

Noise reduction in gravitational-wave data via deep learning

Rich Ormiston,¹ Tri Nguyen^{1,2}, Michael Coughlin^{1,3}, Rana X. Adhikari³, and Erik Katsavounidis²

¹*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

²*LIGO Laboratory, Massachusetts Institute of Technology, 185 Albany Street, Cambridge, Massachusetts 02139, USA*

³*Division of Physics, Math, and Astronomy, California Institute of Technology, Pasadena, California 91125, USA*



(Received 28 May 2020; accepted 19 June 2020; published 14 July 2020)

With the advent of gravitational-wave astronomy, techniques to extend the reach of gravitational-wave detectors are desired. In addition to the stellar-mass black hole and neutron star mergers already detected, many more are below the surface of the noise, available for detection if the noise is reduced enough. Our method (DeepClean) applies machine-learning algorithms to gravitational-wave detector data and data from on-site sensors monitoring the instrument to reduce the noise in the time series due to instrumental artifacts and environmental contamination. This framework is generic enough to subtract linear, nonlinear, and nonstationary coupling mechanisms. It may also provide handles in learning about the mechanisms which are not currently understood to be limiting detector sensitivities. The robustness of the noise-reduction technique in its ability to efficiently remove noise with no unintended effects on gravitational-wave signals is also addressed through software signal injection and parameter estimation of the recovered signal. It is shown that the optimal signal-to-noise ratio (SNR) of the injected signal is enhanced by $\approx 21.6\%$ and the recovered parameters are consistent with the injected set. We present the performance of this algorithm on linear and nonlinear noise sources and discuss its impact on astrophysical searches by gravitational-wave detectors.

DOI: [10.1103/PhysRevResearch.2.033066](https://doi.org/10.1103/PhysRevResearch.2.033066)

I. INTRODUCTION

The recent detections of gravitational waves from binary systems (see Ref. [1] for a summary of the first two observing runs) motivates technological and data analysis improvements to extend the reach of current gravitational-wave detectors. The current network consists of the two advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) interferometers in the United States [2], the Advanced Virgo (adVirgo) interferometer in Italy [3], the GEO-HF interferometer in Germany [4], the Kamioka Gravitational-Wave Detector (KAGRA) interferometer in Japan [5], and eventually the LIGO-India detector in India [6]. Identification of gravitational-wave events from binary systems is subject to transient and periodic noise sources in gravitational-wave instruments. Such noise sources may limit the significance of gravitational wave events and reduce them to a subthreshold level [1,7]. The ability to reduce noise in the instruments, thus enhancing their sensitivity, can enable identification of additional events that would otherwise remain below the threshold.

The ultimate sensitivity of the aLIGO detectors is dictated by the physics inherent to their design, such as shot noise of the laser light or thermal fluctuations of the mirror coatings

and optic suspensions [2]. However, the performance during the observation runs is also influenced by environmental and technical noises which arise from factors such as earthquakes and the instrumentation or control of the interferometer, respectively [8]. The confidence in the significance of any given signal and our ability to extract astrophysical information from it is directly impacted by the noise and sensitivity of the detectors at the time, and so there is a strong need to improve their performance by any available means.

In general, the performance of a single detector is characterized by separately considering the plethora of mechanisms by which nonastrophysical noise sources couple into the strain output of the instrument, such as the shot noise of the light incident on the output photodiodes or thermal motion of the arm cavity mirror surfaces. Once categorized into causally distinct groups, we can predict the instrument's performance from the incoherent sum of these noise mechanisms and compare it to the observed steady-state sensitivity. This is a crucial analysis when working to understand and improve the performance, as a diagnosis of what aspects or subsystems of the detector are the limiting factors. It also shows us where the observed noise exceeds the sum of the budgeted noise sources, and thereby where our understanding of the noise is incomplete.

In this analysis, we concentrate on the frequencies of less than 1000 Hz since this range contains unexplained features. At those greater than 1000 Hz, the detector is generally well understood, and the features mostly correspond to the shot noise at the output photodiode or harmonic lines that are not well witnessed by environmental sensors. At frequencies of less than 100 Hz, there exists some amount of noise that

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Estimate of the noise budget (black) and actual measured strain (red) at the Hanford interferometer for LIGO's second observation run, O2. The strain of the differential arm length (DARM) was calculated using data from July 4, 2017 [10].

remains unexplained, as shown in Fig. 1. There is also significant contribution from technical noise arising from the control of the suspended optics. This is due in part to the control actuation necessary to keep the instrument well aligned over long periods of time and the gradual shifts in the beam spot positions. In addition, the low-frequency sensitivity of the detectors are especially important for detecting high-mass binary mergers because these systems merge at low frequencies. For example, the recorded signal from GW150914 only spent about 200 ms in the sensitive band of the instrument and was resolvable from about 35–250 Hz [9].

II. BACKGROUND

The aLIGO detectors employ numerous subsystems that control different aspects of the instrument and monitor its state. These are coordinated and operated in large part by a distributed digital control system which measures and records a large numbers of signals related to these subsystems, in addition to the main output which measures space-time strain. Thus, numerous signals are synchronously recorded along with the interferometer output, such as those from environmental sensors, mirror suspension actuation, and photodetectors. These auxiliary signals have the potential to witness coupling of unwanted noises into the interferometer and are used in the commissioning of the detector to diagnose and mitigate such couplings.

Seismometer signals have also been used to train feed-forward subtraction filters that are run in real-time to reduce the physical motion of interferometer elements [11]. This manner of online subtraction has the strong benefit of reducing the gain or dynamic range requirements of the length and/or angular feedback systems. However, gradual changes in instrument state, such as alignment or thermal state, can cause changes in the expected couplings during an observation run, when it is preferred to make as few configuration changes to the instrument as possible. This may lead to unwanted noise making its way to the recorded strain data, despite the necessary information required to subtract it being available. Furthermore, there is the possibility of unconsidered noise couplings being present that could in principle be predicted

from other recorded signals. At this point, the only recourse is to revisit previously recorded data and attempt to regress the unwanted noise out.

One technique for reducing the noise in the strain signal *post facto* using auxiliary information is Wiener filtering [12–15], a multiple-input single-output (MISO) algorithm which optimizes the mean squared difference between the subtraction target and the predicted coupled noise from multiple witnesses, taking the correlations between the witnesses themselves into account. Time domain Wiener filtering has been used successfully in terrestrial gravitational wave detectors to enhance the performance of the vibration isolation system [11,16] and reduce the influence of local gravitational field fluctuations [17,18].

Naturally, linear couplings of external disturbances into the gravitational-wave strain readout are a subset of the full dynamics of the detectors. It is therefore worthwhile to extend these regression techniques to more complicated nonlinear and nonstationary couplings. There are many known nonlinear coupling mechanisms, and it is likely that more exist which have not been fully accounted for [19]. The functional forms can vary greatly and even modest uncertainty in the parameters involved can make it impractical to reconstruct and regress the unwanted noise. Machine learning techniques have shown great promise at improving the sensitivity of gravitational-wave data analysis. For example, Gravity Spy combines crowd sourcing with machine learning to aid in the challenging task of categorizing all of the instrumental data transients recorded by the gravitational-wave detectors [20]. Other algorithms have risen to characterize both data transients [21,22] and improve gravitational-wave searches [23–26]. This success has prompted work into developing techniques for performing linear, nonlinear, and nonstationary regression in the interferometer data that does not require precise *a priori* knowledge of all of the system parameters.

In this paper, we describe the DeepClean algorithm in which auxiliary signals are used *post facto* to estimate noise couplings that existed during the recent science runs using machine-learning techniques. We can use simulated gravitational-wave events to characterize the performance of the algorithm and compare the performance of the noise subtraction to that of Wiener filters. We describe how the algorithms were validated to not corrupt or bias the resultant estimates of the astrophysical source parameters. Finally, we detail the sensitivity improvement that results from this subtraction and the consequent improvement in the confidence in the estimates of the source parameters of select detections.

III. NOISE SUBTRACTION PIPELINE

In this section, we present the analysis pipeline applied onto data for the purpose of noise subtraction. The method processes time series corresponding to the gravitational-wave strain data $h(t)$ and a set of auxiliary (“witness”) channels $w_i(t)$. The witness channels may be physical environmental monitors (PEM) or auxiliary interferometric channels that contain information about the witnessed noise and *not* the astrophysical signal. This is a critical convenience afforded to us by our confidence that true astrophysical signals are uniquely present in the main readout signal of the




FIG. 2. Workflow diagram of the noise subtraction pipeline.

interferometer; any attempted noise subtraction from a combination of witness signals may only increase or reduce the influence of noise terms and cannot fundamentally alter any present astrophysical signals. The algorithm employs a one-dimensional convolutional neural network (CNN) which takes in a user-specified set of witness channels at one time and subsequently outputs the predicted noise in $h(t)$. Input witness channels are conditioned before being fed into the CNN. The output from the CNN (i.e., the predicted noise) is also conditioned before being subtracted from $h(t)$. Figure 2 shows the schematic of the noise subtraction pipeline. In what follows, we will discuss the details of the algorithm implemented for this analysis.

A. Formalism and loss function

The gravitational-wave strain $h(t)$ as reconstructed by the gravitational-wave detectors may be written as

$$h(t) = s(t) + n(t), \quad (1)$$

where $s(t)$ is the astrophysical signal that may be present in the data and $n(t)$ is the noise in the detector. We may further subdivide the noise into

$$n(t) = n_w(t) + n_{nw}(t), \quad (2)$$

where $n_w(t)$ represents noise sources coupled into the witness channels $w_i(t)$, and $n_{nw}(t)$ represents noise sources we do not intend to subtract, which include nonremovable (fundamental) noise (e.g., quantum noise, photon shot noise) and noise sources not witnessed by $w_i(t)$.

We design the neural network to take in the witness channels and produce an estimate of the witnessed noise,

which may then be filtered from the gravitational-wave strain data. The neural network can be represented as a function $\mathcal{F}(w_i(t); \vec{\theta})$ which maps the witness channels $w_i(t)$ to the strain output given a set of parameters $\vec{\theta}$. The parameters $\vec{\theta}_i$ are obtained by minimizing a loss function J which quantifies the difference between the predicted noise and the real witnessed noise. Mathematically, the problem may be stated as

$$\vec{\theta} = \arg\min_{\vec{\theta}} J[h(t), \mathcal{F}(w_i(t); \vec{\theta})]. \quad (3)$$

In our analysis, we choose the loss function to be the weighted average of the amplitude spectrum density (ASD) of the residual strain $r(t)$ (i.e. after subtracting the predicted noise) over a frequency interval $[f_1, f_2]$. In other words,

$$J_{\text{asd}} = \frac{1}{f_2 - f_1} \int_{f_1}^{f_2} W(f) \sqrt{S[r, r](f)} df \quad (4)$$

$$r(t) = h(t) - \mathcal{F}(w_i(t); \vec{\theta}), \quad (5)$$

where $W(f)$ is a frequency-dependent weighting function. As in Ref. [27], we choose $W(f)$ to be the reciprocal of the ASD of the target strain $\sqrt{S[h, h](f)}$. Since the ASD typically spans several orders of magnitude and convergence is dependent on the span of the eigenvalues of the correlation matrices of the input [28], this has a whitening effect on the ASD and ensures noisy frequency bins do not dominate the loss function. We also make a modification to set $W(f)$ to be zero at frequencies outside the witnessed noise. This helps the network converge faster and to a more stable solution, especially in cases where the noise source couples into the gravitational-wave strain at multiple frequencies (e.g., the 60-Hz power mains and its harmonics). In the discrete time-series notation, the ASD loss

function [Eq. (4)] can be written as

$$J_{\text{asd}} = \frac{1}{M} \sum_{i=0}^{M-1} \sqrt{\frac{S[r, r][i]}{S[h, h][i]}}, \quad (6)$$

where M is the number of frequency bins.

In addition to the frequency-domain loss function in Eq. (6), we find adding a time-domain loss function further improves subtraction power on spectral lines. We choose the time-domain component of the loss function to be the mean squared error (MSE) across the time series:

$$J_{\text{mse}} = \frac{1}{N} \sum_{i=0}^{N-1} r[i]^2, \quad (7)$$

where N is the number of time-series samples. Combining with the ASD loss in Eq. (6), the final loss function can be written as

$$J = wJ_{\text{asd}} + (1 - w)J_{\text{mse}}, \quad (8)$$

where w is a weighting factor that goes from 0 to 1. It is a hyperparameter which can be adjusted depending on the data set. In general, we found that high values of w perform better on broadband noise, and lower values of w performs better on spectral line noise.

In practice, the MSE loss J_{mse} is calculated from the full discrete time series, while the ASD in the ASD loss J_{asd} is estimated using Welch's method, i.e., averaging over the discrete Fourier transform (DFT) of overlapping time-series segments. The number of frequency bins M is chosen to be four to eight times smaller than the number of time-series samples N in order to obtain a nonbiased estimate of the ASD. Because M is proportional to the frequency resolution, we also make sure N is large enough to reach a frequency resolution of at least 0.5 Hz. M and N remain important hyperparameters to be explored in further studies.

B. Data preprocessing

The witness channels are preprocessed before being fed into the neural network. Because we design the neural network to take in a time series and output time series of the same length, we first resample all witness channels, with appropriate anti-imaging filtering, to the same sampling frequency as the strain. By doing so, we also ensure the input and predicted time-series arrays have coincident start and stop times. In order to save computational time, the sampling frequency of the strain is chosen such that the Nyquist frequency is just above the frequency of the witnessed noise.

Next, we apply an eighth-order Butterworth bandpass filter to mitigate the power outside of the frequency band of the witnessed noise. For nonlinear couplings, the frequency of the noise coupled into the witness channels can be different from the frequency coupled into the strain channel. To account for this, we only bandpass the strain and not the input witnesses.

In machine-learning problems, it is good practice to normalize the data to have zero mean and unit variance. This scaling ensures that any one witness channel does not account for the majority of the error in the loss function, which would prevent the network from correctly learning the coupling of the other channels. In addition, because the ASD of LIGO

strain has a magnitude of order 10^{-23} to 10^{-21} Hz^{-1} , the normalization helps prevent numerical instability when taking the ratio in Eq. (6). We normalize both the strain channel and the witness channels. For each channel, we compute the mean and standard deviation across the time series. All data samples in the time series are then subtracted by the mean and divided by the standard deviation. This ensures normalization does not add any unwanted features to the time series. The procedure is also invertible, so the network's prediction can later be easily converted back into real physical units.

To help the network learn the noise coupling more efficiently, we then divide each time series into smaller overlapping segments. Each training sample consists of segments from multiple witness channels. In machine-learning literature, this step is known as data augmentation. In our analysis, we choose a segment length to be 8 s. This choice is motivated by the frequency resolution of the discrete ASD in Eq. (6). To achieve a resolution of 0.5 Hz, we set the DFT length to 2 s. Using Welch's method with a 1-s DFT overlap, we require the segment length to be at least 8 s. To increase the frequency resolution, we would have to increase the DFT length, the segment length, and by extension, the computational resources of the algorithm in both time and memory. For the noise couplings presented in the results of this paper, we find the length of 8 s gives the best performance in both subtraction power and run time. We choose the overlap duration between segments to be 7.75 s (96.875%) during training and 4 s (50%) during inference. Increasing the overlap duration increases the size of the data set (and thus the computational resources) but also allows the network to make more connections and therefore characterize variants of the noise sources more effectively. The overlap duration during inference is chosen to be smaller than during training to save computational time.

C. Neural network architecture

As mentioned earlier the section, the algorithm employs a one-dimensional CNN which takes in a set of witness channels and predicts the contributed noise in the gravitational wave strain. The typical CNN consists of a number of convolutional layers, each employing a set of discrete window functions, or kernels, with trainable weights. Each layer takes in an input series, which may have multiple channels. During forward pass, each layer slides its kernel across the input and computes the dot product across all channels between the kernels and the input series within the kernel. This locality helps the CNN learn short-term features in the data. The output of each layer is then passed through some nonlinear activation function and becomes the input of the subsequent layer. Because the output of each layer is downsampled, each subsequent layer sees exponentially more of the network's input. This enables deep CNNs (with more layers) to "remember" longer series and learn long-term features in the data. Because the convolutional operator by definition has a built-in spatial (temporal) invariance, the CNN can also detect repeating features in the series, making it suitable to process highly periodic time series. In addition to convolutional layers, we employ transposed convolutional layers in our CNN. Transposed convolutional layers [29] work by exchanging the input and output of convolutional layers. In

other words, they distribute each sample of the input by the weights of the kernels and then add the resulting segments elementwise. As a result, the output of the transposed convolution operator is upsampled instead of being downsampled like in the convolution operator. The upshot is that the output of the network will be the same size as the input, making the subtraction straightforward, and the transposed convolution operator adds a layer of weights which a fully connected layer or traditional upsampling cannot.

Given the above motivations, we employ a fully convolutional autoencoder to map the witness channels to the predicted noise. In more detail, the input witnesses are first passed through an input convolutional layer which extracts their features and maps them to a set of output channels. To preserve the length of the time series, the input layer uses a stride of 1 and applies an appropriate zero padding scheme to the input witnesses. After each layer, the series length is reduced by a factor of 2, and the number of channels increases by a factor of 2 to preserve the time complexity of the layer. To downsample the series, instead of using pooling layers (e.g., max pooling), each convolutional layer uses a stride of 2 with an appropriate zero padding scheme. The output of the downsampling layers is then passed through a series of transposed convolutional layers. After each transposed layer, the series length increases by a factor of 2, and the number of channels is reduced by a factor of 2. This is done by using transposed layers with a stride of 2 and an appropriate zero padding scheme. The output of the transposed convolutional layer is then passed through an output convolutional layer to combine all features into the predicted noise in the gravitational wave strain. Except in the last layer, the output of each layer is normalized using batch normalization [30] and passed through an activation function before going to the subsequent layer.

The symmetrical architecture of the network is motivated by the common knowledge that each convolutional layer learns a different feature level of the input series; while the earlier layer of the CNN learns the low-level features, the deeper layer learns more advanced, high-level features. Therefore, we expect the first downsampling convolutional layer to extract low-level features of the witness channels while the last transposed convolutional layer reconstructs low-level features of the predicted noise. Similarly, the intermediate layers will learn the intermediate-level features, and the last down-sampling convolutional layer and the first transposed convolutional layer will learn the high-level features.

In our analysis, we use three convolutional layers for downsampling (not including the input and output convolutional layers) and three transposed convolutional layers for upsampling. We set the kernel size of all layers to be seven. The numbers of output channels of each layer (including the input and output layers) are 8, 8, 16, 32, 32, 16, 8, 1, respectively. For the activation function, we use Tanh. The input and output dimensions of each layer then depend only on the dimension of the input, i.e., the number of witness channels and length of the time series, which we vary depends on the noise couplings. An example of the network architecture is shown in the Appendix. We have tried different network architectures and other activations such as ReLU and Sigmoid and found that they give similar results.

D. Training and inference

The pipeline can be divided into two parts: training and inference. During training, data are fed into the network by minibatches, wherein each batch consists of 32 training samples. The network then computes the loss function in Eq. (8) (averaging over the minibatch) and its gradient with respect to the network's parameters $\bar{\theta}$. The parameters are updated accordingly using a first-order stochastic gradient descent algorithm. Training terminates when gradient descent converges, which typically takes about 5–10 iterations over the training data ("epochs"), or when it reaches a set maximum number of 50 epochs. The gradient descent algorithm used in our analysis is ADAM [31] with default hyperparameters. Also, we find reducing the learning rate of ADAM by a factor of 10 every five epochs helps the network converge to a lower loss value and improves the subtraction. During inference, the network takes in the witness channels and produces the predicted noise. This is processed further using the procedure described in the subsequent section before being subtracted from the gravitational wave strain. To prevent overfitting, which occurs when the network learns features presented uniquely in the training data without being able to generalize to a broader data set, the training data do not overlap with the inference data. In addition, at the end of every epoch, we compare the loss on the training and inference set and stop the gradient descent if the network is trading performance on the inference set for performance on the training set. The duration of the training data is greatly dependent on the properties of the witnessed noise, such as the complexity of its coupling function to the gravitational-wave strain. More complex coupling functions require longer training data.

Because noise coupling functions in LIGO are typically nonstationary, we design our network to be small such that frequent retraining does not take a significant amount of run time. Using the architecture described above, the training time on 300–1024 s of training data (using 8-s segments with 7.75-s overlapping) takes about 2–6 min (including data preprocessing) on an NVIDIA TITAN X (Pascal) GPU. Once trained, inference on 1024–3600 s of data takes a few seconds on the same GPU. Because both the training and inference time are much less than the duration of the corresponding data set, the algorithm can be applied for both offline and real-time subtraction. Because the algorithm takes in raw data with no featurization, it is also easy to implement in real time. We expect the training time to scale linearly with the duration of the training set. The number of trainable parameters in the network also contributes significantly to the training time. Because the input layer maps the witness channels to a fixed-dimensional subspace, increasing the number of witness channels only increases the run time of the input layer, which is a small fraction of the run time of the network. However, whenever possible, we do not include irrelevant witness channels into the input because they hinder gradient descent convergence and possibly add uncertainty to the subtraction.

E. Output data postprocessing

As mentioned earlier, the output of the neural network is processed before being subtracted from the original target

strain. During inference, the network takes in the witness channels in 8-s segments with a 4-s (50%) overlap between segments and predicts the noise in the strain. We apply a Hann window to each segment before adding them together to reduce edge effects.

Because the network is trained on the normalized strain, the predicted noise will be in the dimensionless units. We convert the noise back to the gravitational-wave strain unit by scaling all data samples up by the standard deviation and adding the mean computed in the normalization step during data preprocessing.

During training, the weighting function $W(f)$ in Eq. (4) is set to zero for frequencies outside the band of the witnessed noise. It is reasonable to assume that prediction made by the network at these frequencies will only add noise to the gravitational-wave strain. We therefore bandpass all these frequencies using an eighth-order Butterworth bandpass filter.

As mentioned earlier, to save computational time, we downsample the strain such that the Nyquist frequency is just above the frequency of the witnessed noise. The predicted noise will have the same sampling frequency as the sampling frequency chosen during training. Because this sampling frequency might not be optimal for detection and parameter estimation of astrophysical signals, we upsample the noise, with an appropriate anti-imaging filter, before subtracting it from the full bandwidth strain.

IV. PIPELINE PERFORMANCE ON LIGO DATA

We applied our noise subtraction pipeline to data collected by the LIGO detectors during their second and third observation runs (O2 and O3). We choose multiple instances of data collection corresponding to different types of couplings present in order to study and quantify the performance of our pipeline.

A. O2 jitter noise

We first explore subtraction on the calibrated output data from the LIGO Hanford detector during the second observation run. These data were subsequently cleaned using a Wiener filter to linearly regress the data over LIGO's second observation run [14]. We aim to compare our method with such Wiener filter subtraction that has been adopted by the LIGO search pipelines. The broadband linearly coupled noise in this analysis comes from fluctuations of the prestabilized laser beam in size and angle [32] and comprises the majority of the noise removed in the frequency band from $\approx 10^2$ to 10^3 Hz [10]. In addition to the beam jitter noise, narrowband spectral features such as the 60-Hz power line and its harmonics, as well as the interferometer's calibration lines are also removed. As input to the network we use an identical set of auxiliary channels, as in Ref. [14]. These channels are known to have strong linear couplings into the output strain




FIG. 3. Comparison of the subtraction performance between a Wiener filter used to create the cleaned LIGO frames (DCH-CLEAN_STRAIN_C02) and the DeepClean neural network shows nearly identical results. A Tanh activation function was used which does not limit the network to simple linear connections. The results demonstrate that DeepClean can reproduce the results of linear subtraction when provided with the same channels and without overfitting the data.




FIG. 4. The increase in the binary neutron star inspiral range after the linear subtraction is seen to be $\approx 20\%$.

data stream. However, we do not limit the network to strictly linear interactions. As mentioned in the previous section, the network instead utilizes a nonlinear Tanh activation function and therefore could in principle discern additional nonlinear couplings within these channels.

We present an example of the O2 linear subtraction in Figs. 3 and 4. In this example, the network is trained on the 1024 s of data starting at GPS time 1182410770 (2017-25-06 07:25:52 UTC) and subtracts on the 1024 s of data starting at GPS 1182411794 (2017-25-06 07:42:56 UTC). The preprocessing and postprocessing procedures, as well as the network architecture, are described in the previous section, with all witness and gravitational-wave strain channels resampled to a sampling frequency of 2048 Hz. In addition to these steps, we remove all spectral lines (power mains and calibration lines) before subtracting the broadband beam jitter noise. For each spectral line (or group of spectral lines, such as the calibration lines near 37 Hz), we train a different network (using the appropriate witness channels) and combine the outputs. Because the lines are at different frequencies, they can be removed simultaneously by parallelizing the training and inference process. To remove the broadband beam jitter noise, a final network is trained using the interferometer strain data with the spectral lines removed. Figure 3 shows the ASD(s) of LIGO Hanford before and after the subtraction, computed based on the inference data (and not the training data). The improvement of the BNS inspiral range due to this linear noise subtraction is shown in Fig. 4. The result of the Wiener filter subtraction from Ref. [14] is also shown for comparison. The range obtained from DeepClean is similar to within 1–2% of that from the Wiener filter result, suggesting that the network has learned the coefficients of the optimal MSE filter and captured physical couplings without overfitting or adding any additional noise.

B. O3 60-Hz sidebands

DeepClean is not limited to linear couplings unlike the Wiener filter. The nonlinear activation function allows the algorithm to learn arbitrarily high order couplings of the input data. One such example of nonlinear and nonstationary couplings is the modulation of a low-frequency signal from LIGO's alignment sensing and control (ASC) system with the 60-Hz line of the power mains. This coupling produces sidebands around the central frequency. In a previous work [27], the ASC system channels and a witness to the power mains have been used to subtract these sidebands during the third observation run, O3.

We have benchmarked our pipeline on the same data set to compare the two methods. Using the same set of witness channels as in Ref. [27], which uses an analytic modulated linear adaptive filter in the frequency domain, we show in Fig. 5 that our network is capable of removing nonlinear and nonstationary couplings such as these. In this example, the neural network is trained on 1024 s of data starting at GPS time 1243926522 (2019-07-06 07:08:24 UTC) and subtracts on the 1024 s of data at GPS time 1243927546 (2019-07-06 07:25:28 UTC). All channels are resampled to a sampling frequency of 1024 Hz. Note that the network subtracts both the linear coupling (central peak at 60 Hz) and the nonlinear and nonstationary coupling (sidebands) at the same time.

V. PARAMETER ESTIMATION AND NETWORK SAFETY

The procedure outlined above is carried out for each gravitational wave time-series separately, i.e., the ones from the LIGO detectors in Hanford and Livingston. If the performance of the trained filters does not add noise and is either consistent with known results from analytic methods or subtracts spectral features of the target channel in a manner consistent




FIG. 5. Since the network is not restricted to linear couplings, nonlinear noise such as the modulation of the 60-Hz power line by the ASC system may be effectively and safely removed. The nonstationary subtraction in Ref. [27] is also shown for comparison.

with the features of the witness channels, then those filters are safe to apply to the strain data during the time of interest. The result of the filters is the production of a new strain time series which should have increased fidelity to the gravitational-wave strain signal incident on the instruments. One way of assessing the ability of our method to denoise gravitational-wave time series is by invoking parameter estimation methods on a set of astrophysical signal wave forms that are injected via software and for signals in which the true astrophysical parameters are known *a priori*. In this way, we can test whether this noise subtraction scheme is legitimately reducing unwanted technical noise without distorting the measured gravitational-wave signals. We use the DeepClean algorithm to filter noise from a stretch of data which contains an astrophysical software injection. Then we check that the resultant posterior parameter estimation distributions are consistent with those from the presubtraction strain signal and not significantly biased away from the known injected parameters.

For this study, we injected nonspinning binary black hole (BBH) signals into the gravitational-wave strain and compared the recovered source parameters from the cleaned and uncleaned time series.

For the O2 beam jitter data set, we injected 10 nonspinning BBH signals. The signals are injected at GPS time 1182411606. As in Ref. [32], each BBH has component masses M_1 and M_2 sampled from a uniform distribution from $[28, 64]M_\odot$ with a mass ratio q constrained to $[0.125, 1.0]$. The sky coordinates and orientation are sampled isotropically,

and the luminosity distance D_L is sampled uniformly in comoving volume from $[70, 1540]$ Mpc. The optimal SNR of the injected signals is ranging from 1.50 to 18.7. In the O3 60-Hz dataset, we injected 12 nonspinning, high-mass BBH signals at GPS time 1243309096 and 1244006580. Each injection time has six signals, which have a mass ratio q and total mass $M = M_1 + M_2$ combination of (0.5, 1) and (100, 150, 300) $M_\odot$ respectively. The high total masses are chosen such that the signals have significant power at around 60 Hz. Similarly to the O2 injections, the sky coordinate and orientation of the signals are sampled isotropically, with the luminosity distance sampled uniformly in comoving volume from $[70, 1540]$ Mpc. This results in an optimal SNR range of 0.722 to 30.3 of the injected events. The wave-form model for both the O2 and O3 injections is generated from the IMRPhenomPv2 wave-form approximant [33–35]. We applied Bayesian statistics and estimated the posterior probability distribution of the source parameters using the dynamic nested sampling algorithm Dynesty [36] implemented in the Bilby library [37]. For this study, the posterior distribution was estimated using only the gravitational-wave strain from a single detector, i.e., LIGO Hanford. We only recovered the masses M_1, M_2 , inclination angle θ_{jn} , and the luminosity distance D_L of the each gravitational-wave signal. All quantities are reconstructed within 3σ of their true values. In the O2 beam jitter dataset, the 90% credible intervals of the reconstructed quantities shrink by approximately 7.25% when comparing them with the ones obtained from the original strain. In




FIG. 6. Corner plot showing the parameter estimation of the O2 data from the Hanford detector after cleaning the data with the DeepClean neural network and using the same auxiliary channel list as was used in the linear cleaning analyses.

addition, we observed an average increase in the recovered optimal SNR of about 21.6%. We attribute this to the improved noise spectrum our method provides. Figure 6 shows a posterior distribution recovered from an example injection

in the O2 linear dataset. In the O3 60-Hz dataset, we did not observe any significant decrease in the 90% credible intervals, or any substantial increase in the recovered SNR. This is to be expected given that subtracting only the 60-Hz line and its sidebands should not significantly change the ASD. In all injections, the parameters recovered from the cleaned strain were consistent with the true values and those recovered from the original strain, suggesting that the network did not introduce any noise or corrupt the gravitational-wave signals.

VI. CONCLUSIONS

Going forward, it is evident that noise regression efforts are worth pursuing further. In addition to analytic methods, neural networks such as DeepClean have the extended advantage of being able to determine linear, nonlinear, and nonstationary couplings into the detector output without previous knowledge of the physical mechanisms of the noise. The ability of the machine-learning algorithm to successfully subtract nonlinear couplings allows for network-derived filters to become a more valuable as Advanced LIGO and future detectors become increasingly sensitive to additional, more complicated noise sources and the hardware engineering limit is approached. Given the great cost associated with the design, construction, commissioning, and analysis of the LIGO interferometers, being able to reliably improve the data quality through semiautomated processes will ensure a greater science return on the investment of the scientific community and the public. Future avenues of application could be to perform training of filters in a low-latency manner, such that a cleaned strain time series could be consistently available not long after




FIG. 7. DeepClean architecture for the O3 60-Hz dataset. The input has 21 witness channels, including 1 PEM 60-Hz channel and 20 ASC channels. Each channel has 8192 data samples (8 s of data at a sample rate of 1024 Hz).

the raw data are recorded. Running regressions in a constant online manner would also facilitate the use of cleaned data in the gravitational-wave search pipelines, which require the use of the entire run's data to properly estimate the statistical significance of events over the background. The prospects of rerunning searches on previous data would be especially promising if successful nonlinear regression routines are developed to capture sources such as scattering noise which was known to be a significant hindrance to the sensitivity of the aLIGO detectors during the first and third observation runs. In this case, it may be possible for marginal event candidates [1] to be promoted to fully confident detections.

ACKNOWLEDGMENTS

Michael Coughlin was supported by NSF Award No. PHY-0757058 and the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. Rich Ormiston was supported in part by the NSF Award No.

PHY-1806630. Tri Nguyen and Erik Katsavounidis work in this was supported by NSF Awards No. OAC-1934700 and No. OAC-1931469. The authors also acknowledge support from the LIGO Laboratory; LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and it operates under Cooperative Agreement No. PHY-1764464.

APPENDIX: NEURAL NETWORK ARCHITECTURE

In this section, we present the network architecture (Fig. 7) used in the O3 60-Hz data set. For different data sets, we keep the network hyperparameters (e.g., kernel size, filter size, etc.) the same. Therefore, the input and output dimension of each layer depends on the dimension of the input, which depends on the number of witness channels and length of the time series.

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo During the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
 - [2] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.* (The LIGO Scientific Collaboration), Advanced LIGO, *Class. Quant. Grav.* **32**, 074001 (2015).
 - [3] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Class. Quant. Grav.* **32**, 024001 (2015).
 - [4] H. Lück, C. Affeldt, J. Degallaix, A. Freise, H. Grote, M. Hewitson, S. Hild, J. Leong, M. Prijatelj, K. A. Strain, B. Willke, H. Wittel, and K. Danzmann, The upgrade of GEO 600, *J. Phys.: Conf. Ser.* **228**, 012012 (2010).
 - [5] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto, Interferometer design of the KAGRA gravitational wave detector, *Phys. Rev. D* **88**, 043007 (2013).
 - [6] C. S. Unnikrishnan, IndIGO and LIGO-INDIA: Scope and plans for gravitational wave research and precision metrology in India, *Int. J. Mod. Phys. D* **22**, 1341010 (2013).
 - [7] K. Riles, Gravitational waves: Sources, detectors, and searches, *Prog. Part. Nucl. Phys.* **68**, 1 (2013).
 - [8] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, M. Adamo, C. Adams, T. Adams, P. Addesso *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Characterization of transient noise in advanced LIGO relevant to gravitational wave signal GW150914, *Class. Quant. Grav.* **33**, 134001 (2016).
 - [9] 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.*, Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [10] R. Macas, Final live noise budget update, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=42950>.
 - [11] R. DeRosa, J. C. Driggers, D. Atkinson, H. Miao, V. Frolov, M. Landry, J. A. Giaime, and R. X. Adhikari, Global feed-forward vibration isolation in a km scale interferometer, *Class. Quant. Grav.* **29**, 215008 (2012).
 - [12] S. V. Vaseghi, *Wiener Filters* (John Wiley & Sons, New York, 2001), pp. 178–204.
 - [13] A. H. Sayed, *Fundamentals of Adaptive Filtering* (John Wiley & Sons, New York, 2003).
 - [14] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall, Improving the sensitivity of advanced LIGO using noise subtraction, *Class. Quant. Grav.* **36**, 055011 (2019).
 - [15] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos *et al.* (The LIGO Scientific Collaboration and the Virgo Collaboration), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, *Class. Quant. Grav.* **37**, 055002 (2020).
 - [16] M. Coughlin, N. Mukund, J. Harms, J. Driggers, R. Adhikari, and S. Mitra, Towards a first design of a Newtonian-noise cancellation system for advanced LIGO, *Class. Quant. Grav.* **33**, 244001 (2016).
 - [17] J. Harms, Terrestrial gravity fluctuations, *Living Rev. Relativ.* **18**, 3 (2015).
 - [18] M. W. Coughlin, J. Harms, J. Driggers, D. J. McManus, N. Mukund, M. P. Ross, B. J. J. Slagmolen, and K. Venkateswara, Implications of Dedicated Seismometer Measurements on Newtonian-Noise Cancellation for Advanced Ligo, *Phys. Rev. Lett.* **121**, 221104 (2018).
 - [19] S. Bose, B. Hall, N. Mazumder, S. Dhurandhar, A. Gupta, and A. Lundgren, Tackling excess noise from bilinear and nonlinear couplings in gravitational-wave interferometers, *J. Phys.: Conf. Ser.* **716**, 012007 (2016).

- [20] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson *et al.*, Gravity spy: Integrating advanced LIGO detector characterization, machine learning, and citizen science, *Class. Quant. Grav.* **34**, 064003 (2017).
- [21] R. Biswas, L. Blackburn, J. Cao, R. Essick, K. A. Hodge, E. Katsavounidis, K. Kim, Y.-M. Kim, E.-O. Le Bigot, C.-H. Lee *et al.*, Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data, *Phys. Rev. D* **88**, 062003 (2013).
- [22] N. Mukund, S. Abraham, S. Kandhasamy, S. Mitra, and N. S. Philip, Transient classification in LIGO data using difference boosting neural network, *Phys. Rev. D* **95**, 104059 (2017).
- [23] D. George and E. A. Huerta, Deep learning for real-time gravitational wave detection and parameter estimation: Results with advanced LIGO data, *Phys. Lett. B* **778**, 64 (2018).
- [24] V. Tiwari, M. Drago, V. Frolov, S. Klimenko, G. Mitselmakher, V. Necula, G. Prodi, V. Re, F. Salemi, G. Vedovato *et al.*, Regression of environmental noise in LIGO data, *Class. Quant. Grav.* **32**, 165014 (2015).
- [25] G. D. Meadors, K. Kawabe, and K. Riles, Increasing ligo sensitivity by feedforward subtraction of auxiliary length control noise, *Class. Quant. Grav.* **31**, 105014 (2014).
- [26] N. Mukund, J. Lough, C. Affeldt, F. Bergamin, A. Bisht, M. Brinkmann, V. Kringel, H. Lück, S. Nadji, and M. Weinert *et al.*, Bilinear noise subtraction at the GEO 600 observatory, *Phys. Rev. D* **101**, 102006 (2020).
- [27] G. Vajente, Y. Huang, M. Isi, J. C. Driggers, J. S. Kissel, M. J. Szczepanczyk, and S. Vitale, Machine-learning nonstationary noise out of gravitational-wave detectors, *Phys. Rev. D* **101**, 042003 (2020).
- [28] B. Farhang-Boroujeny, *Adaptive Filters: Theory and Applications*, 2nd ed. (John Wiley & Sons, New York, 2013).
- [29] V. Dumoulin and F. Visin, A guide to convolution arithmetic for deep learning, [arXiv:1603.07285](https://arxiv.org/abs/1603.07285) [stat.ML] (2016).
- [30] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in *Proceedings of the 32nd International Conference on Machine Learning*, edited by F. Bach and D. Blei, Proceedings of Machine Learning Research (PMLR, Lille, France, 2015), Vol. 37, pp. 448–456.
- [31] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *Third International Conference for Learning Representations, San Diego, 2015*, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [32] J. C. Driggers, S. Vitale, A. P. Lundgren, M. Evans, K. Kawabe, S. E. Dwyer, K. Izumi, R. M. S. Schofield, A. Effler, D. Sigg *et al.* (The LIGO Scientific Collaboration Instrument Science Authors), Improving astrophysical parameter estimation via offline noise subtraction for advanced LIGO, *Phys. Rev. D* **99**, 042001 (2019).
- [33] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [34] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [35] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, *Phys. Rev. D* **93**, 044006 (2016).
- [36] J. S. Speagle, Dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [37] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov *et al.*, Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).