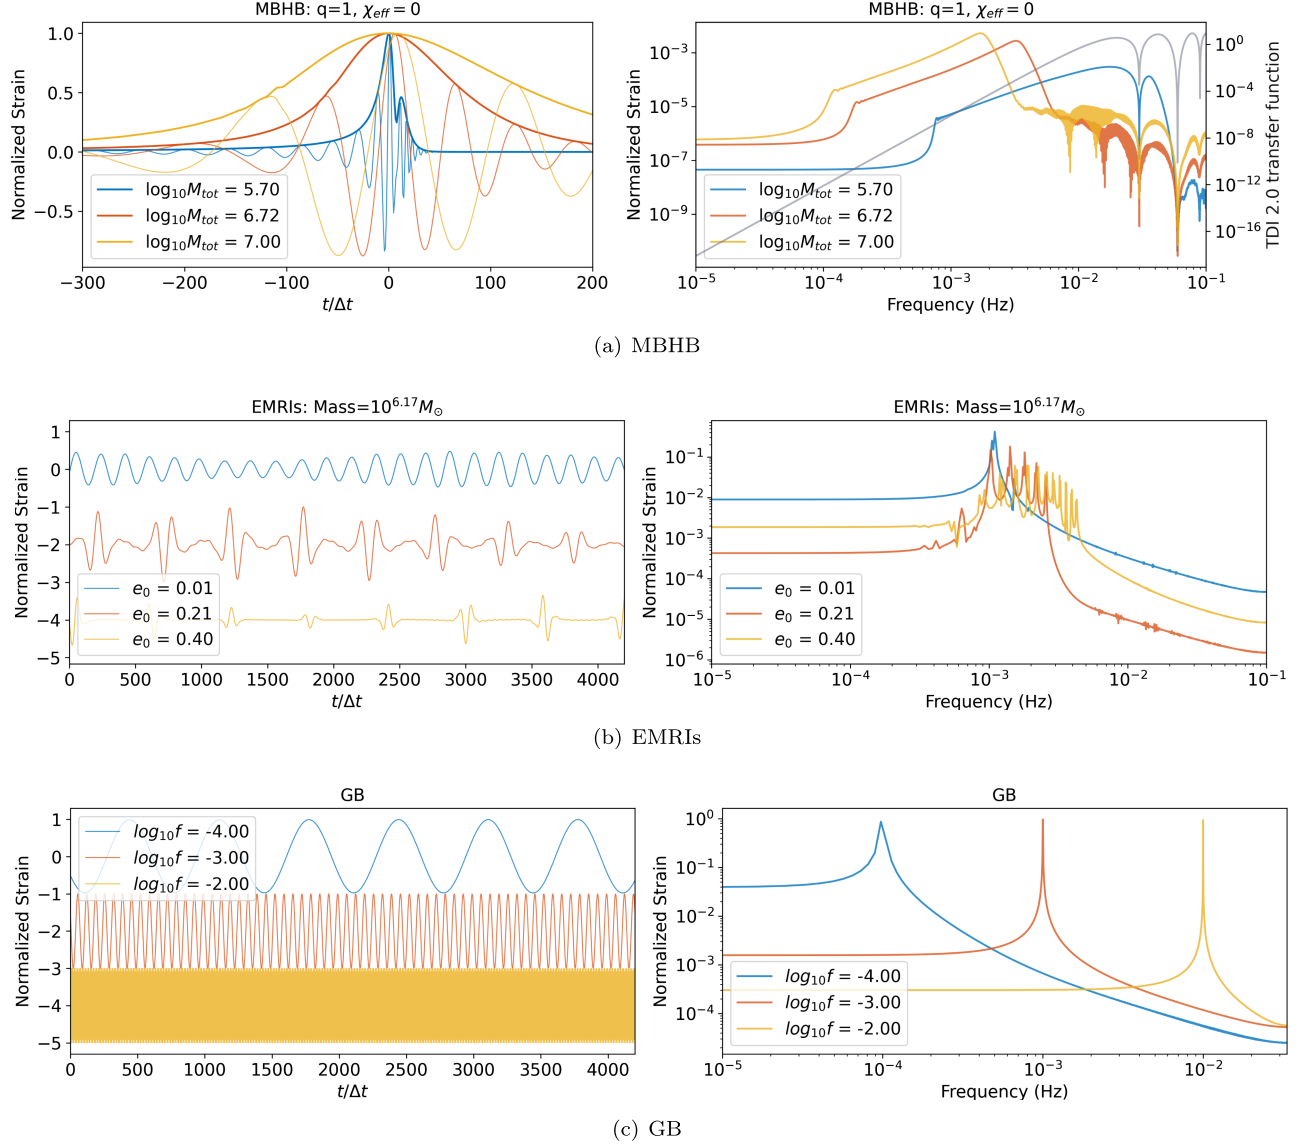


FIG. 2. TDI 2.0 response complicates waveforms. To simplify waveform comparison here, all waveforms were standardized to a maximum amplitude of 1. The Δt in the figure represents the sampling rate. The effects of different parameters on time and frequency domain are shown on the left and right panels. (a) MBHB waveforms at different M_{tot} . At high frequencies, the TDI response function has a greater impact. The gray line represents the TDI 2.0 transfer function in the frequency domain. (b) EMRI waveforms at different e_0 . As the eccentricity increases, the EMRI waveform becomes more and more complex in the frequency domain. (c) GB waveforms at different f . The GB signal is relatively simple and is a single-frequency signal.

we generate MBHB time-series waveforms with a length of 20,000 points with a sampling rate of 5 seconds. We train two models with different input-to-prediction token length ratios. The two scenarios are referred to as 20:1 and 1:1 extrapolation in the subsequent sections. Table II summarizes token information for each source. The 20:1 extrapolation involved predicting the subsequent 200 points after merging with an input sequence of the preceding 4000 points. The 1:1 extrapolation predicted the same 200 points after merging, but with an input sequence limited to the preceding 200 points. During the inference phase, the 4,000 valid points (or 200 valid points) before the merge

time are fed into CBS-GPT to predict the succeeding 200 points, hence achieving a 20:1 extrapolation (or 1:1 extrapolation) prediction of the MBHB waveforms.

2. EMRIs

EMRIs are a kind of black hole binary system with a mass ratio of $m/M \simeq 10^{-4} - 10^{-7}$ and massive black holes (MBHs) that have a mass range of $M \simeq 10^5 - 10^7 M_{\odot}$. EMRI waveforms are able to encapsulate the properties of space-time near a massive black hole. EMRIs are among the primary detection targets for the space-based GW detectors, possessing the

TABLE I. Parameter distribution of training dataset and test dataset.

Parameter	Description	Parameter distribution
(a) Parameter space of MBHB dataset		
M_{tot}	Total mass of massive black hole binaries $m_1 + m_2$	log-uniform $[5.5, 7]M_{\odot}$
q	Mass ratio $\frac{m_2}{m_1}$	Uniform $[0.1, 1]$
S_1^z, S_2^z	Spin parameters of two black holes	Uniform $[-0.99, 0.99]$
i, ψ	The inclination angle and polarization angle	Uniform $[0, \pi]$
ϕ_c	Coalescence phase	Uniform $[0, 2\pi]$
λ	Ecliptic longitude	Uniform $[0, 2\pi]$
β	Ecliptic latitude	Uniform $[0, \pi]$
Parameter	Description	Parameter distribution
(b) Parameter space of the EMRI dataset		
M	The mass of MBH	Uniform $[10^5-10^7]M_{\odot}$
m	The mass of stellar-mass compact	Fix $[10M_{\odot}]$
a	Spin parameter of MBH	Uniform $[10^{-3}, 0.8]$
p_0	Semi-latus rectum	Uniform $[10, 16]$
e_0	Eccentricity	Uniform $[10^{-3}, 0.4]$
i_0	The cosine of the orbit's inclination angle from the equatorial plane	Uniform $[-0.98, 0.98]$
θ_S, θ_K	The polar angles describing the sky location and the orientation of the spin angular momentum vector of the MBH	Uniform $[10^{-3}, \pi]$
θ_S, ϕ_K	The azimuthal angles describing the sky location and the orientation of the spin angular momentum vector of the MBH	Uniform $[10^{-3}, 2\pi]$
$\Phi_{\varphi,0}, \Phi_{\theta,0}, \Phi_{r,0}$	The phase of azimuthal, polar, and radial modes	Fix $[0]$
Parameter	Description	Parameter distribution
(c) Parameter space of GB dataset		
f	Frequency	log-uniform $[-4, -2]$ Hz
\dot{f}	The derivative of f	Fix $[10^{-14}]$
A	Amplitude	Uniform $[10^{-23}, 10^{-21}]$
i_0, ψ, ϕ_0	The inclination angle, polarization angle, and initial phase	Uniform $[0, \pi]$
λ	Ecliptic longitude	Uniform $[0, 2\pi]$
β	Ecliptic latitude	Uniform $[0, \pi]$

TABLE II. The waveform information of different sources.

	20:1 extrapolation	1:1 extrapolation
Input tokens	1000	50
Prediction tokens	50	50
MBHB	4 points/token	4 points/token
EMRIs	4 points/token	16 points/token
GB	4 points/token	32 points/token

potential to unveil new physical phenomena [47–49]. We employ the FASTEMRISWAVEFORMS (FEW) package [50] to generate EMRI waveforms with a sampling rate of 5 s. The EMRI signals with a duration of 1 year are randomly sliced into five waveform segments containing 4,200 points for 20:1 extrapolation (or 1600 points for 1:1 extrapolation). For continuous GWs, the random slice can simulate variations in the phase and amplitude domain of the same signal, enhancing the model's generalization capability.

The parameter space of the EMRI dataset is shown in Table I. The complexity of the EMRI waveform is visible in Fig. 2(b). As the eccentricity increases, there is a corresponding increase in its complexity, which becomes particularly prominent in the frequency domain.

3. Galactic binary

Within the Milky Way galaxy, a substantial population of binary white dwarf systems exists, posing foreground noise challenges for space-based GW detectors. We use the following GB model to generate GB waveforms [51]:

$$\begin{aligned}
 h_+^{\text{src}}(t) &= \mathcal{A}(1 + \cos^2 i) \cos \Phi(t), \\
 h_{\times}^{\text{src}}(t) &= 2\mathcal{A} \sin i \sin \Phi(t), \\
 \Phi(t) &= \phi_0 + 2\pi f_0 t + \pi \dot{f}_0 t^2 + \frac{\pi}{3} \ddot{f}_0 t^3, \\
 \ddot{f}_0 &= \frac{11}{3} \frac{\dot{f}_0^2}{f_0}.
 \end{aligned} \tag{1}$$

Similar to EMRIs, GB waveforms are generated with a duration of 1 year and a sampling rate of 1/15 Hz. Five slices of 4,200 points for 20:1 extrapolation (or 3200 points for 1:1 extrapolation) are randomly truncated for training and inference. The parameter space of the GB dataset is shown in Table I.

4. Detector response and TDI 2.0

After generating the waveform, we project it into the LISA detector [52]. For LISA, the signals will be processed with TDI combination to suppress the overpowering laser noise. The response of space-based GW detectors is more intricate compared to ground-based detectors, accounting for factors such as satellite orbits and arm-length delays. The strain induced on link 12 is

$$H_{12}(t) = h_+^{\text{SSB}}(t) \times \xi_+(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) + h_{\times}^{\text{SSB}}(t) \times \xi_{\times}(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}). \quad (2)$$

The $\xi_{+,\times}$ refers to the antenna pattern:

$$\begin{aligned} \xi_+(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) &= (\hat{\mathbf{u}} \cdot \hat{\mathbf{n}}_{12})^2 - (\hat{\mathbf{v}} \cdot \hat{\mathbf{n}}_{12})^2, \\ \xi_{\times}(\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{n}}_{12}) &= 2(\hat{\mathbf{u}} \cdot \hat{\mathbf{n}}_{12})(\hat{\mathbf{v}} \cdot \hat{\mathbf{n}}_{12}), \end{aligned} \quad (3)$$

where $\hat{\mathbf{n}}_{12}$ is the link unit vector, $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ represent polarization vectors defined as the opposite direction of the polar and azimuthal angles in the Solar System bary-center (SSB) frame, respectively. Because of the longer arm lengths of space-based GW detectors, the influence of arm length needs to be taken into consideration. The time of transmission from spacecraft 2 is denoted as t_2 , and after propagating over the arm length distance to reach spacecraft 1, the reception time is t_1 ,

$$t_1 \approx t_2 + \frac{L_{12}}{c} - \frac{1}{2c} \int_0^{L_{12}} H(x(\lambda), t(\lambda)) d\lambda, \quad (4)$$

where L_{12} represents the arm length of the detector. The variable λ describes the path of the photon, and we approximate t_1 to the first order as $t(\lambda) \approx t_2 + \lambda/c \approx t_2 + L_{12}/c$. Because of the slow motion of the space-based GW detector, the frequency shift is given by

$$y_{12}(t_1) \approx \frac{1}{2(1 - \hat{\mathbf{k}} \cdot \hat{\mathbf{n}}_{12}(t_1))} \left[H_{12}\left(t_1 - \frac{L_{12}(t_1)}{c} - \frac{\hat{\mathbf{k}} \cdot \mathbf{x}_2(t_1)}{c}\right) - H_{12}\left(t_1 - \frac{\hat{\mathbf{k}} \cdot \mathbf{x}_1(t_1)}{c}\right) \right], \quad (5)$$

where $\hat{\mathbf{k}}$ represents the propagation vector of the wave source.

Space-based GW detectors have unequal arm lengths, which results in significant laser frequency noise. To mitigate this issue, TDI techniques are commonly

employed to suppress laser frequency noise [53,54]. The first and second generation Michelson combinations, X1 and X2, are defined by [53]

$$X_1 = y_{13} + \mathbf{D}_{13}y_{31} + \mathbf{D}_{131}y_{12} + \mathbf{D}_{1312}y_{21} - [y_{12} + \mathbf{D}_{12}y_{21} + \mathbf{D}_{121}y_{13} + \mathbf{D}_{1213}y_{31}], \quad (6)$$

$$\begin{aligned} X_2 = X_1 &+ \mathbf{D}_{13121}y_{12} + \mathbf{D}_{131212}y_{21} + \mathbf{D}_{1312121}y_1 \\ &+ \mathbf{D}_{13121213}y_{31} - [\mathbf{D}_{12131}y_{13} + \mathbf{D}_{121313}y_{31} \\ &+ \mathbf{D}_{1213131}y_{12} + \mathbf{D}_{12131312}y_{21}], \end{aligned} \quad (7)$$

where the delay operators are defined by

$$\mathbf{D}_{i_1, i_2, \dots, i_n} x(t) = x\left(t - \sum_{k=1}^{n-1} L_{i_k i_{k+1}}(t)\right). \quad (8)$$

The detector response and TDI 2.0 response of GW are calculated using FASTLISARESPONSE [52]. TDI 2.0 generates three channels: X, Y, and Z. The variables Y and Z may be produced via cyclic permutation of the indices in Eq. (7). A more detailed derivation can be found in Sec. IV of Ref. [52]. By combining X, Y, and Z, three independent channels A, E, and T are obtained,

$$\begin{aligned} A &= (Z - X)/\sqrt{2}, \\ E &= (X - 2Y + Z)/\sqrt{6}, \\ T &= (X + Y + Z)/\sqrt{3}. \end{aligned} \quad (9)$$

The incorporation of response functions and TDI 2.0 combination introduces increased complexity to the waveform, especially in the high-frequency part. As depicted in Fig. 2, MBHB waveforms exhibit significant differences at various parameter values.

B. CBS-GPT model

Transformers [55] are a class of deep learning models that have exhibited excellent performance in various tasks, such as NLP [43] and CV [44]. We incorporate the masked self-attention mechanism and feed-forward neural network to build our CBS-GPT model.

Patching. First, the input waveform is preprocessed by standardization, which facilitates the model in capturing waveform information more effectively:

$$I = \text{standard}(s) = \frac{s - \text{mean}(s)}{\text{std}(s)}, \quad (10)$$

where $s = \{s_i | i \in [0, N]\}$ represents the input waveform, $\mu = \text{mean}(s) = \frac{1}{N} \sum_{i=1}^N s_i$ and $\text{std}(s) = \sqrt{\frac{\sum_{i=1}^N (s_i - \mu)^2}{N}}$ represent the mean and standard deviation of the waveform, respectively. The standardization centers the original data

to a mean of 0 and a standard deviation of 1, which makes features have equal weight in various analyses and is more suitable for machine learning algorithms that are sensitive to feature scales. Then, $I = \{x_i | i \in [0, N]\}$ is divided into nonoverlapping patches, and we refer to each patch as a ‘‘token’’ here. In our 20:1 extrapolation experiment for example, we have an input waveform with $N = 4200$ sampling points, which is segmented into $num = 1,050$ tokens, and each token contains four points. Each token is treated as a vector, after patching, the standardized waveform I is processed into the input matrix $I' \in \mathbb{R}^{[1050,4]}$.

Hybrid embedding. The hybrid embedding module is utilized in our model, because each token contains richer physical information and cannot be tokenized by simple tokenizers as in NLP. As Fig. 1 shows, it is combined with a token embedding layer and a positional embedding layer [Eq. (11)]. The token embedding layer performs linear projection to achieve dimension matching with following encoder blocks, which meanwhile preserves the entire information of the input waveform. The positional embedding is also a linear layer that encodes positional relationships between tokens, which is rather important in improving prediction accuracy [55]:

$$\begin{aligned} E_e &= I'W_e \\ E_p &= \mathbf{I}_{num}W_p \\ E_{\text{hybrid}} &= E_e + E_p, \end{aligned} \quad (11)$$

where $W_e \in \mathbb{R}^{[4, d_{\text{model}}]}$, $d_{\text{model}} = 2048$ and $W_p \in \mathbb{R}^{[num, d_{\text{model}}]}$ are both learnable parameters, and \mathbf{I}_{num} represents an identity matrix with shape $[num, num]$.

Encoder block. The encoder contains $n_{\text{block}} = 36$ blocks. Each block mainly consists of an attention module and a feed-forward neural network. As for the attention module, masked multiheads self-attention (MMHSA) is adopted in our work, which enables information to be projected into matrices in different ways, thereby enhancing the expressive capacity of the model. The computation process of the attention module is as follows:

$$\begin{aligned} Q_{ji} &= W_{ji}^Q x_j \\ K_{ji} &= W_{ji}^K x_j \\ V_{ji} &= W_{ji}^V x_j \end{aligned} \quad (12)$$

$$\text{head}_{ji}(Q_{ji}, K_{ji}, V_{ji}) = \text{softmax}\left(\frac{Q_{ji}K_{ji}^T \cdot \text{mask}}{\sqrt{d}}\right)V_{ji}, \quad (13)$$

$$\text{MMHSA}_j(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_H)W_j^E, \quad (14)$$

$$H'_j = \text{LayerNorm}(\text{MMHSA}(Q, K, V)) + x_j. \quad (15)$$

In each encoder block, there are $H = d_{\text{model}}/64 = 32$ heads. $W_{ji}^Q, W_{ji}^K, W_{ji}^V$ represent learnable query, key, and

value parameters of i th attention head and j th encoder block, respectively, and mask is a lower triangular standard matrix:

$$x_j = \begin{cases} E_{\text{hybrid}}, & j = 0 \\ y'_{j-1}, & 0 < j < n_{\text{block}}, \end{cases} \quad (16)$$

where x_j represents the hybrid embedding or the output of the previous encoder block. The feed-forward network (FFN) is composed of two dense layers and is connected to each attention module. We employ the residual connection [Eq. (18)], which is helpful to alleviate the gradient-vanishing problem:

$$\text{Inter}(H'_j) = \text{GeLU}(H'_jW_{j1} + b_{j1})W_{j2} + b_{j2}, \quad (17)$$

$$y_j = \text{FFN}(H'_j) = \text{LayerNorm}(\text{Inter}(H'_j)) + H'_j, \quad (18)$$

where $W_{j1}, b_{j1}, W_{j2}, b_{j2}$ are learnable parameters and GeLU is an activation function. Finally, the output of the last encoder block is inversely projected to the same shape of I' :

$$y' = y_{n_{\text{block}}}W_i^T, \quad y' \in \mathbb{R}^{[1050,4]}. \quad (19)$$

Loss function. Next-token-prediction error is adopted to train CBS-GPT, which means that the predicted token y'_m is designed to match the input token (I'_{m+1}) at position $m + 1$. Hence, only $num - 1$ tokens are taken into account when calculating the training loss. Specifically, the mean squared error loss is used to measure the difference between the predictions:

$$\mathcal{L} = \frac{1}{(num - 1) \times 4} \sum_{m=0}^{num-2} \sum_{t=0}^3 \|y'_{m,t} - I'_{m+1,t}\|^2. \quad (20)$$

C. Training and inference

During training, the Adam [56] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ is used, and the initial learning rate is $2e-4$. There are 1.6 million waveforms in the training dataset of each model, and the parameter of each waveform is randomly selected from its corresponding parameter space. After passing through the LISA response, each waveform is divided into three TDI channels (A, E, and T). In this study, the E channel is selected to train the model. The model was trained on two NVIDIA V100 GPUs for approximately 30 hours. During inference, for each signal source, 10,000 waveforms are generated to test CBS-GPT’s performance. For each waveform, the initial input contains 1,000/50 valid tokens and 50 masked tokens that are masked with zero, whose corresponding value in the mask matrix also equals zero, which guarantees that no attention is paid to the to-be-extrapolated token. In the first step, the 1,001st/51st token

is predicted and replaces the previous 1,001st/51st token, and so forth, 50 successive tokens are predicted based on 1,000/50 valid input tokens.

III. RESULTS AND DISCUSSION

During inference, overlap is defined to evaluate the extrapolation accuracy of the predicted waveform. Overlap is calculated between the target waveform and the predicted waveform generated by CBS-GPT as stated in Eq. (21). The overlap \mathcal{O} ranges between $[0, 1]$, with values closer to 1 indicating that the predicted waveform is more similar to the target waveform:

$$\mathcal{O}(h_t, h_p) = \max_{t_c} (\hat{h}_t | \hat{h}_p [t_c])^{1/2}, \quad (21)$$

with

$$(h|s) = 2 \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}^*(f)\tilde{s}(f) + \tilde{h}(f)\tilde{s}^*(f)}{S_n(f)} df, \quad (22)$$

$$\hat{h} = \frac{h}{\sqrt{(h|h)}},$$

where t_c represents time shifted, and we set $S_n(f) = 1$.

Overall, in the context of 20:1 extrapolation tasks targeting MBHB, GB, and EMRIs signals, CBS-GPT has demonstrated remarkable efficacy, with over 50% of the overlaps exceeding 0.99. Figures 3 and 4 showcase the prediction performance of each waveform under varying parameter conditions, revealing that the CBS-GPT model

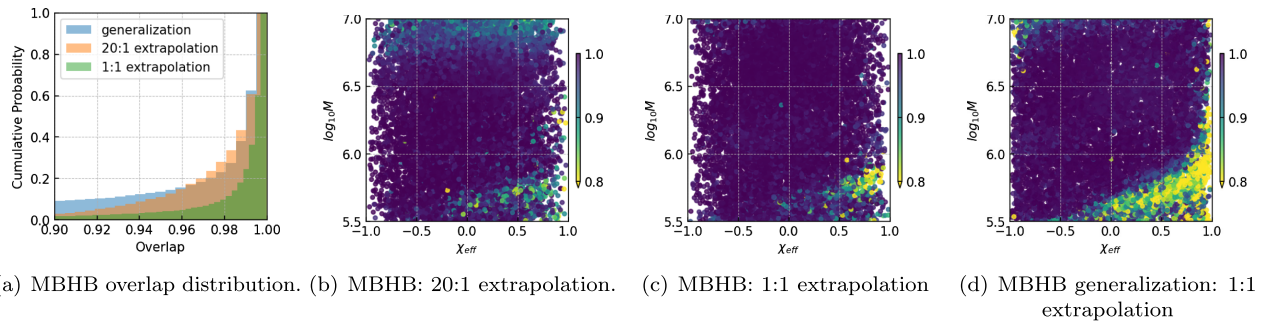


FIG. 3. The overlap distribution of MBHB is shown in (a). Parts (b), (c), and (d) portray the heat maps of M_{tot} and χ_{eff} parameters, which have the greatest impact on overlap. A darker color corresponds to a higher overlap value.

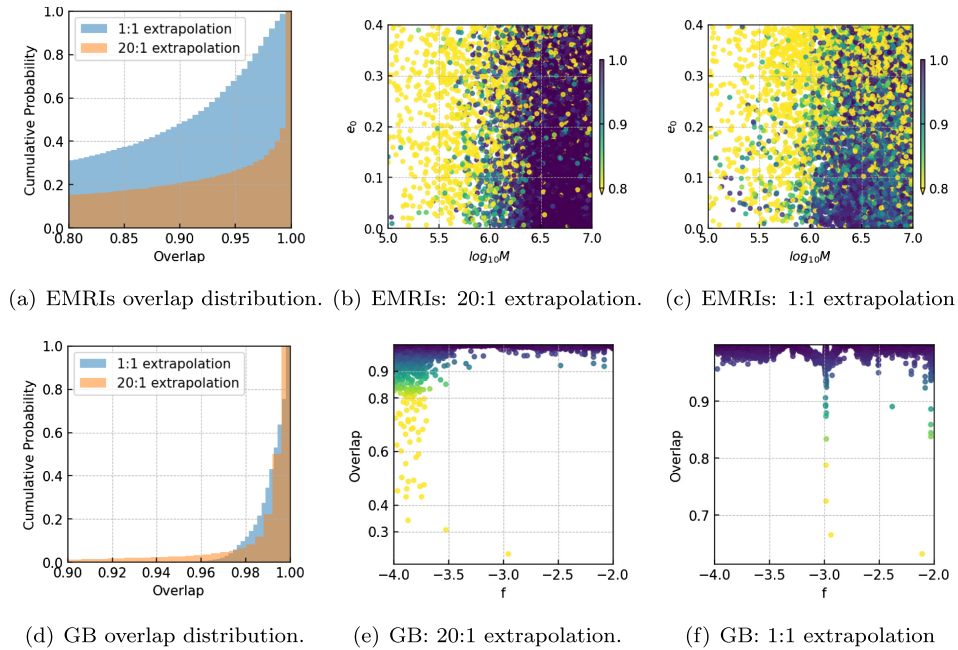


FIG. 4. The overlap distributions of EMRIs and GB are shown in (a) and (d). Parts (b) and (c) portray the heat maps of e_0 and M parameters, which have the greatest impact on overlap of EMRIs. Similarly, (e) and (f) portray the heat maps of frequency parameter f , which have the greatest impact on overlap of GB. A darker color corresponds to a higher overlap value.

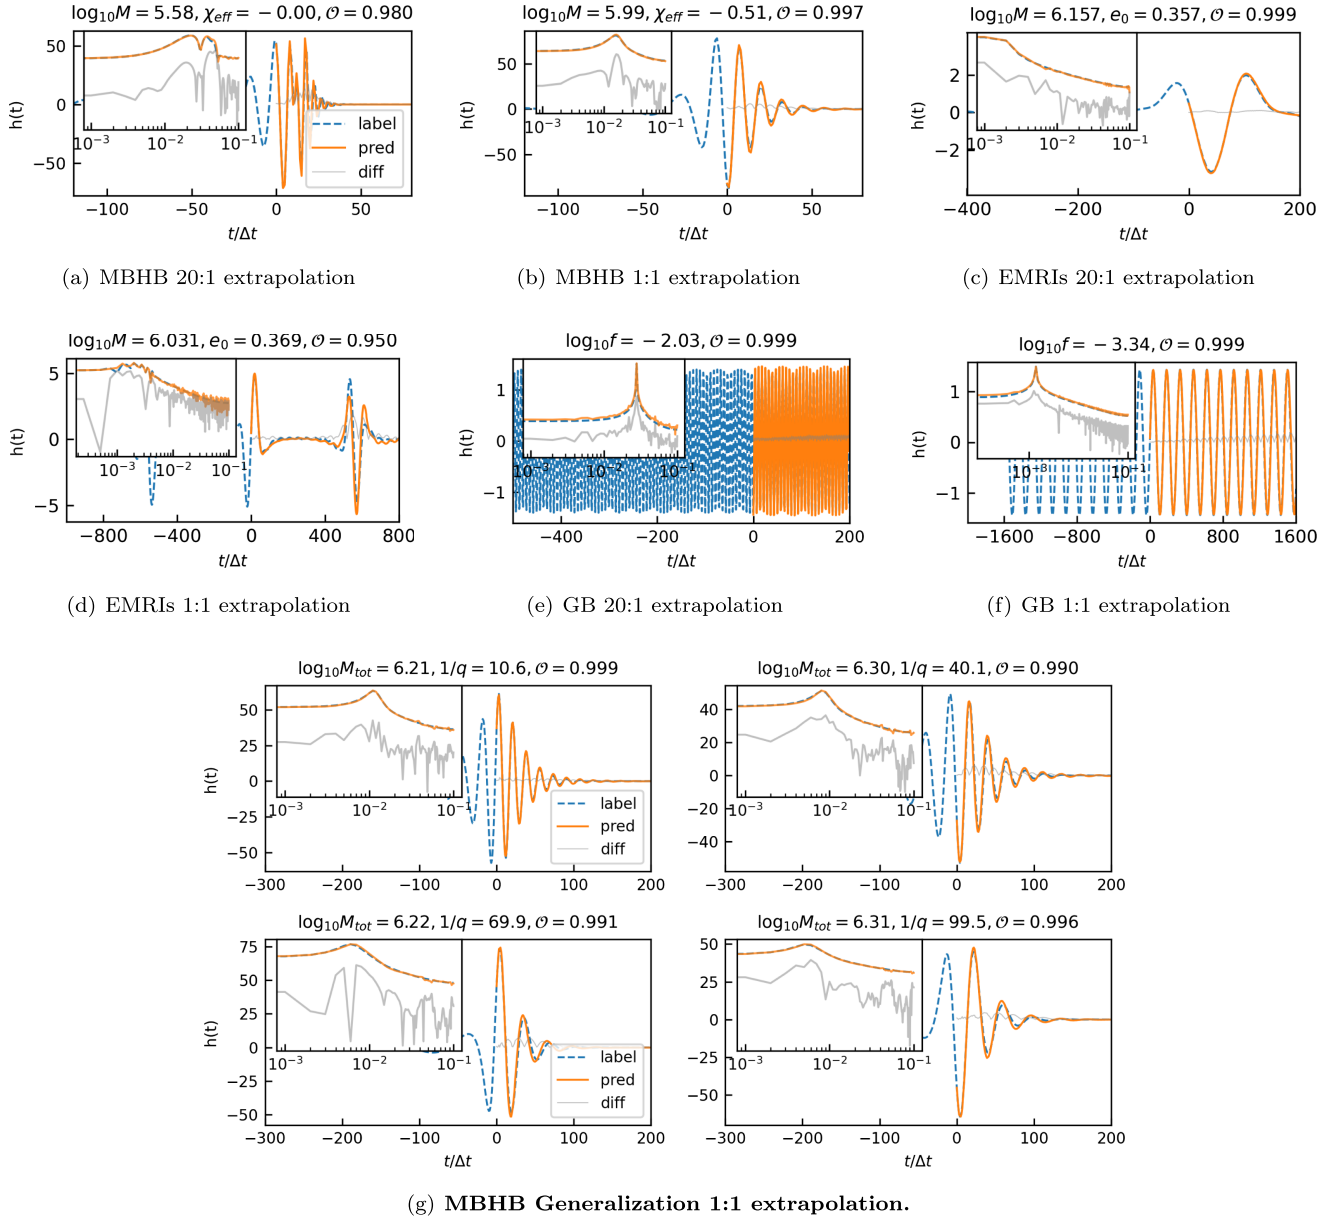


FIG. 5. CBS-GPT prediction results. (a),(b) MBHB results. (c),(d) EMRIs results. (e),(f) GB results. (g) Generalization results of MBHB waveform with $1/q \approx 10, 40, 70$, and 100 , respectively. We set the predicted starting point at time zero. The blue line represents the conjunction of the last part of the input waveform and target label, the orange line is the predicted waveform, and the gray line is the difference between the predicted and target waveform. The inset in each part represents the anticipated and target waveforms in the frequency domain, as well as the differences between them.

can learn waveform features with a wide range of parameters. Figures 5 and 6 demonstrate the generalization and potent interpretability of CBS-GPT.

A. Results of MBHB

The results of MBHB overlap are shown in Table III. The CBS-GPT model is sensitive to total mass, mass ratio, and spin parameters. Here we use χ_{eff} to represent the spin parameter [57]:

$$\chi_{\text{eff}} = \frac{S_1^z}{1+q} + \frac{qS_2^z}{1+q}. \quad (23)$$

20:1 extrapolation. The overlap distribution and waveform examples are shown in Figs. 3(a) and 5(a), with mean and median overlaps equal to 0.981 and 0.992, respectively. The overlap results reveal that CBS-GPT can forecast the waveform of the merge-ringdown phase based on the

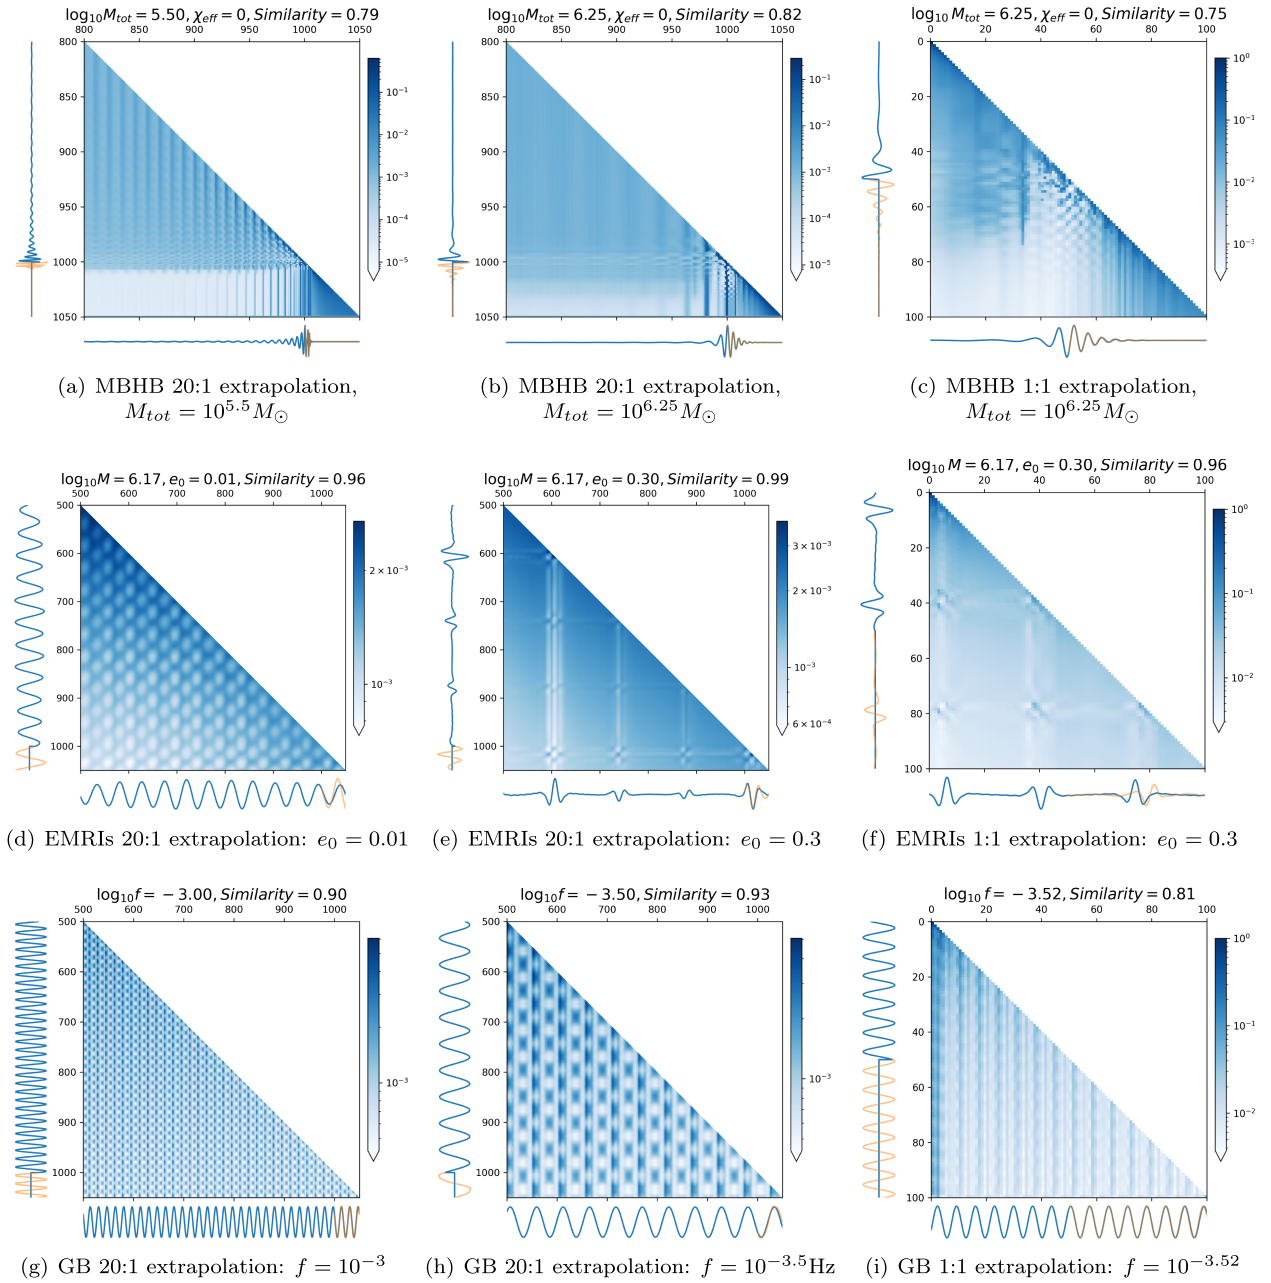


FIG. 6. Attention maps of the last encoder layer. For a clear presentation, only part of the attention map is displayed. The blue lines on the left and bottom panels represent the input waveforms, whose 1001st (or 101st) to 1050th (or 50th) tokens are padded with zero value during inference, and the orange line represents the waveform predicted by CBS-GPT. The term “Similarity” in the title of each figure denotes the correlation coefficient between the waveform and the attention map.

inspiral phase characteristics. CBS-GPT exhibits optimal inference performance when the total mass is approximately $10^{6.5} M_{\odot}$ as shown in Fig. 3(b). This phenomenon has also been observed in other signal sources. The overlap is lower for waveforms with low total mass and high effective spin χ_{eff} . Comparing low- and high-mass situations to those involving intermediate masses, the performance of midfrequency band prediction is the best. Since TDI 2.0 transfer functions in the high-frequency part are

more complex [58,59], the waveform is also more complex. Consequently, the model’s performance experiences a slight decrease. But even under such less ideal circumstances, CBS-GPT can still successfully recover a significant portion of the signals.

1:1 extrapolation. We find that the model pays little attention to the early-stage waveform and mainly concentrates on the late-stage inspiral waveform when forecasting the merging waveform of an MBHB (detailed explanation

is in Sec. III C). This demonstrates the marginal contribution of early-stage inspiral waveforms to subsequent waveform generation. Hence, we retrained a model, whose input only contains 200 points before merge time and predicted the subsequent 200 points, thus achieving a 1:1 extrapolation. The average and median overlap achieved 0.990 and 0.996, respectively. The results are slightly better than the previous 20:1 extrapolation, which validates our former conclusion. In Table III, we observe a noticeable improvement in overlap for cases with masses greater than $10^6 M_\odot$, which illustrates that shorter input waveforms allow the model's attention to be more focused, leading to improved inference performance. In Fig. 5(b), we showcase the predictive performance of CBS-GPT in the 1:1 extrapolation scenario.

Generalization ability refers to the performance of a model when applied to data that has not been seen before. To evaluate the generalization capability of CBS-GPT, we selected MBHB signals with mass ratios ranging from 1:10 to 1:100 in the 1:1 extrapolation model. Figure 5(g) showcases the waveform examples of generalization ability. The average overlap achieved 0.970, with more than half of the overlaps surpassing 0.993, which demonstrated the strong generalization ability of our method. The model's performance on generalization experiment also illustrates its ability to learn the essence of the data.

B. Results of continuous waveform: EMRIs and GB

The overlap distributions of the EMRIs and GB are shown in Fig. 4 and their mean and median values are displayed in Table III. Examples of predicted EMRIs and GB waveforms are shown in Figs. 5(c)–5(f).

20:1 extrapolation. Regarding GB, its mean and median overlap both exceed 0.99. The mean and median overlap of EMRIs is equal to 0.912 and 0.997, respectively. While the mean overlap of EMRIs is slightly lower, its median overlap aligns with that observed in MBHB and GB waveforms.

Specifically, the overlap distribution of EMRIs is significantly influenced by the mass parameters and eccentricity parameters. As depicted in Table III, when e_0 is less than 0.1, the majority of overlaps remain below 0.9. As the eccentricity increases, the waveform features become more complex in waveform amplitude. Therefore, when the eccentricity is higher, the corresponding overlap tends to decrease.

In contrast to MBHB and EMRIs signals, the GB signal presents a comparatively straightforward, single-frequency waveform. As for GB, the frequency parameter has the greatest impact on the waveform. When the frequency is larger than $10^{-3.5}$ Hz, the overlap is basically higher than 0.9. The result of GB signals demonstrates the model's sensitivity over frequency, with the distinct preference for

TABLE III. The overlap results.

MBHB	20:1	1:1	Generalization
(a) The overlap results of MBHB. Groups (iii) and (iv) correspond to the mean overlap values for $M_{\text{tot}} < 10^6 M_\odot$ and $M_{\text{tot}} \geq 10^6 M_\odot$, respectively.			
(i) Mean	0.981	0.990	0.970
(ii) Median	0.992	0.996	0.993
(iii) Mass $< 10^6 M_\odot$	0.980	0.979	0.938
(iv) Mass $\geq 10^6 M_\odot$	0.982	0.995	0.986
EMRIs	20:1	1:1	
(b) The overlap results of EMRIs. Groups (iii) and (iv) correspond to the mean overlap values for $e_0 < 0.1$ and $e_0 \geq 0.1$, respectively.			
(i) Mean	0.912	0.807	
(ii) Median	0.997	0.910	
(iii) $e_0 < 0.1$	0.962	0.905	
(iv) $e_0 \geq 0.1$	0.896	0.778	
GB	20:1	1:1	
(c) The overlap results of GB. Groups (iii) and (iv) correspond to the mean overlap values for $f < 10^{-3}$ and $f \geq 10^{-3}$, respectively.			
(i) Mean	0.991	0.992	
(ii) Median	0.996	0.994	
(iii) $f < 10^{-3}$	0.987	0.990	
(iv) $f \geq 10^{-3}$	0.995	0.993	

learning the characteristics associated with intermediate frequency signals.

In this scenario, the mean and median overlaps for EMRIs were found to be 0.807 and 0.910, while for GB, the mean and median overlaps were 0.992 and 0.994, respectively.

The performance impact was negligible for GB, but there was a significant decrease in EMRI waveforms. This can be attributed to the larger eccentricity and wider range of scales exhibited by EMRIs, as well as their continuous periodic transitions. Because of the high complexity of EMRI waveforms, shorter waveforms fail to capture the waveform features. Therefore, in the case of complex waveforms, CBS-GPT requires longer input waveforms to learn more distinctive features.

C. Interpretability

The attention map [Eq. (24)] allows us to understand the extrapolation process and attention mechanism while forecasting waveforms, making it easier to gain insight into how CBS-GPT interprets GW data:

$$A = \frac{1}{H} \sum_{i=1}^H \text{softmax} \left(\frac{Q_{ji} K_{ji}^T \cdot \text{mask}}{\sqrt{d}} \right), \quad (24)$$

where H represents all attention heads of the last encoder block. In Fig. 6, the vertical axis represents the model input waveform, and the horizontal axis represents the predicted waveform.

When predicting continuous gravitational waveforms (EMRIs and GB), the attention maps [Figs. 6(d)–6(i)] exhibit gridlike patterns that are closely related to the phase of the waveforms, with the scale of the grid expanding as the frequency decreases. In order to measure the similarity between the attention map and the input waveform, we introduce the correlation coefficient (with details described in the Appendix). Overall, the average correlation coefficient of the continuous waveform exceeds 0.8, which demonstrates that the model can accurately match the waveform’s frequency and phase information. This mode assists CBS-GPT in successfully extrapolating waveforms.

As showcased in Figs. 6(a)–6(c), during the prediction of the merge-ringdown phase of MBHB waveforms, attention primarily focuses on near-diagonal elements. In contrast to continuous GW signals, the amplitude of MBHB reaches zero after the merge ringdown, and the main focus of attention mechanism lies in the merging stage and the stage after the merge, with relatively less attention paid to inspiral phase.

D. Potential applications

Complex waveform generation. Currently, waveform generation for high mass ratio binary black holes remains a challenging problem because of high computational cost. Our approach can partially alleviate this problem since CBS-GPT that trained on low mass ratio waveforms with relatively low computational cost can be applied to high mass ratio waveform generation. This generalization characteristic, as shown in Fig. 5(g), demonstrates that the model can learn intrinsic features and can be applied to waveform extrapolation of a broader parameter space. By incorporating simulations based on numerical relativity, we may build a waveform template bank by extrapolating more complex and computation-intensive waveforms. For burst wave sources such as MBHB, waveform generation time of CBS-GPT for a single waveform is less than 100 ms on a single NVIDIA V100 GPU. With the rapid development of GPU computing power, CBS-GPT presents the potential for high-speed template waveform generation.

Gap imputation. In space-based GW detectors, the presence of data gaps due to data transmission, satellite attitude adjustments, and unidentified glitches can significantly impact the precision of waveform parameter estimation. Our waveform extrapolation method is promising to accomplish the task of waveform imputation, and by integrating with successive denoising models [33–37,37,38,60], parameter estimation accuracy can be further enhanced [61].

Model design guidance. We established a more convenient method for visualizing and quantifying attention maps,

offering guidance for transformer-based models design in the GW research realm. Our results also demonstrate that attention mechanism can be leveraged to establish more robust deep learning models that are specifically tailored for GW astronomy.

IV. CONCLUSION

In this paper, we introduce the CBS-GPT model, consisting of hybrid embedding and encoder blocks. The CBS-GPT is applied to predict GW waveforms after the TDI 2.0 response. We investigated two scenarios of different extrapolation ratios between input and predicted waveform length. Different models are trained for MBHB, EMRIs, and GB. In the 20:1 and 1:1 extrapolation scenarios, the average overlaps between the predicted waveform and the target waveform of MBHB, EMRIs, and GB reach 0.981, 0.912, 0.991, and 0.990, 0.807, 0.991, respectively. EMRIs exhibited poorer performance in the 1:1 extrapolation due to their complex waveform patterns and rich amplitude variations caused by eccentricity. We also proved the strong generalization of CBS-GPT on MBHB waveforms.

Moreover, we introduced a correlation coefficient and found that the correlation between hidden parameters of CBS-GPT and waveform was relatively high, which indicated that the model could learn the waveform’s phase information extremely well. Overall, our results show that CBS-GPT has the ability to comprehend detailed waveform properties and make predictions over varied frequencies. We are confident that, in the future, large AI models such as CBS-GPT can be applied to GW data processing tasks including complex waveform generation and gap imputation.

ACKNOWLEDGMENTS

This research was supported by the Peng Cheng Laboratory and Peng Cheng Cloud-Brain. This work was also supported in part by the National Key Research and Development Program of China Grant No. 2021YFC2203001 and in part by the NSFC (No. 11920101003 and No. 12021003). Z.C. was supported by the “Interdisciplinary Research Funds of Beijing Normal University” and CAS Project for Young Scientists in Basic Research YSBR-006.

APPENDIX: CORRELATION COEFFICIENT BETWEEN WAVEFORM AND HIDDEN PARAMETERS

To evaluate the correlation between the attention map’s gridlike pattern and the waveform, we introduce the correlation coefficient between the waveform and hidden parameters (or attention map). This coefficient assesses the level of correlation and demonstrates the attention map’s ability to capture phase information. First, we compute the mean value of each token of the patched waveform to get

the sequence M . Subsequently, the outer product of M is computed, resulting in the autocorrelation matrix. As the attention map A [Eq. (24)] is processed by masking and normalization, we do a similar adjustment to the autocorrelation matrix:

$$\begin{aligned} R_{\text{mask}} &= \text{Mask}(M \otimes M - \min(M \otimes M)), \\ R_{\text{Norm}} &= \text{RowNorm}(R_{\text{mask}}), \end{aligned} \quad (\text{A1})$$

where $\text{RowNorm}(\cdot)$ denotes the normalization of each row of the matrix and $\text{Mask}(\cdot)$ is consistent with the mask method of Sec. II C.

To assess the correlation between the two matrices, we calculate the Pearson correlation coefficient between the

flattened attention map A and flattened R_{Norm} :

$$\begin{aligned} \rho_{A, R_{\text{Norm}}} &= \rho\{\text{Flatten}(A), \text{Flatten}(R_{\text{Norm}})\} \\ &= \frac{n \sum_1^n A_{Fi} R_{Fi} - \sum_1^n A_{Fi} \sum_1^n R_{Fi}}{\sqrt{n \sum_1^n A_{Fi}^2 - (\sum_1^n A_{Fi})^2} \sqrt{n \sum_1^n R_{Fi}^2 - (\sum_1^n R_{Fi})^2}}, \end{aligned} \quad (\text{A2})$$

where $\text{Flatten}(\cdot)$ denotes flattening the matrix into one dimension, A_F and R_F represent the flattened vector of A and R_{Norm} , respectively, and n represents the length after flattening. Finally, $\rho_{A, R_{\text{Norm}}}$ is defined as the correlation coefficient between waveform and hidden parameters.

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, *Phys. Rev. D* **93**, 122003 (2016).
- [3] B. F. Schutz, *Nature (London)* **323**, 310 (1986).
- [4] B. S. Sathyaprakash and B. F. Schutz, *Living Rev. Relativity* **12**, 2 (2009).
- [5] A. Klein, E. Barausse, A. Sesana, A. Petiteau, E. Berti, S. Babak, J. Gair, S. Aoudia, I. Hinder, F. Ohme *et al.*, *Phys. Rev. D* **93**, 024003 (2016).
- [6] N. Tamanini, C. Caprini, E. Barausse, A. Sesana, A. Klein, and A. Petiteau, *J. Cosmol. Astropart. Phys.* **04** (2016) 002.
- [7] N. Bartolo, C. Caprini, V. Domcke, D. G. Figueroa, J. Garcia-Bellido, M. C. Guzzetti, M. Liguori, S. Matarrese, M. Peloso, A. Petiteau *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2016) 026.
- [8] C. Caprini, M. Hindmarsh, S. Huber, T. Konstandin, J. Kozaczuk, G. Nardini, J. M. No, A. Petiteau, P. Schwaller, G. Servant *et al.*, *J. Cosmol. Astropart. Phys.* **04** (2016) 001.
- [9] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
- [10] K. G. Arun, E. Belgacem, R. Benkel, L. Bernard, E. Berti, G. Bertone, M. Besancon, D. Blas, C. G. Böhm, R. Brito *et al.*, *Living Rev. Relativity* **25**, 4 (2022).
- [11] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration), *Phys. Rev. X* **13**, 041039 (2023).
- [12] G. Agazie, A. Anumarlapudi, A. M. Archibald, Z. Arzoumanian, P. T. Baker, B. Bécsy, L. Blecha, A. Brazier, P. R. Brook, S. Burke-Spolaor *et al.*, *Astrophys. J. Lett.* **951**, L8 (2023).
- [13] H. Xu, S. Chen, Y. Guo, J. Jiang, B. Wang, J. Xu, Z. Xue, R. N. Caballero, J. Yuan, Y. Xu *et al.*, *Res. Astron. Astrophys.* **23**, 075024 (2023).
- [14] EPTA Collaboration and InPTA Collaboration; J. Antoniadis, P. Arumugam, S. Arumugam, S. Babak, M. Bagchi, A.-S. Bak Nielsen, C. G. Bassa, A. Bathula, A. Berthereau *et al.*, *Astron. Astrophys.* **678**, A50 (2023).
- [15] D. J. Reardon, A. Zic, R. M. Shannon, G. B. Hobbs, M. Bailes, V. Di Marco, A. Kapur, A. F. Rogers, E. Thrane, J. Askew *et al.*, *Astrophys. J. Lett.* **951**, L6 (2023).
- [16] M. Bailes, B. K. Berger, P. R. Brady, M. Branchesi, K. Danzmann, M. Evans, K. Holley-Bockelmann, B. R. Iyer, T. Kajita, S. Katsanevas *et al.*, *Nat. Rev. Phys.* **3**, 344 (2021).
- [17] F. Matichard, B. Lantz, R. Mittleman, K. Mason, J. Kissel, B. Abbott, S. Biscans, J. McIver, R. Abbott, S. Abbott *et al.*, *Classical Quantum Gravity* **32**, 185003 (2015).
- [18] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bortoluzzi *et al.*, *arXiv:1702.00786*.
- [19] W.-R. Hu and Y.-L. Wu, *Natl. Sci. Rev.* **4**, 685 (2017).
- [20] Z. Ren, T. Zhao, Z. Cao, Z.-K. Guo, W.-B. Han, H.-B. Jin, and Y.-L. Wu, *Front. Phys.* **18**, 64302 (2023).
- [21] J. Luo, L.-S. Chen, H.-Z. Duan, Y.-G. Gong, S. Hu, J. Ji, Q. Liu, J. Mei, V. Milyukov, M. Sazhin *et al.*, *Classical Quantum Gravity* **33**, 035010 (2016).
- [22] H. Wang, S. Wu, Z. Cao, X. Liu, and J.-Y. Zhu, *Phys. Rev. D* **101**, 104003 (2020).
- [23] D. George and E. Huerta, *Phys. Lett. B* **778**, 64 (2018).
- [24] D. George and E. A. Huerta, *Phys. Rev. D* **97**, 044039 (2018).
- [25] H. Gabbard, M. Williams, F. Hayes, and C. Messenger, *Phys. Rev. Lett.* **120**, 141103 (2018).
- [26] C. Badger, K. Martinovic, A. Torres-Forné, M. Sakellariadou, and J. A. Font, *Phys. Rev. Lett.* **130**, 091401 (2023).
- [27] X.-T. Zhang, C. Messenger, N. Korsakova, M. L. Chan, Y.-M. Hu, and J.-d. Zhang, *Phys. Rev. D* **105**, 123027 (2022).
- [28] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [29] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, *Nat. Phys.* **18**, 112 (2022).

- [30] H. Wang, Z. Cao, Y. Zhou, Z.-K. Guo, and Z. Ren, *Big Data Mining and Anal.* **5**, 53 (2022).
- [31] W. Wei and E. Huerta, *Phys. Lett. B* **800**, 135081 (2020).
- [32] R. E. Colgan, K. R. Corley, Y. Lau, I. Bartos, J. N. Wright, Z. Márka, and S. Márka, *Phys. Rev. D* **101**, 102003 (2020).
- [33] A. Torres-Forné, A. Marquina, J. A. Font, and J. M. Ibáñez, *Phys. Rev. D* **94**, 124040 (2016).
- [34] A. Akhshi, H. Alimohammadi, S. Baghran, S. Rahvar, M. R. R. Tabar, and H. Arfaei, *Sci. Rep.* **11**, 20507 (2021).
- [35] Z. Ren, H. Wang, Y. Zhou, Z.-K. Guo, and Z. Cao, [arXiv:2212.14283](https://arxiv.org/abs/2212.14283).
- [36] T. Zhao, R. Lyu, H. Wang, Z. Cao, and Z. Ren, *Commun. Phys.* **6**, 212 (2023).
- [37] W. Wei and E. Huerta, *Phys. Lett. B* **800**, 135081 (2020).
- [38] C. Chatterjee, L. Wen, F. Diakogiannis, and K. Vinsen, *Phys. Rev. D* **104**, 064046 (2021).
- [39] J. Lee, S. H. Oh, K. Kim, G. Cho, J. J. Oh, E. J. Son, and H. M. Lee, *Phys. Rev. D* **103**, 123023 (2021).
- [40] A. Khan, E. A. Huerta, and H. Zheng, *Phys. Rev. D* **105**, 024024 (2022).
- [41] A. J. K. Chua, C. R. Galley, and M. Vallisneri, *Phys. Rev. Lett.* **122**, 211101 (2019).
- [42] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [43] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, [arXiv:2005.14165](https://arxiv.org/abs/2005.14165).
- [44] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, [arXiv:2103.14030](https://arxiv.org/abs/2103.14030).
- [45] L. P. Collaboration, *Phys. Rev. D* **106**, 062001 (2022).
- [46] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond *et al.*, *Phys. Rev. D* **95**, 044028 (2017).
- [47] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
- [48] A. Maselli, N. Franchini, L. Gualtieri, T. P. Sotiriou, S. Barsanti, and P. Pani, *Nat. Astron.* **6**, 464 (2022).
- [49] P. Amaro-Seoane, J. Andrews, M. Arca Sedda, A. Askar, Q. Baghi, R. Balasov, I. Bartos, S. S. Bavera, J. Bellovary, C. P. L. Berry *et al.*, *Living Rev. Relativity* **26**, 2 (2023).
- [50] M. L. Katz, A. J. K. Chua, L. Speri, N. Warburton, and S. A. Hughes, *Phys. Rev. D* **104**, 064047 (2021).
- [51] X.-H. Zhang, S.-D. Zhao, S. D. Mohanty, and Y.-X. Liu, *Phys. Rev. D* **106**, 102004 (2022).
- [52] M. L. Katz, J.-B. Bayle, A. J. K. Chua, and M. Vallisneri, *Phys. Rev. D* **106**, 103001 (2022).
- [53] M. Tinto, F. B. Estabrook, and J. W. Armstrong, *Phys. Rev. D* **69**, 082001 (2004).
- [54] M. Tinto and S. V. Dhurandhar, *Living Rev. Relativity* **17**, 6 (2014).
- [55] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, *Advances in Neural Information Processing Systems* (Curran Associates, Inc, 2017).
- [56] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [57] L. Barack, V. Cardoso, S. Nissanke, T. P. Sotiriou, A. Askar, C. Belczynski, G. Bertone, E. Bon, D. Blas, R. Brito *et al.*, *Classical Quantum Gravity* **36**, 143001 (2019).
- [58] S. L. Larson, W. A. Hiscock, and R. W. Hellings, *Phys. Rev. D* **62**, 062001 (2000).
- [59] S. Babak, M. Hewitson, and A. Petiteau, [arXiv:2108.01167](https://arxiv.org/abs/2108.01167).
- [60] R. Ormiston, T. Nguyen, M. Coughlin, R. X. Adhikari, and E. Katsavounidis, *Phys. Rev. Res.* **2**, 033066 (2020).
- [61] K. Dey, N. Karnesis, A. Toubiana, E. Barausse, N. Korsakova, Q. Baghi, and S. Basak, *Phys. Rev. D* **104**, 044035 (2021).