

Boosting the Efficiency of Parametric Detection with Hierarchical Neural Networks

Jingkai Yan^{1,4}, Robert Colgan^{2,4}, John Wright^{1,4}, Zsuzsa Márka^{4,5}, Imre Bartos⁶, and Szabolcs Márka^{3,4}

¹*Department of Electrical Engineering, Columbia University in the City of New York, 500 W. 120th St., New York, NY 10027, USA*

²*Department of Computer Science, Columbia University in the City of New York, 500 W. 120th St., New York, NY 10027, USA*

³*Department of Physics, Columbia University in the City of New York, 538 W. 120th St., New York, NY 10027, USA*

⁴*Data Science Institute, Columbia University in the City of New York, 550 W. 120th St., New York, NY 10027, USA*

⁵*Columbia Astrophysics Laboratory, Columbia University in the City of New York, 538 W. 120th St., New York, NY 10027, USA*

⁶*Department of Physics, University of Florida, PO Box 118440, Gainesville, FL 32611-8440, USA*

Gravitational wave astronomy is a vibrant field that leverages both classic and modern data processing techniques for the understanding of the universe. Various approaches have been proposed for improving the efficiency of the detection scheme, with hierarchical matched filtering being an important strategy. Meanwhile, deep learning methods have recently demonstrated both consistency with matched filtering methods and remarkable statistical performance. In this work, we propose Hierarchical Detection Network (HDN), a novel approach to efficient detection that combines ideas from hierarchical matching and deep learning. The network is trained using a novel loss function, which encodes simultaneously the goals of statistical accuracy and efficiency. We discuss the source of complexity reduction of the proposed model, and describe a general recipe for initialization with each layer specializing in different regions. We demonstrate the performance of HDN with experiments using open LIGO data and synthetic injections, and observe with two-layer models a 79% efficiency gain compared with matched filtering at an equal error rate of 0.2%. Furthermore, we show how training a three-layer HDN initialized using two-layer model can further boost both accuracy and efficiency, highlighting the power of multiple simple layers in efficient detection.

I. INTRODUCTION

The study of gravitational wave (GW) signals [1–7] is a vibrant field that constantly expands our understanding of gravitational phenomena and our universe. In the detection schemes currently employed by gravitational-wave detectors [8–10], such as KAGRA [11], GEO600 [12], Virgo [13], and LIGO [14–15], the core algorithmic methods consists of excess energy-based burst searches [16–22] and matched filtering searches [23–42]. Matched filtering is a classic signal processing technique, which computes the correlation of the time-delayed input signal with a bank of templates. While the matched filtering method has outstanding performance and statistical rigor, it potentially has room for improvement in terms of computational complexity [43]. In fact, millions of templates are being used in the LIGO matched filtering pipeline, with the number still expanding with the scope of the search [5].

As we deepen our search for gravitational wave signals, the issue of computational efficiency (namely, the number of basic operations required by a computer) is becoming increasingly prominent. Detection methods that excel in both statistical performance and computational efficiency can significantly boost our capacities for exploring wider and higher-dimensional parameter spaces, and even other families of eccentric waveforms [44]. This in turn will help with uncovering more astrophysical events, potentially unveiling novel astrophysical phenomena, as well as reducing the carbon footprint associated with searching for these events.

In the literature, a promising approach to reducing the

complexity of matched filtering searches has been to apply a two-step hierarchical search, which seeks to rapidly reject most negative samples [45]. Later, [46] expands the hierarchy to involve temporal multi-scale approach. Some other meritorious extensions include using geometric template placing [47] and hierarchy based on chirp times [48]. The work of [49] applies two-step detection within the PyCBC framework, and compares the performance on simulated data. A recent work of [50] further combines the two-step method with dimensionality reduction in the template space using principle component analysis (PCA). All the above examples demonstrate improvements relative to basic matched filtering in various settings for GW detection.

Similar ideas have also been widely explored and applied in machine learning contexts. For example, [51, 52] consider hierarchical matching of image features in both spatial domain and feature domain for image classification. [53] considers image classification using hierarchical matching in the spatial domain. In natural language processing, hierarchical model have also been used in sentiment classification [54]. More specific applications of this idea include medical imaging [55], human detection and segmentation [56], and crime classification [57].

In the meantime, with the growing literature of applying deep learning and neural networks on GW detection, it is tempting to leverage deep learning’s power to reduce complexity. Indeed, various neural network architectures have been shown to perform tasks such as GW detection, parameter estimation, noise transients identification and data denoising [58–93], at performance levels comparable to that of matched filtering. Furthermore, it has been

shown that matched filtering is generally suboptimal for parametric signal detection [43, 94], and the performance can be improved by optimizing the templates using deep learning techniques [43]. This can be achieved by setting up a neural network that is formally equivalent to matched filtering, and then training on data. Inspired by the flexibility of deep learning models, it is conceptually appealing to explicitly incorporate computational efficiency into the neural network objectives, aim to achieve “the best of both worlds.”

In this work, we propose a novel neural network architecture, named Hierarchical Detection Network (HDN), which takes the form of a multi-layer matched filtering with trainable parameters. In order to achieve the dual goal of accuracy and efficiency, we constructed a *novel loss function* that explicitly incorporates computational complexity. We demonstrate the efficiency gains on data with open LIGO noise data and synthetic GW signal injections. As a quick glance at the performance gains, when tested on synthetically injected data at SNR = 9, compared with matched filtering, two-layer HDN can achieve false positive and false negative rates 0.2% with 79% lower complexity, and reduces error rates by 88% when at equal complexity equivalent to 100 templates, for instance. Experimental details are described in Section V.

Yet, the two-layer networks do not reveal the full power of the proposed model. We further show that by training a three-layer model with careful initialization, it is capable of achieving even better accuracy at lower complexity. We also provide some intuitive insights into the mechanism behind multi-layer hierarchical models and their construction.

The rest of the paper is organized as follows. Section II reviews the problem of parametric detection and some relevant models. Section III introduces Hierarchical Detection Networks, including the setup, complexity and training process. Section IV further discusses the complexity reduction from HDN. Section V presents experimental results of applying HDN on real LIGO data and synthetic injections. We discuss some further implications and future steps of this work in Section VI.

II. PRELIMINARIES

In this section, we describe the problem setup for a single gravitational-wave detector, and review some related detection algorithms, including matched filtering and a closely related model MNet [43].

A. Parametric Detection and Matched Filtering

Consider the problem of detecting gravitational waves in a single gravitational-wave detector data stream. Formally, assume we observe detector data $\mathbf{x} \in \mathbb{R}^n$, and need

to decide whether or not $\mathbf{x}$ contains any gravitational-wave signal in addition to noise. Throughout this paper, we model the inputs as having fixed dimension, where time-domain convolution can be applied for generic time-series inputs. The signals of interest are modeled as belonging to a parametric set

$$S_\Gamma = \{\mathbf{s}_\gamma \mid \gamma \in \Gamma\} \quad (1)$$

where the parameters γ can represent properties of the objects that generate the gravitational wave, such as masses, orbits and spins. We model noise as a random vector $\mathbf{z} \in \mathbb{R}^n$ with distribution ρ_0 and independent of the signal. The detection problem can thus be written as the following hypothesis test:

$$H_0 : \mathbf{x} = \mathbf{z}, \quad (2)$$

$$\text{or } H_1 : \mathbf{x} = \mathbf{s}_\gamma + \mathbf{z} \text{ for some } \gamma \in \Gamma, \quad (3)$$

For this testing problem, we need to design decision rules $\delta : \mathbb{R}^n \rightarrow \{0, 1\}$ that ideally achieve both good statistical performance and high computational efficiency.

One natural approach to this problem is matched filtering [95–97], a classical method for signal detection. A matched filter is the optimal linear filter for maximizing the signal-to-noise ratio (SNR) in stationary Gaussian white noise [95]. When detecting a single target signal $\mathbf{s}$, matched filtering takes the following form of an inner product test

$$\delta(\mathbf{x}) = 1 \text{ iff } \langle \mathbf{s}, \mathbf{x} \rangle \geq \tau, \quad (4)$$

where τ is a cutoff threshold. Its natural extension to detecting a parametric family of signals is to take the maximum over template samples $\mathbf{s}_{\gamma_1}, \dots, \mathbf{s}_{\gamma_k}$ in the signal space, namely

$$\delta(\mathbf{x}) = 1 \text{ iff } \max_{i=1, \dots, k} \langle \mathbf{s}_{\gamma_i}, \mathbf{x} \rangle \geq \tau. \quad (5)$$

As noted in the literature [43], matched filtering is suboptimal for the generic parametric detection problem in terms of statistical performance. To address such suboptimality, one alternative approach that has been described is the MNet architectures.

B. MNet Architectures

The MNet architectures, including MNet-Deep and MNet-Shallow, are two machine learning model for parametric detection based on neural networks [43]. Formally, an MNet is a neural network initialized to exactly replicate matched filtering, and then trained on data for improved performance. Such improvement is mainly due to its structural capability of handling non-convex decision boundaries and non-Gaussian noise distributions. The shallow and deep versions differ in how the replication of matched filtering is constructed. MNet-Shallow is structured identically to matched filtering, allowing for

a more direct comparison, whereas MNet-Deep uses a multi-layer pairwise-max structure for higher flexibility and statistical performance.

In this work, with computational complexity in mind, we will focus on the MNet-Shallow model, which can be expressed as:

$$f_{\text{MNet-Shallow}}(\mathbf{x}) = \max_i \langle \mathbf{s}_i, \mathbf{x} \rangle. \quad (6)$$

Here the weights $\mathbf{s}_i$ are initialized as templates and then trained over data, making it advantageous over classic matched filtering. If we compare its computational efficiency against matched filtering (measured in terms of the number of operations required to achieve a target error rate), the strict performance improvement with identical architecture suggests that one can expect a strict efficiency improvement as well.

However, the structural similarity between MNet-Shallow and matched filtering implies that such efficiency gains may typically be very limited. In order to achieve efficiency gains on higher orders of magnitude, we may need to reconsider the parametric detection problem, and innovate on the basic matched filtering rule. As we present in the next section, one solution is to arrange the templates in a multi-layer hierarchy, so that significant proportions of negative-labeled data are subject to early rejections.

III. HIERARCHICAL DETECTION NETWORKS

In this section, we present the Hierarchical Detection Network (HDN), which improves over matched filtering and MNet-Shallow to simultaneously maximize statistical performance and computational efficiency.

The main idea behind HDN is intuitive. If an input segment clearly contains no gravitational wave signals, we may not need to subject it to millions of templates to tell that. A small number of “gatekeeper” templates may be sufficient for confidently rejecting these “obviously wrong” instances. Once these inputs have been ruled out, we can apply a more refined test using possibly more templates, and reject a larger portion of the input space. This procedure can be repeated, until in the very last step, we employ our full template bank for a full diagnosis on the remaining instances which all previous tests failed to reject. Since the overwhelming majority part of the gravitational wave strain data contains noise only, most instances will likely be addressed by the initial simple layers of the model, saving the need for the full template bank. In addition, different layers of the HDN may be designed to specialize in different parts of the input space, such that the available parameter space of the potentially allowed waveforms are successively restricted as the hierarchical process progresses from later to layer, allowing for further efficiency gains.

A. Architecture of HDN

We first formally define a hierarchical detection network (HDN). Generally speaking, a HDN is a hierarchical template matching model trained as a neural network, as illustrated in Fig 1. Let L be the number of layers in the hierarchical structure, and let $\{\mathbf{s}_i\}_{i=1}^n$ be the entire set of n templates used by the model. For each layer $\ell = 1, \dots, L$, only the first n_ℓ of these templates are used in that layer, where $0 < n_1 < \dots < n_L = n$. Let the threshold associated with template i at layer ℓ be $t_{i,\ell}$, $i = 1, \dots, n_\ell$. Here we let layer ℓ reuse all templates from the previous layer(s), but assign independent threshold values to the reused templates at different layers, in order to reduce computation complexity.




FIG. 1. Illustration of a hierarchical detection network.

Following conventions of the machine learning literature on binary classification, we call an input $\mathbf{x}$ *positive* if it contains a GW signal, and *negative* if it only contains noise. For a given input $\mathbf{x}$, the model processes it using the following procedure:

Algorithm 1 The HDN algorithm

Parameters: $L, n, \{n_\ell\}, \{\mathbf{s}_i\}, \{t_{i,\ell}\}$

Input: $\mathbf{x}$

```

ℓ ← 1
while ℓ ≤ L do
  y_ℓ ← max_{i ≤ n_ℓ} ⟨x, s_i⟩ - t_{i,ℓ}
  if y_ℓ < 0 then
    return 0
  end if
  ℓ ← ℓ + 1
end while
return 1
  
```

More formally, for a given input $\mathbf{x}$, let

$$y_\ell = \max_{i \leq n_\ell} \langle \mathbf{x}, \mathbf{s}_i \rangle - t_{i,\ell} \quad (7)$$

be the matching output at layer ℓ , $\ell = 1, \dots, L$. Let

$$I_\ell = \begin{cases} \mathbb{1}[y_\ell > 0], & \text{if } I_{\ell-1} = 1 \\ 0, & \text{if } I_{\ell-1} = 0 \end{cases} \quad (8)$$

be the indicator of whether the input passes layer ℓ of the model on to the later layer(s), $\ell = 1, \dots, L-1$, and define $I_L \equiv 0$. With these notations, the overall output of the model can be written as

$$y(\mathbf{x}) = \sum_{\ell=1}^L y_\ell (1 - I_\ell) \prod_{k=1}^{\ell-1} I_k. \quad (9)$$

Note that both matched filtering and MNet-Shallow can be unified under the framework of HDN, viewed as a model with a single layer. In the meantime, some of the existing two-step MF methods [45, 47, 49] can also be interpreted under this framework. Furthermore, the HDN architecture is not restricted to the typical two-layer hierarchy of coarse and fine searches, but can utilize multiple layers which specialize in different parts of the signal space. The use of multiple layers and their setup is further discussed in Section IV.

B. Measure of Computational Complexity

With matched filtering and HDN unified under the same framework, we can provide a formal definition of computational complexity to facilitate our discussion. We are often most concerned about the execution efficiency of the model in deployment rather than in training, since it determines the real-time processing abilities. In the meantime, any computational cost of setting up the parameters of the model, including template selection for matched filtering and training for neural networks, is a one-time cost and can be conducted offline. Therefore it is natural to define complexity based on test time.

Also, since for the vast majority of time the input strain does not contain gravitational wave events, we can capture the computational complexity solely by its performance on negative data. This leads to the following definition of complexity:

Definition 1 (Complexity). *The (computational) complexity of a HDN model is defined as the expected number of template matching (inner product) operations conducted to evaluate a negative input.*

Formally, we can write the complexity as

$$Z = \mathbb{E}_{\mathbf{x} \sim F_-} [z(\mathbf{x})], \quad (10)$$

where

$$z(\mathbf{x}) = \sum_{\ell=1}^L n_\ell (1 - I_\ell) \prod_{k=1}^{\ell-1} I_k \quad (11)$$

is the number of matching operations required for evaluating an input $\mathbf{x}$.

To illustrate this measure of complexity, note that for matched filtering and MNet-Shallow models, the complexity simply equals the number of templates used in the model. For a two-layer HDN, assuming only a proportion p of negative data enters the second layer, the complexity for the model will be $n_1 + p \cdot n_2$. Intuitively, if the initial layer contains fewer templates while being able to reject a significant portion of negative inputs, these inputs will not need to undergo the entire model, hence reducing the complexity of the model. This straightforward idea forms the basis of HDN, upon which we further leverage the power of data through training for an additional boost in performance.

C. Training of HDN

So far, we have described the behavior of HDN at test/deployment time, and now we turn our attention to the training process. Conceptually, we want to set up a loss function as an appropriate combination of classification error and model complexity, so that minimizing the loss would achieve simultaneously accuracy and efficiency. However, a loss function directly based on the above expressions (9) and (11) is undesirable because of non-differentiability. Instead, we use soft surrogates for the indicators I_ℓ . Define

$$\hat{I}_\ell = \phi(y_\ell) \quad (12)$$

for $\ell = 1, \dots, L-1$ where $\phi(x) := \frac{1}{1+e^{-x}}$ is the sigmoid function, serving as a soft surrogate of the step function. Also let $\hat{I}_L \equiv 0$. Note that during training we can simply compute y_ℓ for all layers regardless of whether previous layers were passed, since this will only be a one-time off-line cost. Define the soft surrogates for $y(\mathbf{x})$ and $z(\mathbf{x})$ accordingly:

$$\hat{y}(\mathbf{x}) = \sum_{\ell=1}^L y_\ell (1 - \hat{I}_\ell) \prod_{k=1}^{\ell-1} \hat{I}_k, \quad (13)$$

$$\hat{z}(\mathbf{x}) = \sum_{\ell=1}^L n_\ell (1 - \hat{I}_\ell) \prod_{k=1}^{\ell-1} \hat{I}_k. \quad (14)$$

Assume the training dataset is $\{(\mathbf{x}_i, y_i^*)\}_{i=1}^N$, with N_+ positive entries and N_- negative entries. The loss function can be formulated as

$$\mathcal{L} = \frac{1}{N} \sum_i \ell_i^{\text{accu}} + \lambda \cdot \frac{1}{N_-} \sum_{i: y_i^*=0} \ell_i^{\text{plx}}, \quad (15)$$

where

$$\ell_i^{\text{accu}} = y_i^* \log p_i + (1 - y_i^*) \log (1 - p_i) \quad (16)$$

with $p_i = \frac{1}{1+e^{-\hat{y}_i}}$ is equivalent to the cross-entropy loss for binary classification, and

$$\ell_i^{\text{plx}} = \hat{z}_i \quad (17)$$

is the soft approximate for the complexity of evaluate the negative inputs. With the loss function (15) defined above, we can then train the model parameters $\{s_i\}$ and $\{t_{i,\ell}\}$ using first order optimization methods.

Experimental results of the HDN architecture will be shown in Section V

IV. COMPLEXITY REDUCTION FROM MULTIPLE LAYERS

Here we provide some theoretical insights into why the hierarchical model achieves reduced complexity at similar target performance levels, particularly with more layers.

Consider as an example a two-layer hierarchical model with n_1, n_2 templates respectively on the layers. Let α_ℓ and β_ℓ denote respectively the FPR and FNR of layer ℓ conditioned on data that reaches the corresponding layer. The overall FPR, FNR and the complexity z can then be represented as:

$$\alpha_{\text{all}} = \alpha_1 \alpha_2 \quad (18)$$

$$\beta_{\text{all}} = \beta_1 + (1 - \beta_1)\beta_2 \quad (19)$$

$$z_{\text{all}} = n_1 + \alpha_1(n_2 - n_1) \quad (20)$$

To understand why an improvement in complexity can be expected, we consider the following example setup of parametric detection as shown in FIG. 2. The probability density of the two labeled classes are ρ_0 and ρ_1 respectively. Imagine a baseline MF model with decision




FIG. 2. An example of the complexity advantage of hierarchical detection models.

boundary as shown by the green curve, at the cost of n templates, where n has to be relatively large to approximate the smoothly curved boundary. Then we can construct the following hierarchical model to achieve a significantly lower complexity with identical statistical accuracy. To do this, we construct a simple two-layer hierarchical model, with the first layer decision boundary as shown by the dotted blue line, and the second layer decision boundary coinciding with that of the MF model. Notice that the first layer features a very low complexity

n_1 (with $n_1 = 2$ in this example), and in the meantime has a fairly high true negative rate $1 - \alpha$. Since the second layer reproduces the MF decision boundary, the overall decision rule of the hierarchical model is identical to that of the MF model, and hence they share exactly the same ROC (receiver operating characteristic) curves. However, the complexity of the HDN model is $n_1 + \alpha_1(n_2 - n_1)$, which is significantly smaller than n_2 provided n_1 is small compared with n_2 and α_1 is not too close to 1.

This example provides inspirations for a general recipe for designing hierarchical models with reduced complexity. For any decision rule given by a MF model, we can construct a sequence of preceding layers whose negative decision regions all lie inside the negative decision region of the MF, and finally let the very last layer be equivalent to the original MF. The resulting hierarchical model will again have exactly the same overall decision rule and hence ROC curve, but with significantly reduced complexity. An illustrating example is shown in FIG. 3




FIG. 3. A hierarchical model with more simple layers that lie inside the overall negative decision region.

More generally, such constructions of hierarchical models can serve as good initializations for a HDN. One practical initialization scheme for an L -layer HDNs is the following: first train a separate $L - 1$ -layer model that only consists of the latter $L - 1$ layers of the desired model. Then we initialize the latter $L - 1$ layers of the original model with the trained network, and initialize the first layer with small $t_{i,1}$ values such that almost all inputs pass. This gives an initialization of the L -layer model which at initialization essentially replicates the $L - 1$ -layer model. From there, we train the initialized L -layer model on data, which will leverage the higher architectural capacity for further improved performance and complexity. An experiment that illustrates this approach is shown in Section V

V. SIMULATION AND EXPERIMENTS

A. Data Generation

In the experiments, we use open L1 strain data from LIGO Livingston's O2 run between August 1 and August 25, 2017 with ANALYSIS_READY flag [98]. The total duration of the frame files is 389.12 hours. We downsample the strain data from the original 4096Hz to 2048Hz for processing efficiency. The downsampled data is then divided into segments of 2 seconds, with each segment overlapping with 50% of its preceding segment.

To evaluate the accuracy of detection models, we need both positive and negative datasets. For the negative datasets, the strain data itself is used. For the positive datasets, due to the very limited number of confirmed detections of GW events, we generate positive data by injecting synthetic waveforms into the noise strains, at a preset SNR value.

The entire L1 strain dataset is first divided into two sets to be used in training and test respectively, such that any segment in the training set does not overlap with any segment in the test set. For training and test respectively, a positive and a negative dataset are generated. For the positive datasets, synthetic waveforms are generated with masses m_1, m_2 uniformly drawn from $[20, 50]$ and 3-dimensional spins drawn from an isotropic distribution and with spin dimensionless magnitudes drawn from a uniform distribution within $[0, 1]$. The injected waveforms are aligned such that the peak is located at 0.95 second, and the injection amplitude is chosen such that the signal-to-noise ratio (SNR) after preprocessing is constant at 9. The preprocessing is applied to all data (after injection if applicable) by using an FIR bandpass filter with cutoff frequency 30Hz and 400Hz, whitening with power spectral density estimated from the L1 strain data, and truncating to only keep the center 1 second.

B. Two-Layer Networks

In this experiment, we limit our HDN models to two layers and $n_2 = 10n_1$. At initialization, the templates $\mathbf{s}_i$ are chosen as random gravitational waveforms from the same parameter space, and the thresholds $t_{i,\ell}$ are set to the same within each of the two layers. The parameter λ is the loss function is fixed at $\lambda = 10^{-4}$. For the optimization procedure of the network, we use the Adam optimizer and a constant learning rate of 10^{-4} .

FIG. 4 shows the comparison of the complexity-performance trade-offs of MF and HDN models, where the HDN models are two-layer architectures structured as described above. The horizontal axis plots the logarithm of the complexity measure defined in this paper, and the vertical axis plots the logarithm of error rates at the point on the ROC curve where $\text{FPR} = \text{FNR}$. This choice of measure eliminates the arbitrariness of choosing FNR at a fixed FPR level. For each architecture, 10




FIG. 4. Complexity-performance trade-off of matched filtering and the hierarchical neural network.




FIG. 5. Proportion of error rate reduced by using HDN over matched filtering.

independent runs are conducted, and the one with lowest accuracy measure is shown. The red curve for HDN is cut off early due to memory limitations of training the model. We see that HDN consistently achieves a lower complexity than MF at equal accuracy.

C. Three-Layer Networks

We further demonstrate the power of the proposed model with a deeper three-layer network. Conceptually, since adding more layers strictly improves model expressability, it should never hurt performance provided that the parameters are initialized or trained appropriately.

In this experiment, we construct a 3-layer HDN with layer sizes $(n_1, n_2, n_3) = (30, 100, 1000)$ in the following way. First, a shallower 2-layer model with layer sizes $(100, 1000)$ is trained, and we use these trained parameters to initialize the latter two layers of the 3-layer model. We then initialize the first layer of the 3-layer model, setting the per-template thresholds $t_{i,1}$ as the same value for all i , such that all training data passes this first layer at

initialization. This scheme ensures that at initialization, the 3-layer model essentially replicates the performance of the trained 2-layer model, giving it a head start before entering the training phase. The training is done in the same way as described before. FIG. 6 illustrates the architecture of this 3-layer model, along with the output densities of test data that reaches that layer, separated by the true class labels. Namely, the densities of layer 1 involves all input data, and the densities of layers 2 and 3 involve only the data that pass the previous layers. We see that most of the negative data are successfully intercepted by the initial layers, with very few of them reaches the final layer, which corroborates our intuition.




FIG. 6. Illustration of the 3-layer architecture, and the output densities on the test data from each layer. Only data entries that reach a given layer is shown. We see that each layer successfully rejects the vast majority of incoming negative data, and barely any negative data reaches the last layer.

Here when evaluating the ROC curve, we adopt a slightly different approach that is more consistent with deep hierarchical models. Notice from equation (8) that the model uses a built-in threshold 0 to control the passing of each layer. When generating the ROC curve using a varying threshold, such a threshold should be applied at all layers instead of only the last layer. Therefore, at test time only, we replace the threshold 0 in equation (8) with a variable threshold $t \in \mathbb{R}$ which is constant for all layers, and compute the test outputs using (9) as before for each t value. Varying this threshold t produces the ROC curve. Also note that t determines which test entries would pass the layers, hence it also affects the model complexity evaluated on the negative test dataset. In actual deployment, the threshold t should be fixed at some

level that gives the desired trade-off between FPR and FNR, so this is only for demonstration purpose.




FIG. 7. Comparison of ROC curves between three models. The numbers in parentheses show the complexity of the model.

FIG. 7 shows the comparison of ROC curves between a matched filtering model with 100 templates, a 2-layer HDN model with complexity 100.5 (used for initializing the 3-layer model), and a 3-layer HDN model with complexity 37.1 at the point of equal FPR and FNR. While the complexity of the 3-layer model depends on the specific point chosen on the ROC curve, it does not exceed 65 for the entire segment of ROC curve shown in the figure, and is thus always lower than the 2-layer model. We see that the deeper 3-layer model excels at both accuracy and efficiency compared with the 2-layer model, and significantly more so if compared with the matched filtering model. This further showcases the power of depth in hierarchical models, and corroborates our discussion in Section IV.

VI. DISCUSSION

In this paper, we showed that by leveraging ideas from classical matched filtering and modern machine learning, we are able to design systems for GW detection that simultaneously optimize statistical accuracy and efficiency. This general conceptual idea of trainable hierarchical matched filtering can be applied upon a wide range of existing proposals for efficient detection pipelines.

While the proposed HDN model conducts hierarchical rejection on the data, an alternative can be proposed to conduct hierarchical acceptance, namely to progressively label parts of the data as positive rather than negative. This has the advantage of aligning better with

the matched filtering routine, since it suffices to use one matching template to confirm a signal. In the specific problem of GW signal detection, due to the class imbalance from the scarcity of actual GW events, the majority of computational complexity hinges on the classification of negative data, and therefore a hierarchical rejection model will have much more significant efficiency gains. In more general signal detection problems, hierarchical acceptance constructions can also be deployed in similar fashions as HDN.

The proposed HDN can potentially have wider applications within the field of gravitational wave science. For example, in the task of glitch detection and identification [80, 99, 103], one can combine existing constructions of machine learning based models with hierarchical models, to improve on both efficiency and accuracy.

One aspect to be further explored is how to select the number of layers in the hierarchy. While having more layers can boost model expressability and further leverage the efficiency gains, excessive hierarchy may offer diminishing returns, and also make training increasingly difficult. In the paper we demonstrated how a 3-layer model excels over a 2-layer one, and there may be a “sweet spot” number of layers for a given signal detection setups. Another promising direction would be to incorporate prior knowledge about the signal domain such as low-dimensionality and representative features into the detection model, which may be able to further outperform these current models agnostic of the signal space properties.

ACKNOWLEDGMENTS

This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

We acknowledge computing resources from Columbia University’s Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant 1G20RR030893-01, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010.

The authors are grateful for the LIGO Scientific Collaboration for the careful review of the paper. This paper is assigned a LIGO DCC number of LIGO-LIGO-P2200203. The authors acknowledge the LIGO Labora-

tory and Scientific Collaboration for the detectors, data, and the game changing computing resources (National Science Foundation Grants PHY-0757058 and PHY-0823459). The authors would like to thank colleagues of the LIGO Scientific Collaboration and the Virgo Collaboration and Columbia University for their help and useful comments, in particular the CBC group, Andrew Williamson, Stefan Countryman, William Tse, Nicolas Beltran, Asif Malik, Sireesh Gururaja, and Thomas Dent which we hereby gratefully acknowledge. SM thanks David Spergel, Rainer Weiss, Rana Adhikari, and Kipp Cannon for the motivating general discussions related to the role of machine learning and data analysis.

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center [104, 105] (<https://www.gwopenscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. This material is in part based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation, specifically some of the testing used LIGO detector noise publicly released by the LIGO Scientific Collaboration. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain.

The authors thank the University of Florida and Columbia University in the City of New York for their generous support. The authors are grateful for the generous support of the National Science Foundation under grant CCF-1740391. The authors thank Sharon Sputz of Columbia University for her effort in facilitating this collaboration.

I.B. acknowledges the support of the National Science Foundation under grant PHY-1911796 and the Alfred P. Sloan Foundation.

-
- [1] Albert Einstein. Näherungsweise Integration der Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, pages 688–696, January 1916.
 - [2] Albert Einstein. Über Gravitationswellen. *Sitzungs-*

berichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), pages 154–167, January 1918.

- [3] LIGO and Virgo Collaborations. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.
- [4] B. P. Abbott and et al. Gwtc-1: A gravitational-wave

- transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Physical Review X*, 9(3), Sep 2019.
- [5] Virgo LSC and KAGRA. Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *arXiv e-prints*, page arXiv:2010.14527, October 2020.
- [6] B. P. Abbott and et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.
- [7] B. P. Abbott and et al. Multi-messenger observations of a binary neutron star merger. *Astrophysical Journal Letters*, 848(2):L12, October 2017.
- [8] C Affeldt, K Danzmann, K L Dooley, H Grote, M Hewitson, S Hild, J Hough, J Leong, H Lück, M Prijatelj, S Rowan, A Rüdiger, R Schilling, R Schnabel, E Schreiber, B Sorazu, K A Strain, H Vahlbruch, B Willke, W Winkler, and H Wittel. Advanced techniques in GEO 600. *Classical and Quantum Gravity*, 31(22):224002, nov 2014.
- [9] F. Acernese, M. Agathos, L. Aiello, A. Allocca, A. Amato, S. Ansoldi, S. Antier, M. Arène, N. Arnaud, S. Ascenzi, P. Astone, F. Aubin, S. Babak, P. Bacon, F. Badaracco, M. K. M. Bader, J. Baird, F. Baldacchini, G. Ballardín, G. Baltus, C. Barbieri, P. Barneo, F. Barone, M. Barsuglia, D. Barta, A. Basti, M. Bawaj, M. Bazzan, M. Bejger, I. Belahcene, S. Bernuzzi, D. Bersanetti, A. Bertolini, M. Bisch, M. Bitossi, M. A. Bizouard, F. Bobba, M. Boer, G. Bogaert, F. Bondu, R. Bonnand, B. A. Boom, V. Boschi, Y. Bouffanaïs, A. Bozzi, C. Bradaschia, M. Branchesi, M. Breschi, T. Briant, F. Brighenti, A. Brillet, J. Brooks, G. Bruno, T. Bulik, H. J. Bulten, D. Buskulic, G. Cagnoli, E. Calloni, M. Canepa, G. Carapella, F. Carbognani, G. Carullo, J. Casanueva Diaz, C. Casentini, J. Castañeda, S. Caudill, F. Cavalier, R. Cavalieri, G. Cella, P. Cerdá-Durán, E. Cesarini, O. Chaibi, E. Chassande-Mottin, F. Chiadini, R. Chierici, A. Chincarini, A. Chiummo, N. Christensen, S. Chua, G. Ciani, P. Ciecielag, M. Cieřlar, R. Ciofi, F. Cipriano, A. Cirone, S. Clesse, F. Cleva, E. Coccia, P.-F. Cohadon, D. Cohen, M. Colpi, L. Conti, I. Cordero-Carrión, S. Corezzi, D. Corre, S. Cortese, J.-P. Coulon, M. Croquette, J.-R. Cudell, E. Cuoco, M. Curylo, B. D’Angelo, S. D’Antonio, V. Dattilo, M. Davier, J. Degallaix, M. De Laurentis, S. Deléglise, W. Del Pozzo, R. De Pietri, R. De Rosa, C. De Rossi, T. Dietrich, L. Di Fiore, C. Di Giorgio, F. Di Giovanni, M. Di Giovanni, T. Di Girolamo, A. Di Lieto, S. Di Pace, I. Di Palma, F. Di Renzo, M. Drago, J.-G. Ducoin, O. Durante, D. D’Urso, M. Eisenmann, L. Errico, D. Estevez, V. Fafone, S. Farinon, F. Feng, E. Fenyvesi, I. Ferrante, F. Fidecaro, I. Fiori, D. Fiorucci, R. Fittipaldi, V. Fiumara, R. Flaminio, J. A. Font, J.-D. Fournier, S. Frasca, F. Frasconi, V. Frey, G. Fronzè, F. Garufi, G. Gemme, E. Genin, A. Gennai, Archisman Ghosh, B. Giacomazzo, M. Gosselin, R. Gouaty, A. Grado, M. Granata, G. Greco, G. Grignani, A. Grimaldi, S. J. Grimm, P. Gruning, G. M. Guidi, G. Guixé, Y. Guo, P. Gupta, O. Halim, T. Harder, J. Harms, A. Heidmann, H. Heitmann, P. Hello, G. Hemming, E. Hennes, T. Hinderer, D. Hoffman, D. Huet, V. Hui, B. Idzkowski, A. Iess, G. Intini, J.-M. Isac, T. Jacqmin, P. Jaranowski, R. J. G. Jonker, S. Katsanevas, F. Kéfélian, I. Khan, N. Khetan, G. Koekoek, S. Koley, A. Królak, A. Kutynia, D. Laghi, A. Lamberts, I. La Rosa, A. Lartaux-Vollard, C. Lazzaro, P. Leaci, N. Leroy, N. Letendre, F. Linde, M. Llorens-Monteagudo, A. Longo, M. Lorenzini, V. Loriette, G. Losurdo, D. Lumaca, A. Macquet, E. Majorana, I. Maksimovic, N. Man, V. Mangano, M. Mantovani, M. Mapelli, F. Marchesoni, F. Marion, A. Marquina, S. Marsat, F. Martelli, V. Martinez, A. Masserot, S. Mastrogiovanni, E. Mejuto Villa, L. Mereni, M. Merzougui, R. Metzdrorff, A. Miani, C. Michel, L. Milano, A. Miller, E. Milotti, O. Minazzoli, Y. Minenkov, M. Montani, F. Morawski, B. Mours, F. Muciaccia, A. Nagar, I. Nardecchia, L. Naticchioni, J. Neilson, G. Nelemans, C. Nguyen, D. Nichols, S. Nissanke, F. Nocera, G. Oganessian, C. Olivetto, G. Pagano, G. Pagliaroli, C. Palomba, P. T. H. Pang, F. Pannarale, F. Paoletti, A. Paoli, D. Pascucci, A. Pasqualetti, R. Passaquieti, D. Passuello, B. Patricelli, A. Perego, M. Pegoraro, C. Pérois, A. Perreca, S. Perriès, K. S. Phukon, O. J. Piccinni, M. Pichot, M. Piendibene, F. Piergiovanni, V. Pierro, G. Pillant, L. Pinard, I. M. Pinto, K. Piotrkowski, W. Plastino, R. Poggiani, P. Popolizio, E. K. Porter, M. Prevedelli, M. Principe, G. A. Prodi, M. Punturo, P. Puppo, G. Raaijmakers, N. Radulesco, P. Rapagnani, M. Razzano, T. Regimbau, L. Rei, P. Rettengo, F. Ricci, G. Riemenschneider, F. Robinet, A. Rocchi, L. Rolland, M. Romanelli, R. Romano, D. Rosińska, P. Ruggi, O. S. Salafia, L. Salconi, A. Samajdar, N. Sanchis-Gual, E. Santos, B. Sassolas, O. Sauter, S. Sayah, D. Sentenac, V. Sequino, A. Sharma, M. Sieniawska, N. Singh, A. Singhal, V. Sipala, V. Sordini, F. Sorrentino, M. Spera, C. Stachie, D. A. Steer, G. Stratta, A. Sur, B. L. Swinkels, M. Tacca, A. J. Tanasijczuk, E. N. Tapia San Martin, S. Tiwari, M. Tonelli, A. Torres-Forné, I. Tosta e Melo, F. Travasso, M. C. Tringali, A. Trovato, K. W. Tsang, M. Turconi, M. Valentini, N. van Bakel, M. van Beuzekom, J. F. J. van den Brand, C. Van Den Broeck, L. van der Schaaf, M. Vardaro, M. Vasúth, G. Vedovato, D. Verkindt, F. Vetrano, A. Viceré, J.-Y. Vinet, H. Vocca, R. Walet, M. Was, A. Zadrożny, T. Zelenova, J.-P. Zendri, Henning Vahlbruch, Moritz Mehmet, Harald Lück, and Karsten Danzmann. Increasing the astrophysical reach of the advanced virgo detector via the application of squeezed vacuum states of light. *Phys. Rev. Lett.*, 123:231108, Dec 2019.
- [10] M. Tse, Haocun Yu, N. Kijbunchoo, A. Fernandez-Galiana, P. Dupej, L. Barsotti, C. D. Blair, D. D. Brown, S. E. Dwyer, A. Effler, M. Evans, P. Fritschel, V. V. Frolov, A. C. Green, G. L. Mansell, F. Matichard, N. Mavalvala, D. E. McClelland, L. McCuller, T. McRae, J. Miller, A. Mullavey, E. Oelker, I. Y. Phinney, D. Sigg, B. J. J. Slagmolen, T. Vo, R. L. Ward, C. Whittle, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, J. S. Areeda, Y. Asali, S. M. Aston, C. Austin, A. M. Baer, M. Ball, S. W. Ballmer, S. Banagiri, D. Barker, J. Bartlett, B. K. Berger, J. Betzwieser, D. Bhattacharjee, G. Billingsley, S. Biscans, R. M. Blair, N. Bode, P. Booker, R. Bork, A. Bramley, A. F. Brooks, A. Buikema, C. Cahillane, K. C. Cannon, X. Chen, A. A. Ciobanu, F. Clara, S. J. Cooper, K. R. Corley, S. T. Countryman, P. B. Covas, D. C. Coyne, L. E. H. Datrier, D. Davis, C. Di Fronzo,

- J. C. Driggers, T. Etzel, T. M. Evans, J. Feicht, P. Fulda, M. Fyffe, J. A. Giaime, K. D. Giardina, P. Godwin, E. Goetz, S. Gras, C. Gray, R. Gray, Anchal Gupta, E. K. Gustafson, R. Gustafson, J. Hanks, J. Hanson, T. Hardwick, R. K. Hasskew, M. C. Heintze, A. F. Helmling-Cornell, N. A. Holland, J. D. Jones, S. Kandhasamy, S. Karki, M. Kasprzack, K. Kawabe, P. J. King, J. S. Kissel, Rahul Kumar, M. Landry, B. B. Lane, B. Lantz, M. Laxen, Y. K. Lecoecueche, J. Leviton, J. Liu, M. Lormand, A. P. Lundgren, R. Macas, M. MacInnis, D. M. Macleod, S. Márka, Z. Márka, D. V. Martynov, K. Mason, T. J. Massinger, R. McCarthy, S. McCormick, J. McIver, G. Mendell, K. Merfeld, E. L. Merilh, F. Meylahn, T. Mistry, R. Mittleman, G. Moreno, C. M. Mow-Lowry, S. Mozzon, T. J. N. Nelson, P. Nguyen, L. K. Nuttall, J. Oberling, R. J. Oram, B. O'Reilly, C. Osthelder, D. J. Ottaway, H. Overmier, J. R. Palamos, W. Parker, E. Payne, A. Pele, C. J. Perez, M. Pirello, H. Radkins, K. E. Ramirez, J. W. Richardson, K. Riles, N. A. Robertson, J. G. Rollins, C. L. Romel, J. H. Romie, M. P. Ross, K. Ryan, T. Sadecki, E. J. Sanchez, L. E. Sanchez, T. R. Saravanan, R. L. Savage, D. Schaetzl, R. Schnabel, R. M. S. Schofield, E. Schwartz, D. Sellers, T. J. Shaffer, J. R. Smith, S. Soni, B. Sorazu, A. P. Spencer, K. A. Strain, L. Sun, M. J. Szczepańczyk, M. Thomas, P. Thomas, K. A. Thorne, K. Toland, C. I. Torrie, G. Traylor, A. L. Urban, G. Vajente, G. Valdes, D. C. Vander-Hyde, P. J. Veitch, K. Venkateswara, G. Venugopalan, A. D. Viets, C. Vorvick, M. Wade, J. Warner, B. Weaver, R. Weiss, B. Willke, C. C. Wipf, L. Xiao, H. Yamamoto, M. J. Yap, Hang Yu, L. Zhang, M. E. Zucker, and J. Zweizig. Quantum-enhanced advanced ligo detectors in the era of gravitational-wave astronomy. *Phys. Rev. Lett.*, 123:231107, Dec 2019.
- [11] T. Akutsu et al. Overview of KAGRA: Detector design and construction history. *Progress of Theoretical and Experimental Physics*, 08 2020. ptaa125.
- [12] K L Dooley, J R Leong, T Adams, C Affeldt, A Bisht, C Bogan, J Degallaix, C Gräf, S Hild, J Hough, A Khalaidovski, N Lastzka, J Lough, H Lück, D Macleod, L Nuttall, M Prijatelj, R Schnabel, E Schreiber, J Slutsky, B Sorazu, K A Strain, H Vahlbruch, M Was, B Willke, H Wittel, K Danzmann, and H Grote. GEO 600 and the GEO-HF upgrade program: successes and challenges. *Classical and Quantum Gravity*, 33(7):075009, mar 2016.
- [13] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32:024001, Jan 2015.
- [14] Alex Abramovici, William E. Althouse, Ronald W. P. Drever, Yekta Gursel, Seiji Kawamura, Frederick J. Raab, David Shoemaker, Lisa Sievers, Robert E. Spero, Kip S. Thorne, Rochus E. Vogt, Rainer Weiss, Stanley E. Whitcomb, and Michael E. Zucker. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science*, 256(5055):325–333, April 1992.
- [15] B. P. Abbott et al. Advanced LIGO. *Classical and Quantum Gravity*, 32:074001, Apr 2015.
- [16] Warren G. Anderson, Patrick R. Brady, Jolien D. E. Creighton, and Éanna É. Flanagan. A Power Filter for the Detection of Burst Sources of Gravitational Radiation in Interferometric Detectors. *International Journal of Modern Physics D*, 9(3):303–307, January 2000.
- [17] Laura Cadonati and Szabolcs Márka. CorrPower: a cross-correlation-based algorithm for triggered and untriggered gravitational-wave burst searches. *Classical and Quantum Gravity*, 22(18):S1159–S1167, September 2005.
- [18] Warren G. Anderson, Patrick R. Brady, Jolien D. Creighton, and Éanna É. Flanagan. Excess power statistic for detection of burst sources of gravitational radiation. *Phys. Rev. D*, 63(4):042003, February 2001.
- [19] Andrea Viceré. Optimal detection of burst events in gravitational wave interferometric observatories. *Phys. Rev. D*, 66(6):062002, September 2002.
- [20] S. Klimenko and G. Mitselmakher. A wavelet method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 21(20):S1819–S1830, October 2004.
- [21] M. Rakhmanov and S. Klimenko. A cross-correlation method for burst searches with networks of misaligned gravitational-wave detectors. *Classical and Quantum Gravity*, 22(18):S1311–S1320, September 2005.
- [22] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher. A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 25(11):114029, June 2008.
- [23] S. W. Hawking and W. Israel. *Three Hundred Years of Gravitation*. 1989.
- [24] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, Jun 1999.
- [25] Curt Cutler, Theodoros A. Apostolatos, Lars Bildsten, Lee Smauel Finn, Eanna E. Flanagan, Daniel Kennefick, Dragoljub M. Markovic, Amos Ori, Eric Poisson, Gerald Jay Sussman, and Kip S. Thorne. The last three minutes: Issues in gravitational-wave measurements of coalescing compact binaries. *Phys. Rev. Lett.*, 70(20):2984–2987, May 1993.
- [26] Curt Cutler and Éanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Phys. Rev. D*, 49(6):2658–2697, March 1994.
- [27] Éanna É. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown. *Phys. Rev. D*, 57(8):4535–4565, April 1998.
- [28] Éanna É. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences. II. The waves’ information and its extraction, with and without templates. *Phys. Rev. D*, 57(8):4566–4587, April 1998.
- [29] B. et al Abbott. Analysis of LIGO data for gravitational waves from binary neutron stars. *Phys. Rev. D*, 69(12):122001, June 2004.
- [30] Pycbc software releases. <https://github.com/gwastro/pycbc/releases>
- [31] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FIND-CHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85(12):122006, June 2012.
- [32] C. M. Biwer, Collin D. Capano, Soumi De, Miriam Cabero, Duncan A. Brown, Alexander H. Nitz, and

- V. Raymond. PyCBC Inference: A Python-based Parameter Estimation Toolkit for Compact Binary Coalescence Signal. *PASP*, 131(996):024503, February 2019.
- [33] Bruce Allen. χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71(6):062001, March 2005.
- [34] Tito Dal Canton, Alexander H. Nitz, Andrew P. Lundgren, Alex B. Nielsen, Duncan A. Brown, Thomas Dent, Ian W. Harry, Badri Krishnan, Andrew J. Miller, Karl Wette, Karsten Wiesner, and Joshua L. Willis. Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors. *Phys. Rev. D*, 90(8):082004, October 2014.
- [35] Samantha A. Usman, Alexander H. Nitz, Ian W. Harry, Christopher M. Biwer, Duncan A. Brown, Miriam Cabero, Collin D. Capano, Tito Dal Canton, Thomas Dent, Stephen Fairhurst, Marcel S. Kehl, Drew Keppel, Badri Krishnan, Amber Lenon, Andrew Lundgren, Alex B. Nielsen, Larne P. Pekowsky, Harald P. Pfeiffer, Peter R. Saulson, Matthew West, and Joshua L. Willis. The PyCBC search for gravitational waves from compact binary coalescence. *Classical and Quantum Gravity*, 33(21):215004, November 2016.
- [36] Alexander H. Nitz, Tito Dal Canton, Derek Davis, and Steven Reyes. Rapid detection of gravitational waves from compact binary mergers with PyCBC Live. *Phys. Rev. D*, 98(2):024050, July 2018.
- [37] Cody Messick, Kent Blackburn, Patrick Brady, Patrick Brockill, Kipp Cannon, Romain Cariou, Sarah Caudill, Sydney J. Chamberlin, Jolien D. E. Creighton, Ryan Everett, Chad Hanna, Drew Keppel, Ryan N. Lang, Tjonnie G. F. Li, Duncan Meacher, Alex Nielsen, Chris Pankow, Stephen Privitera, Hong Qi, Surabhi Sachdev, Laleh Sadeghian, Leo Singer, E. Gareth Thomas, Leslie Wade, Madeline Wade, Alan Weinstein, and Karsten Wiesner. Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data. *Phys. Rev. D*, 95(4):042001, February 2017.
- [38] Surabhi Sachdev, Sarah Caudill, Heather Fong, Rico K. L. Lo, Cody Messick, Debnandini Mukherjee, Ryan Magee, Leo Tsukada, Kent Blackburn, Patrick Brady, Patrick Brockill, Kipp Cannon, Sydney J. Chamberlin, Deep Chatterjee, Jolien D. E. Creighton, Patrick Godwin, Anuradha Gupta, Chad Hanna, Shasvath Kapadia, Ryan N. Lang, Tjonnie G. F. Li, Duncan Meacher, Alexander Pace, Stephen Privitera, Laleh Sadeghian, Leslie Wade, Madeline Wade, Alan Weinstein, and Sophia Liting Xiao. The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs. *arXiv e-prints*, page arXiv:1901.08580, January 2019.
- [39] Chad Hanna, Sarah Caudill, Cody Messick, Amit Reza, Surabhi Sachdev, Leo Tsukada, Kipp Cannon, Kent Blackburn, Jolien D. E. Creighton, Heather Fong, Patrick Godwin, Shasvath Kapadia, Tjonnie G. F. Li, Ryan Magee, Duncan Meacher, Debnandini Mukherjee, Alex Pace, Stephen Privitera, Rico K. L. Lo, and Leslie Wade. Fast evaluation of multidetector consistency for real-time gravitational wave searches. *Phys. Rev. D*, 101(2):022003, January 2020.
- [40] Qi Chu. *Low-latency detection and localization of gravitational waves from compact binary coalescences*. PhD thesis, The University of Western Australia, 2017.
- [41] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Classical and Quantum Gravity*, 33(17):175012, September 2016.
- [42] Alexander H Nitz, Thomas Dent, Tito Dal Canton, Stephen Fairhurst, and Duncan A Brown. Detecting binary compact-object mergers with gravitational waves: Understanding and improving the sensitivity of the pycbc search. *The Astrophysical Journal*, 849(2):118, 2017.
- [43] Jingkai Yan, Mariam Avagyan, Robert E Colgan, Doğa Veske, Imre Bartos, John Wright, Zsuzsa Márka, and Szabolcs Márka. Generalized approach to matched filtering using neural networks. *arXiv preprint arXiv:2104.03961*, 2021.
- [44] V Gayathri, J Healy, J Lange, B O'Brien, M Szczepańczyk, Imre Bartos, M Campanelli, S Klimenko, CO Lousto, and R O'Shaughnessy. Eccentricity estimate for black hole mergers with numerical relativity simulations. *Nature Astronomy*, pages 1–6, 2022.
- [45] SD Mohanty and SV Dhurandhar. Hierarchical search strategy for the detection of gravitational waves from coalescing binaries. *Physical Review D*, 54(12):7108, 1996.
- [46] Anand S Sengupta, Sanjeev Dhurandhar, and Albert Lazzarini. Faster implementation of the hierarchical search algorithm for detection of gravitational waves from inspiraling compact binaries. *Physical Review D*, 67(8):082004, 2003.
- [47] Benjamin J Owen and Bangalore Suryanarayana Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Physical Review D*, 60(2):022002, 1999.
- [48] Anand S Sengupta, Sanjeev V Dhurandhar, Albert Lazzarini, and Tom Prince. Extended hierarchical search (ehs) algorithm for detection of gravitational waves from inspiralling compact binaries. *Classical and Quantum Gravity*, 19(7):1507, 2002.
- [49] Bhooshan Gadre, Sanjit Mitra, and Sanjeev Dhurandhar. Hierarchical search strategy for the efficient detection of gravitational waves from nonprecessing coalescing compact binaries with aligned-spins. *Physical Review D*, 99(12):124035, 2019.
- [50] Rahul Dhurkunde, Henning Fehrmann, and Alexander H Nitz. A hierarchical approach to matched filtering using a reduced basis. *arXiv preprint arXiv:2110.13115*, 2021.
- [51] Kristen Grauman and Trevor Darrell. The pyramid match kernel: Discriminative classification with sets of image features. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1458–1465. IEEE, 2005.
- [52] Qiang Chen, Zheng Song, Yang Hua, Zhongyang Huang, and Shuicheng Yan. Hierarchical matching with side information for image classification. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3426–3433. IEEE, 2012.
- [53] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, volume 2, pages 2169–2178.

- IEEE, 2006.
- [54] Chenlin Shen, Changlong Sun, Jingjing Wang, Yangyang Kang, Shoushan Li, Xiaozhong Liu, Luo Si, Min Zhang, and Guodong Zhou. Sentiment classification towards question-answering with hierarchical matching network. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3654–3663, 2018.
 - [55] Marc Vaillant and Christos Davatzikos. Hierarchical matching of cortical features for deformable brain image registration. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 182–195. Springer, 1999.
 - [56] Zhe Lin, Larry S Davis, David Doermann, and Daniel DeMenthon. Hierarchical part-template matching for human detection and segmentation. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
 - [57] Pengfei Wang, Yu Fan, Shuzi Niu, Ze Yang, Yongfeng Zhang, and Jiafeng Guo. Hierarchical matching network for crime classification. In *proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 325–334, 2019.
 - [58] Timothy Gebhard, Niki Kilbertus, Giambattista Parascandolo, Ian Harry, and Bernhard Schölkopf. Convwave: Searching for gravitational waves with fully convolutional neural nets. In *Workshop on Deep Learning for Physical Sciences (DLPS) at the 31st Conference on Neural Information Processing Systems (NIPS)*, pages 1–6, 2017.
 - [59] Daniel George and EA Huerta. Deep learning for real-time gravitational wave detection and parameter estimation: Results with advanced ligo data. *Physics Letters B*, 778:64–70, 2018.
 - [60] Hunter Gabbard, Michael Williams, Fergus Hayes, and Chris Messenger. Matching matched filtering with deep networks for gravitational-wave astronomy. *Physical review letters*, 120(14):141103, 2018.
 - [61] Daniel George and EA Huerta. Deep neural networks to enable real-time multimessenger astrophysics. *Physical Review D*, 97(4):044039, 2018.
 - [62] XiLong Fan, Jin Li, Xin Li, YuanHong Zhong, and Jun-Wei Cao. Applying deep neural networks to the detection and space parameter estimation of compact binary coalescence with a network of gravitational wave detectors. *SCIENCE CHINA Physics, Mechanics & Astronomy*, 62(6):969512, 2019.
 - [63] Filip Morawski, Michal Bejger, and Pawel Ciecielag. Convolutional neural network classifier for the output of the time-domain-statistic all-sky search for continuous gravitational waves. *Machine Learning: Science and Technology*, 1(2):025016, 2020.
 - [64] Christoph Dreissigacker, Rahul Sharma, Chris Messenger, Ruining Zhao, and Reinhard Prix. Deep-learning continuous gravitational waves. *Physical Review D*, 100(4):044009, 2019.
 - [65] Plamen G Krastev. Real-time detection of gravitational waves from binary neutron stars using artificial neural networks. *Physics Letters B*, page 135330, 2020.
 - [66] Bai-Jiong Lin, Xiang-Ru Li, and Wo-Liang Yu. Binary neutron stars gravitational wave detection based on wavelet packet analysis and convolutional neural networks. *Frontiers of Physics*, 15(2):1–7, 2020.
 - [67] Yu-Chiung Lin and Jiun-Huei Protty Wu. Detection of gravitational waves using bayesian neural networks. *arXiv preprint arXiv:2007.04176*, 2020.
 - [68] Christopher Bresten and Jae-Hun Jung. Detection of gravitational waves using topological data analysis and convolutional neural network: An improved approach. *arXiv preprint arXiv:1910.08245*, 2019.
 - [69] P Astone, P Cerdá-Durán, I Di Palma, M Drago, F Muciaccia, C Palomba, and F Ricci. New method to observe gravitational waves emitted by core collapse supernovae. *Physical Review D*, 98(12):122002, 2018.
 - [70] Takahiro S Yamamoto and Takahiro Tanaka. Use of excess power method and convolutional neural network in all-sky search for continuous gravitational waves. *arXiv preprint arXiv:2011.12522*, 2020.
 - [71] Christoph Dreissigacker and Reinhard Prix. Deep-learning continuous gravitational waves: Multiple detectors and realistic noise. *Physical Review D*, 102(2):022005, 2020.
 - [72] Roberto Corizzo, Michelangelo Ceci, Eftim Zdravetski, and Nathalie Japkowicz. Scalable auto-encoders for gravitational waves detection from time series data. *Expert Systems with Applications*, 151:113378, 2020.
 - [73] Andrew L Miller, Pia Astone, Sabrina D’Antonio, Sergio Frasca, Giuseppe Intini, Iuri La Rosa, Paola Leaci, Simone Mastrogiovanni, Federico Muciaccia, Andonis Mitidis, et al. How effective is machine learning to detect long transient gravitational waves from neutron stars in a real search? *Physical Review D*, 100(6):062005, 2019.
 - [74] Joe Bayley, Chris Messenger, and Graham Woan. Robust machine learning algorithm to search for continuous gravitational waves. *Physical Review D*, 102(8):083024, 2020.
 - [75] Plamen G Krastev, Kiranjyot Gill, V Ashley Villar, and Edo Berger. Detection and parameter estimation of gravitational waves from binary neutron-star mergers in real ligo data using deep learning. *arXiv preprint arXiv:2012.13101*, 2020.
 - [76] Hua-Mei Luo, Wenbin Lin, Zu-Cheng Chen, and Qing-Guo Huang. Extraction of gravitational wave signals with optimized convolutional neural network. *Frontiers of Physics*, 15(1):1–6, 2020.
 - [77] Gerson R Santos, Marcela P Figueiredo, Antonio de Pádua Santos, Pavlos Protopapas, and Tiago AE Ferreira. Gravitational wave detection and information extraction via neural networks. *arXiv preprint arXiv:2003.09995*, 2020.
 - [78] Man Leong Chan, Ik Siong Heng, and Chris Messenger. Detection and classification of supernova gravitational wave signals: A deep learning approach. *Physical Review D*, 102(4):043022, 2020.
 - [79] Heming Xia, Lijing Shao, Junjie Zhao, and Zhou-jian Cao. Improved deep learning techniques in gravitational-wave data analysis. *Physical Review D*, 103(2):024040, 2021.
 - [80] Rahul Biswas, Lindy Blackburn, Junwei Cao, Reed Essick, Kari Alison Hodge, Erotokritos Katsavounidis, Kyungmin Kim, Young-Min Kim, Eric-Olivier Le Bigot, Chang-Hwan Lee, et al. Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data. *Physical Review D*, 88(6):062003, 2013.
 - [81] Daniel George, Hongyu Shen, and EA Huerta. Deep transfer learning: A new deep learning glitch classification method for advanced ligo. *arXiv preprint*

- arXiv:1706.07446*, 2017.
- [82] Nikhil Mukund, Sheelu Abraham, Shivaraj Kandhasamy, Sanjit Mitra, and Ninan Sajeeth Philip. Transient classification in ligo data using difference boosting neural network. *Physical Review D*, 95(10):104059, 2017.
 - [83] Massimiliano Razzano and Elena Cuoco. Image-based deep learning for classification of noise transients in gravitational wave detectors. *Classical and Quantum Gravity*, 35(9):095016, 2018.
 - [84] S Coughlin, S Bahaadini, N Rohani, M Zevin, O Patane, M Harandi, C Jackson, V Noroozi, S Allen, J Areeda, et al. Classifying the unknown: Discovering novel gravitational-wave detector glitches using similarity learning. *Physical Review D*, 99(8):082002, 2019.
 - [85] Robert E Colgan, K Rainer Corley, Yenson Lau, Imre Bartos, John N Wright, Zsuzsa Márka, and Szabolcs Márka. Efficient gravitational-wave glitch identification from environmental data through machine learning. *Physical Review D*, 101(10):102003, 2020.
 - [86] Hiroyuki Nakano, Tatsuya Narikawa, Ken-ichi Oohara, Kazuki Sakai, Hisa-aki Shinkai, Hirotaka Takahashi, Takahiro Tanaka, Nami Uchikata, Shun Yamamoto, and Takahiro S Yamamoto. Comparison of various methods to extract ringdown frequency from gravitational wave data. *Physical Review D*, 99(12):124032, 2019.
 - [87] Stephen R Green, Christine Simpson, and Jonathan Gair. Gravitational-wave parameter estimation with autoregressive neural network flows. *Physical Review D*, 102(10):104057, 2020.
 - [88] Juan Pablo Marulanda, Camilo Santa, and Antonio Enea Romano. Deep learning merger masses estimation from gravitational waves signals in the frequency domain. *Physics Letters B*, 810:135790, 2020.
 - [89] A Caramete, AI Constantinescu, LI Caramete, T Popescu, RA Balasov, D Felea, MV Rusu, P Stefanescu, and OM Tintareanu. Characterization of gravitational waves signals using neural networks. *arXiv preprint arXiv:2009.06109*, 2020.
 - [90] Arnaud Delaunoy, Antoine Wehenkel, Tanja Hinderer, Samaya Nissanke, Christoph Weniger, Andrew R Williamson, and Gilles Louppe. Lightning-fast gravitational wave parameter inference through neural amortization. *arXiv preprint arXiv:2010.12931*, 2020.
 - [91] Hongyu Shen, Daniel George, Eliu A Huerta, and Zhizhen Zhao. Denoising gravitational waves with enhanced deep recurrent denoising auto-encoders. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3237–3241. IEEE, 2019.
 - [92] Wei Wei and EA Huerta. Gravitational wave denoising of binary black hole mergers with deep learning. *Physics Letters B*, 800:135081, 2020.
 - [93] Timothy D Gebhard, Niki Kilbertus, Ian Harry, and Bernhard Schölkopf. Convolutional neural networks: A magic bullet for gravitational-wave detection? *Physical Review D*, 100(6):063015, 2019.
 - [94] Thomas Dent and John Veitch. Optimizing gravitational-wave searches for a population of coalescing binaries: Intrinsic parameters. *Physical Review D*, 89(6):062002, 2014.
 - [95] George Turin. An introduction to matched filters. *IRE transactions on Information theory*, 6(3):311–329, 1960.
 - [96] Philip M Woodward. *Probability and information theory, with applications to radar: international series of monographs on electronics and instrumentation*, volume 3. Elsevier, 2014.
 - [97] Bernard F Schutz. Gravitational wave astronomy. *Classical and Quantum Gravity*, 16(12A):A131, 1999.
 - [98] Gwosc summary. <https://ldas-jobs.ligo-la.caltech.edu/~detchar/summary/day/20170801/>
 - [99] Marco Cavaglia, Kai Staats, and Teerth Gill. Finding the origin of noise transients in ligo data with machine learning. *arXiv preprint arXiv:1812.05225*, 2018.
 - [100] Robert E. Colgan, K. Rainer Corley, Yenson Lau, Imre Bartos, John N. Wright, Zsuzsa Márka, and Szabolcs Márka. Efficient gravitational-wave glitch identification from environmental data through machine learning. *Phys. Rev. D*, 101(10):102003, May 2020.
 - [101] Robert E. Colgan, Jingkai Yan, Zsuzsa Márka, Imre Bartos, Szabolcs Márka, and John N. Wright. Architectural Optimization and Feature Learning for High-Dimensional Time Series Datasets. *arXiv e-prints*, page arXiv:2202.13486, February 2022.
 - [102] Robert E. Colgan, Zsuzsa Márka, Jingkai Yan, Imre Bartos, John N. Wright, and Szabolcs Márka. Detecting and Diagnosing Terrestrial Gravitational-Wave Mimics Through Feature Learning. *arXiv e-prints*, page arXiv:2203.05086, March 2022.
 - [103] Elena Cuoco, Jade Powell, Marco Cavaglià, Kendall Ackley, Michał Bejger, Chayan Chatterjee, Michael Coughlin, Scott Coughlin, Paul Easter, Reed Essick, Hunter Gabbard, Timothy Gebhard, Shaon Ghosh, Leïla Haegel, Alberto Iess, David Keitel, Zsuzsa Márka, Szabolcs Márka, Filip Morawski, Tri Nguyen, Rich Ormiston, Michael Pürrer, Massimiliano Razzano, Kai Staats, Gabriele Vajente, and Daniel Williams. Enhancing gravitational-wave science with machine learning. *Machine Learning: Science and Technology*, 2(1):011002, dec 2020.
 - [104] R. Abbott et al. Open data from the first and second observing runs of advanced ligo and advanced virgo. *SoftwareX*, 13:100658, 2021.
 - [105] GWOSC. Gravitational wave open science center.