Machine learning algorithm for minute-long burst searches

Vincent Boudart[*] and Maxime Fays[†]

STAR Institute, Bâtiment B5, Université de Liège, Sart Tilman B4000 Liège, Belgium

Minute-long gravitational wave (GW) transients are events lasting from a few to hundreds of seconds. In opposition to compact binary mergers, their GW signals cover a wide range of poorly understood astrophysical phenomena such as accretion disk instabilities and magnetar flares. The lack of accurate and rapidly generated gravitational-wave emission models prevents the use of matched filtering methods. Such events are thus probed through the template-free excess-power method, consisting in searching for a local excess of power in the time-frequency space correlated between detectors. The problem can be viewed as a search for high-value clustered pixels within an image, which has been generally tackled by deep learning algorithms such as convolutional neural networks (CNNs). In this work, we use a CNN as a anomaly detection tool for the long-duration searches. We show that it can reach a pixel-wise detection despite trained with minimal assumptions, while being able to retrieve both astrophysical signals and noise transients originating from instrumental coupling within the detectors. We also note that our neural network can extrapolate and connect partially disjoint signal tracks in the time-frequency plane.

I. INTRODUCTION

The first gravitational wave (GW) event, coming from the coalescence of two black holes, was detected by the Advanced LIGO [1] interferometers on September 14, 2015 [2]. Two years later, on August 17, 2017, the first binary neutron star (BNS) merger was observed [3], paving the way for the search for longer-duration gravitational waves. These two events are part of a broader class of signals called compact binary coalescence (CBC), which further includes recently observed black hole-neutron star collisions [4]. To this date, only CBC events, coming from powerful sources and being accurately modelled systems, have been detected [5]. However, new sources are expected to be observed given the planned sensitivity improvement of the Advanced LIGO and Advanced Virgo detectors [6]. Among the proposed candidates, unmodeled GW transients, known as bursts, cover a wide range of poorly understood astrophysical phenomena for which accurate waveforms are not available. Bursts include accretion-disk instabilities [7], nonaxisymmetric deformations in magnetars [8], supernovae [9], gamma-ray bursts [10], fallback accretion events [11]. The above-mentioned family of possible GW progenitor events contains both short- (< 2 seconds) and long-duration signals (between 2 and several hundreds of seconds). In this paper, we present a novel machine learning algorithm targeted at anomaly detection for long-duration burst searches (ALBUS).

Minute-long burst searches usually consist in finding an excess of power in the cross-correlated data of two or more detectors, known as correlated spectrograms or time-frequency (TF) maps. Algorithms that are searching for gravitational waves in these spectrograms can be classified in two categories: seed-based or seedless. Seed-based methods aim at clustering pixels above a predefined threshold while seedless algorithms are processing pixels derived from generic models. The current generation pipelines are the long-duration configuration of coherent WaveBurst (cWB) [12], the two different versions of the Stochastic Transient Analysis Multi-detector Pipeline—All Sky (STAMP-AS), Zebragard and Lonetrack [13,14], PySTAMPAS [15], cocoA [16], and X-SphRad [17]. cWB, PySTAMPAS, and X-SphRad are seed-based algorithms while the two STAMP-AS pipelines, Zebragard and Lonetrack, use seed-based and seedless clustering algorithms respectively. As well as Lonetrack, cocoA is a seedless pipeline. In opposition to CBC searches, no machine-learning based algorithm has yet been applied to long-duration searches. The uncertainties in the existing physical models of long-duration transients lead us to make minimal assumptions on progenitor sky-position, inclination, time-of-arrival and GW waveform characteristics. The waveform models cannot consequently be taken as accurate patterns to be recognized and are thus used as tests for algorithms and pipelines rather than actual targets of the search. The present work aims at circumventing this problem by mimicking template database of generic long-duration signals, allowing us to take

[*]vboudart@uliege.be
[†]Maxime.Fays@uliege.be

advantage of the speed and the robustness of convolutional neural networks.

In Sec. II, we describe the data-generation method as well as the strategy to mimic long-duration signals. In Sec. III, details about the architecture and the training method are given. We present the results of the training and the detection performances both on injected signals and on background spectrograms in Sec. IV. Section V is finally dedicated to discussions and conclusions.

II. DATA

A. Time-frequency map generation

One of the detection methods of long-duration GWs is based on the excess of power method [18]. It consists in correlating the output of a least two different detectors, under the assumption that their noise is uncorrelated. This produces coherent spectrograms where GW signals are represented as minute-long high-correlation patterns. The problem of detecting long-duration GWs can therefore be reduced to finding a cluster of pixels with high intensity in background noise.

The exact formulation of the coherence between two signals x and y is:

$$C_{xy}(f) = \frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)}, \quad (1)$$

where $G_{xy}(f)$ is the cross-spectral density (CSD) between signals x and y, and $G_{xx}(f)$ and $G_{yy}(f)$ the power spectral density (PSD) of x and y respectively. The cross-spectral density is defined as:

$$G_{xy}(f) = \int_{-\infty}^{\infty} \left[\lim_{T \to \infty} \frac{1}{T} \int_{-\infty}^{\infty} x_T(t-\tau) y_T(t) dt \right] \times e^{-i2\pi f \tau} d\tau. \quad (2)$$

$G_{xx}(f)$ and $G_{yy}(f)$ are then a particular case of (2) where respectively y is replaced by x and conversely. As the latter expression is intractable for a real signal, we employ Welch's method [19] as an approximation method. Evaluating (1) at several sampled frequencies then leads to the generation of a vector of coherence values versus frequency bins. To generate a full time-frequency array, we apply Welch's method to small subsets of our original signal. This is equivalent to updating the coherence value after some time and compiling this time evolution as a single map. The resulting spectrogram is finally whitened along the time axis by summing all the values of a frequency bin along the time axis and dividing by the total sum. An example of the obtained spectrogram is shown in Fig. 1.

To constitute a sufficient number of background spectrograms, we use time-slides [20] where the detector data are shifted by time delays larger than the time of flight of GWs between detectors. The time delays are also larger than the duration of target signals (≥ 500 s) to guarantee our cross-correlated data to contain only detector noise.

The time and frequency resolutions of the generated spectrograms have an impact on the sensitivity of the search. The longer the time segments, the more GW energy will be accumulated in a single pixel, leading to a higher coherence. As the noise appears to be coherent on very small timescales ($\ll 1$ second) [21], increasing the length of the time segments allows to reduce the coherence of the noise versus the coherence of hypothetical signals, yielding an increased signal-to-noise ratio (SNR). However, longer time bins will
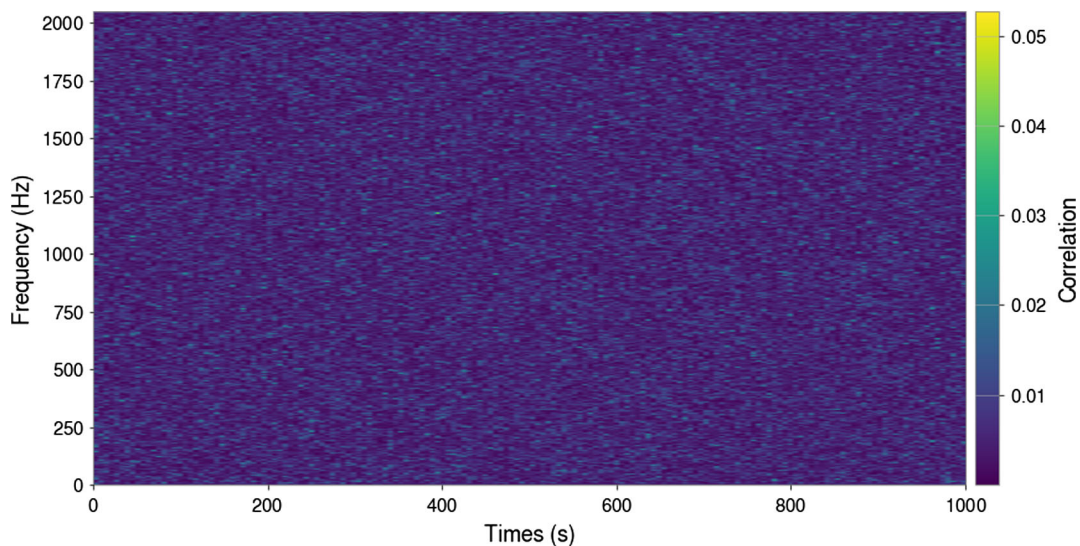

FIG. 1.   Typical time-frequency map of cross-correlated O3a background noise. The GPS time at the start of the Advanced LIGO Hanford (H1) and Advanced LIGO Livingston (L1) data is respectively 1246174396 and 1246108524.

cause short signals (∼10 seconds) to fall into very few pixels, making them harder to detect. An identical reasoning can be made for the frequency resolution. In this work, we use a 6-second time resolution combined with a 2 Hertz frequency bin as a good compromise. Taking a 1000s data stream from the Advanced LIGO interferometers spanning frequencies up to 2048 Hz (see Fig. 1), results in spectrograms with dimensions of 166×1025 pixels.

The time-frequency arrays are heavy to store in the default 32-bit *Python* precision because of their large amount ($> 10^5$) of pixels. Each TF map weights about 13 megabytes. This substantial storage space can reduce the maximum batch size, i.e., the number of training samples that are passed through the network in a single forward pass, accessible for the training, eventually increase the training time and downgrade the test performance [22]. It is also recommended to use batch sizes above 10 to avoid a highly noisy gradient descent, and to take advantage of the speed-up of matrix-matrix products over matrix-vector products [22]. A noisy gradient descent is obtained when the gradients obtained after each training iteration poorly generalizes to all the samples in the dataset and in fine extend the time needed for the network to converge. Saving the spectrograms as RGB images (8-bit integers) allows to reduce the storage needed at the cost of a small loss of precision in the values of the array. The maximum loss of precision for values falling right in between two integer levels is $(1/2^8)/2$, which is less than 0.2%. We thus have 3 channels displaying different information depending on the color map used to draw the initial array. We choose the "cubehelix" colormap from the *Matplotlib Python* library [23], displaying the GW signals clearly in all 3 channels, as seen in Fig. 2. The final reduction factor

in memory is roughly 26 compared to the 32-bit *NumPy* arrays [24], ultimately giving access to larger training batches. A further argument in favour of using RGB images is the wide use of the format in deep learning applications, e.g., some neural networks even require a 3-channel image as input [25–27].

Therefore, the final dataset that we consider is made up of RGB time-frequency images generated from the coherence between the data from the Advanced LIGO Hanford (H1) and Advanced LIGO Livingston (L1) detectors, gathered during the first phase of the third advanced LIGO observing run (O3a).

B. A new measure of excess power

In order to form our dataset, we need to inject signals so that our neural network can learn to recognize their patterns in spectrograms. The current way of adding long-duration burst signals to background noise makes use of the root sum squared value of the strain $h(t)$ to gauge the strength of the injection. It is defined as:

$$h_{rss} = \sqrt{\int h(t)^2 dt}. \tag{3}$$

However, this expression only depends on the strain value of the injected signal, and makes the human eye visibility of the injection vary with the local noise level of the background map. Figure 3 shows a waveform model injected with a h_{rss} value of 510^{-22} in two different O3a background spectrograms. When the noise level is sufficiently high, as in the right panel of Fig. 3, the signal can be buried in the background noise, preventing a clear detection for the
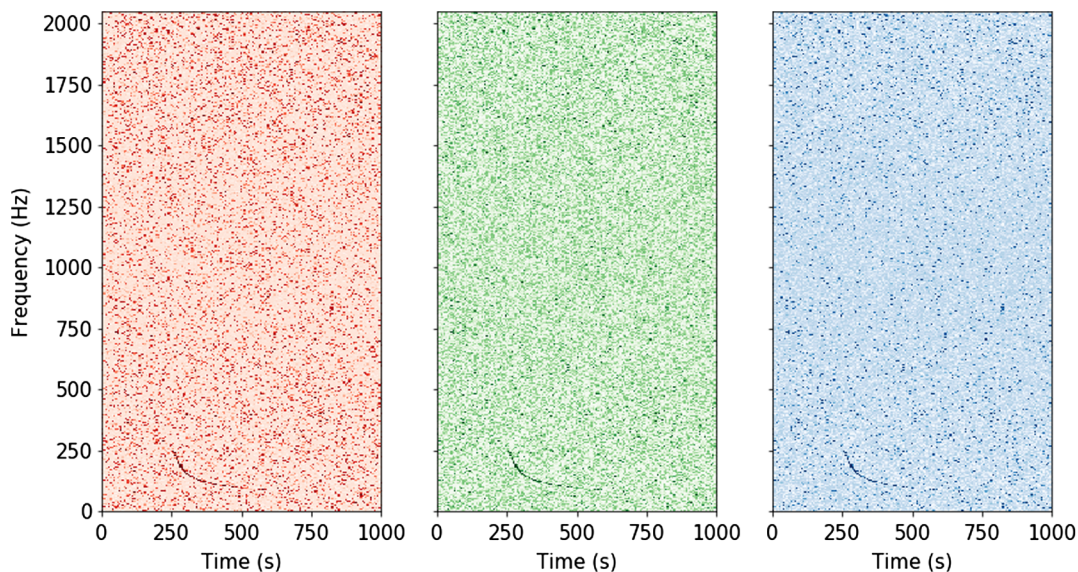

FIG. 2.   RGB channels of an 03a time-frequency background map where a *GRBplateau* waveform [28] is injected. Note how the signal is visible in all 3 channels, from 250s to 550s. The GPS time at the start of the H1 and L1 data is respectively 1246174396 and 1246108524.
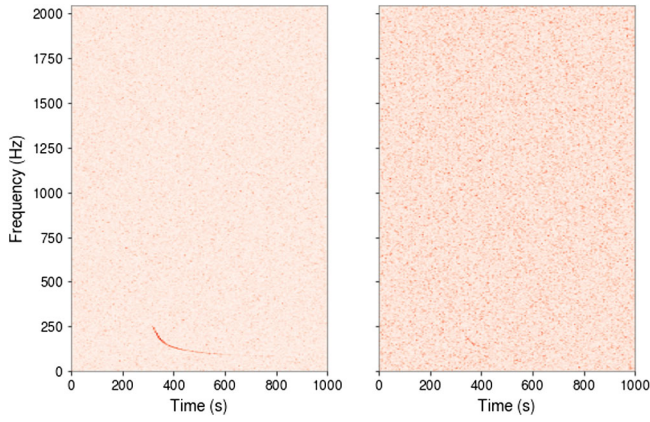
FIG. 3.   Injection of a *GRBplateau* waveform from [28] at 300s with a h_{rss} value of 510^{-22} in two different O3a background spectrograms. The GPS time at the start of the H1 and L1 data in the left panel is respectively 1246174396 and 1246108524, while for the right panel it is 1248305006 and 1248273184.

network. Such samples can fool a neural network during a supervised learning and eventually cause the training to be badly conditioned [29]. A new criterion is therefore required to form a healthy dataset.

To be visible in time-frequency representations, an injected signal has to stand out of the local noise level. Defining a noise-only spectrogram N_{ij} and the same spectrogram to which a signal has been injected by S_{ij}, the new criterion is

$$V = \sum_{i,j}(S_{ij} - N_{ij}),   \quad (4)$$

where the sum is carried over all the pixels (i, j) in the map. We call this new criterion "visibility". The pixel-to-pixel difference allows to fine tune how much signal emerges from the local noise level and form a dataset with different levels of intensity. This is particularly useful when training procedures like curriculum learning [30] are used.

C. Mimicking long-duration signals

Neural networks are particularly good at recognizing and classifying shapes and objects they have seen in training (YOLO [25], AlexNet [31], GoogLeNet [32], etc). They are therefore well suited for detecting signals in time-frequency images if the training set is sufficiently close to the expected long-duration GWs. Figure 4 shows the models used to test the long-duration pipelines for the third advanced LIGO-Virgo observing run. The waveforms generally show a *chirp up* or *chirp down* behavior. This property can be easily mimicked thanks to the SCIPY PYTHON package [33]. Specifically, it allows us to draw time series with varying parameters such as the duration, the frequency bandwidth, and the frequency evolution. The frequency evolution can either be linear, hyperbolic, quadratic or logarithmic. Once the chirp signal is generated, the energy distribution is adjusted thanks to a Kaiser filter [34] as seen in Fig. 5. For this, a window twice as long as the signal is generated and the chirp signal is multiplied by either the first or second half of the window. When the first half is implied, the final chirp shows a increasing energy distribution and the reverse holds
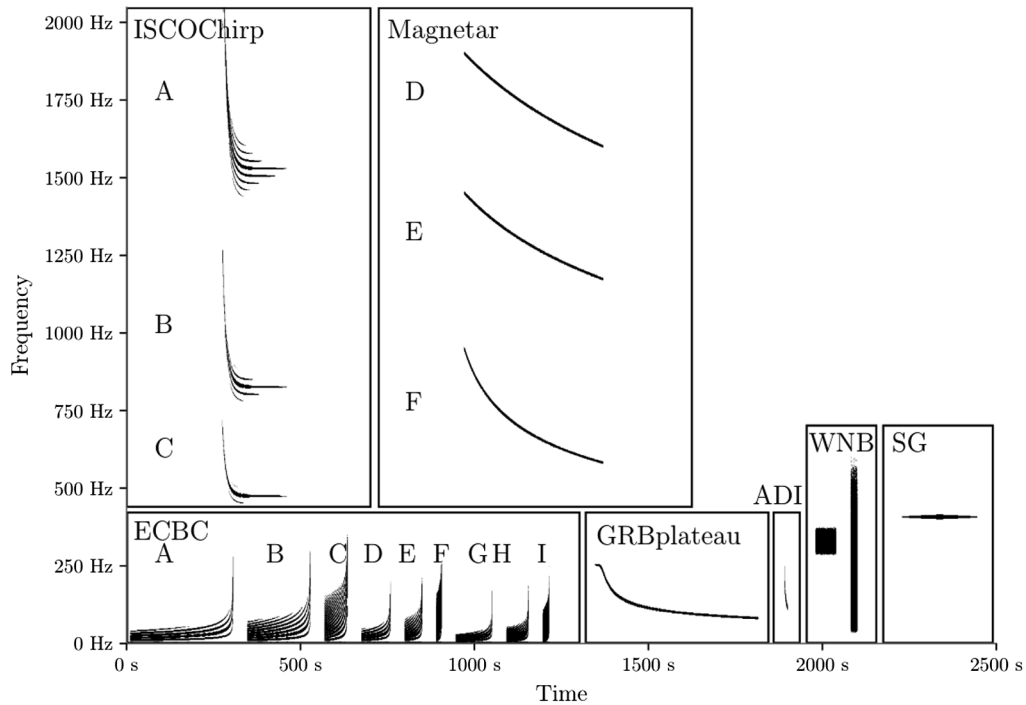

FIG. 4.   Signal models (waveforms) targeted by long-duration pipelines for the third observing run [35].
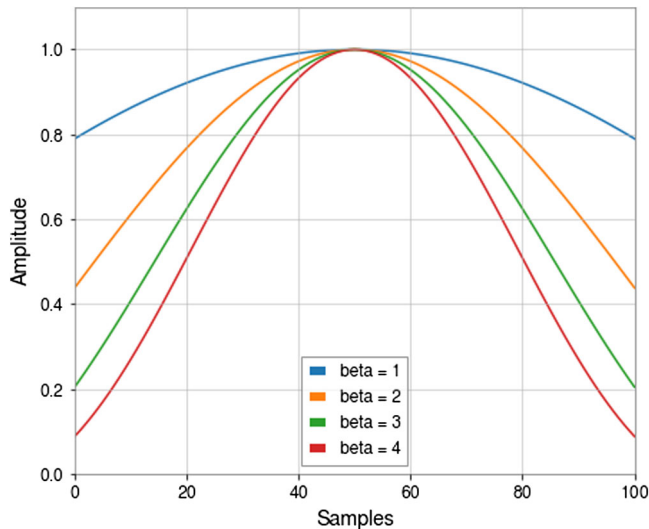
FIG. 5. Examples of Kaiser windows with different values for the β parameter.

when the second half is concerned. A shape parameter β is used to control the width of the Kaiser window. A larger beta parameter implies a narrower Kaiser window leading to a more unbalanced energy distribution, the signal showing more intensity in the beginning or in the end of the chirp structure.

Figure 6 shows examples of chirp signals. These chirps are nearly identical to the models that have been selected to test long-duration algorithms. However, the harmonics that appear in GW emission models like [36] or [37] cannot be reproduced with the Scipy library [33]. These harmonics come from emission mechanisms such as multiple mass moments in the torus around black holes [38] or eccentricity

oscillations in eccentric compact binary coalescences [37]. They usually show less power than the main component of the gravitational wave and show up exclusively for high amplitude injections, usually easily recognized in time-frequency images. We will thus make use of [33] to generate our dataset and train our neural network without *a priori* knowledge of the expected burst signals.

III. DEEP LEARNING APPROACH

A. Neural network in GW applications

Neural networks and machine learning techniques have been recently applied to gravitational wave physics [39]. Among the variety of networks used, convolutional neural networks (CNNs) [40], being very good at pattern and shape detection [25,31], have shown successful applications in the detection of black hole collisions [41], identification of the GW counterpart from supernovae [42], binary neutron star detection [43,44] and estimation of their parameters [45], as well as in classification of detector glitches [46]. CNNs have also been used in short burst detection with [47].

The strong capability of CNNs to recognize patterns can either be applied to one-dimensional (time series) [41,43,44] or two-dimensional data [32,46]. Their efficiency to detect shapes has even been adapted to generative adversarial networks (GAN) [48]. Such networks have recently been used to generate short duration bursts [49].

Given the performance, speed and wide application of machine learning techniques in the gravitational wave domain, we decide to apply it to the long duration search. Specifically, we use CNNs to detect and precisely localize chirp signals in the time-frequency space.
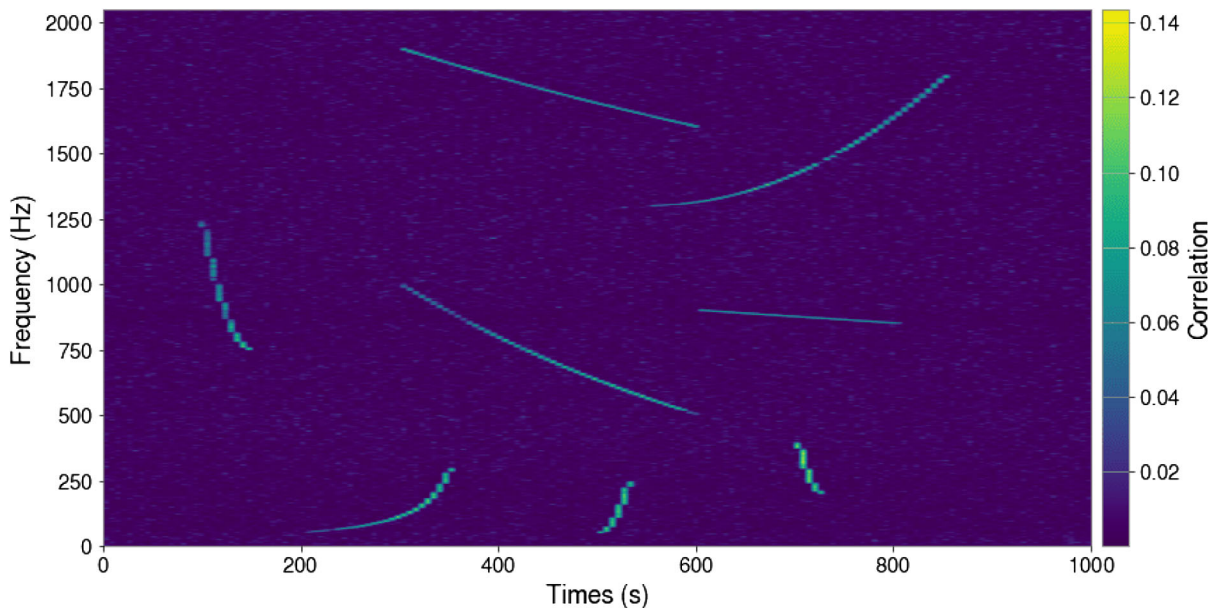

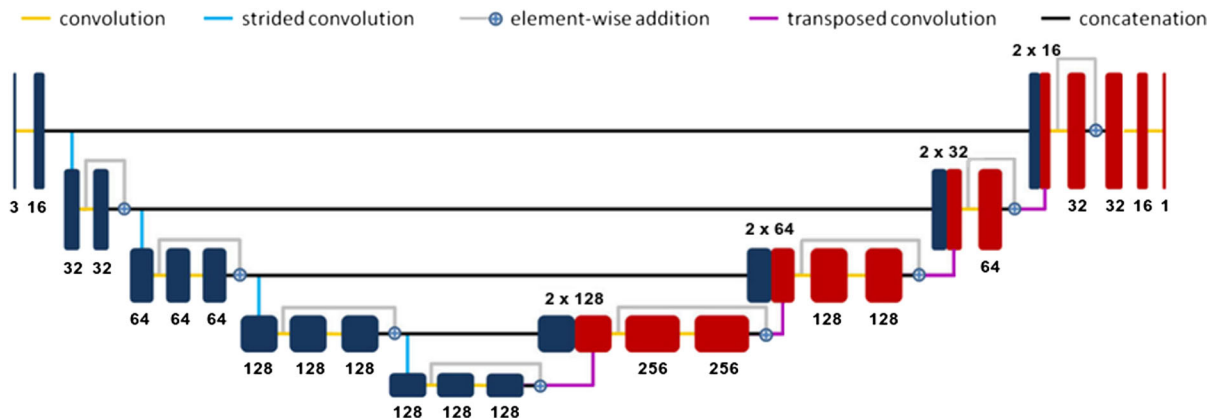FIG. 6. Examples of chirp signals produced for the training of our neural network.

FIG. 7. Architecture of ALBUS, modified from [50]. The downscaling and upscaling parts are represented in blue and red respectively. These two parts are coupled thanks to skipped connections, represented as concatenation lines. The numbers in black indicate the number of filters used at each stage of the network.

B. ALBUS

In order to detect signatures of long-duration bursts, we will also make use of CNNs. Most of the CNNs used to detect patterns and objects also involve a classification task [25,26,31]. However, we want to highlight the pixels that could resemble to burst signals rather than assigning a label to the whole input map. We found that the convolutional network built in [50] returns a pixel-by-pixel localization map. They associated a boolean target map to every image in the training set. The training method then consists in minimizing a loss between the output map and the target map, so that the former keeps approaching the latter as the training progresses. We will follow the same strategy but also bring some modifications, notably on the network architecture and on the definition of the target map.

The network, shown in Fig. 7, is made up of two parts, a downscaling part that keeps the useful information through its different layers, and an upscaling part that aims at localizing precisely this information in a map with the same dimensions as the input of the network. The connections between the downscaling and upscaling parts help the network to learn how to precisely position the signals. The number of filters at every step, indicated below each layer in Fig. 7, has been divided by 2 compared to [50], reducing both the training time and the memory usage.

The target map definition has also been modified. The localization map alone is not sufficient to rank the spectrograms based on their content. As an example, a score can be defined as the sum of the pixels in the localization map, which helps distinguishing GW candidates from noise-only images. If we keep the definition of [50], all the pixels that will be highlighted by the network will be put to 1 in the localization map. Therefore, summing up all the 1's in the image can lead to a high score, even if the pixels are scattered in the map. In such a case, it becomes harder to identify GW events through a unique score. We rather need a definition that can output both high and low values

depending on the intensity of the signal injected in the image. We set a threshold on the spectrogram pixels corresponding to the 99th percentile of the values. This is equivalent to keeping the top 1% pixels showing the highest values. We then normalize the spectrogram. This procedure leads to a target map that follows the intensity evolution of the signals through the input map. An example of a spectrogram containing a chirp signal and its corresponding target map can be seen in Fig. 8.

The last adaptation concerns the loss used to train our neural network. The loss defined by [50] is a weighted mean squared error (MSE):

$$L = \frac{1}{2}\sum_{i,j}(T_{ij} + \lambda \bar{T})(T_{ij} - O_{ij})^2, \qquad (5)$$

where T and O respectively stand for the target and output maps, $\bar{T}$ is the mean of the target map and λ is a control
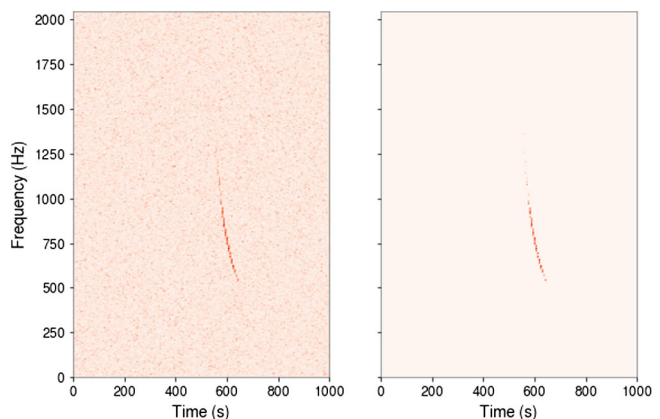

FIG. 8. Background spectrogram in which a chirp signal is injected at 100s (left) and its associated target map used for the training phase (right). The energy distribution of the chirp signal can also be seen in the target map, providing adequate samples for the training phase.

parameter. However, our dataset is also composed of background images where no signal has been injected, leading to empty target maps. It is essential to add them in the training loop so that the network will not automatically try to find something in the input TF maps. With empty maps, the left parenthesis in (5) is null, preventing the loss to give any feedback to the network in the backpropagation loop, which does not happen with the classical MSE loss, defined as:

$$MSE = \frac{1}{2}\sum_{i,j}(T_{ij} - O_{ij})^2. \tag{6}$$

As the MSE loss still gives a nonzero response with background images for which T_{ij} is zero everywhere, we decided to choose it for the training of our neural network.

IV. RESULTS

A. Training ALBUS

The dataset for the training is composed of 4500 background images and 4500 chirp images. The chirp signals were injected with 9 levels of *visibility* (500 samples for each intensity). All the parameters for injecting chirp signals are summarized in Table I. The validation set is made of 10% of both the background and injection dataset. The delay indicates the time from the start of the spectrogram where signals are injected. We set a low frequency threshold at 30 Hz because of the high noise level of the Advanced LIGO detectors at lower frequencies [21]. The chirp signals being drawn at a chosen h_{rss} value, the visibility level is evaluated after every injection. An iterative loop then allows to obtain the desired visibility levels by adapting accordingly the initial h_{rss} value. Because of this iterative loop, we tolerate a $\pm 10\%$ range around the selected values to cover a wider space and to converge faster.

The training algorithms have all been coded with PYTORCH [51]. The ADAM optimizer [52] has been chosen with a learning rate of 10^{-4}. The batch size is set to 20 where one half is taken from the background images and the other half from the injection images. The training and

TABLE I. Parameters used to inject chirp signals in the TF maps. All the parameters are uniformly drawn from their range of values.

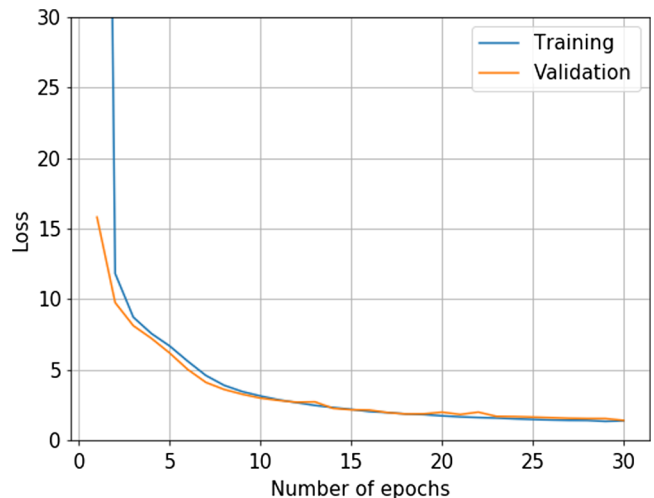|	Range of values
Duration	10–500 s
Delay	0–500 s
Frequency range	30–2000 Hz
Frequency evolution	Linear quadratic, logarithmic or hyperbolic
β parameter	1–4
Visibility levels	12, 14, 16, 18, 20, 30, 40, 50, 60

FIG. 9. Training and validation losses for a 30-epoch training of ALBUS.

validation losses for a training phase of 30 epochs are shown in Fig. 9. The training loss decreases monotonically which suggests that the learning progresses evenly. The validation loss remains in close vicinity of the training loss ruling out any overfitting on the training data. We decided to stop the training after 30 epochs because both losses started to reach a plateau, indicating that no major improvements are made by the network.

B. Detection performance

Figure 10 shows the output of ALBUS for 4 different waveforms from the selected models in [35]. The simulated signals are well recognized and the variation of intensity in the input map is also seen in the localization map.

An additional remark can be made concerning the upper right panel of Fig. 10, where a few pixels above the curve (around 600 Hz and 300 seconds) are highlighted in the output map. This behavior is also observed in the lower right panel. Indeed, our network is not only looking at the pixels having a high value but also at the connectivity between these pixels. It then naturally looks prolongs the main structure to catch pixels following the general trend of the signal. Such a propriety can be a relevant tool to reject background images showing isolated hot pixels.

Another detection capability is observed with transient noises called glitches. Glitches are appearing in the detector data in abundant quantities, produced due to several sources such as instruments or the environment [54,55]. Several classes of these artefacts have been identified through machine learning algorithms [46] and all show particular time-frequency morphologies. As most of the glitches last less than 6 seconds [21], they show up in our TF maps as straight vertical lines. The cross-correlation method reduces their impact since a glitch
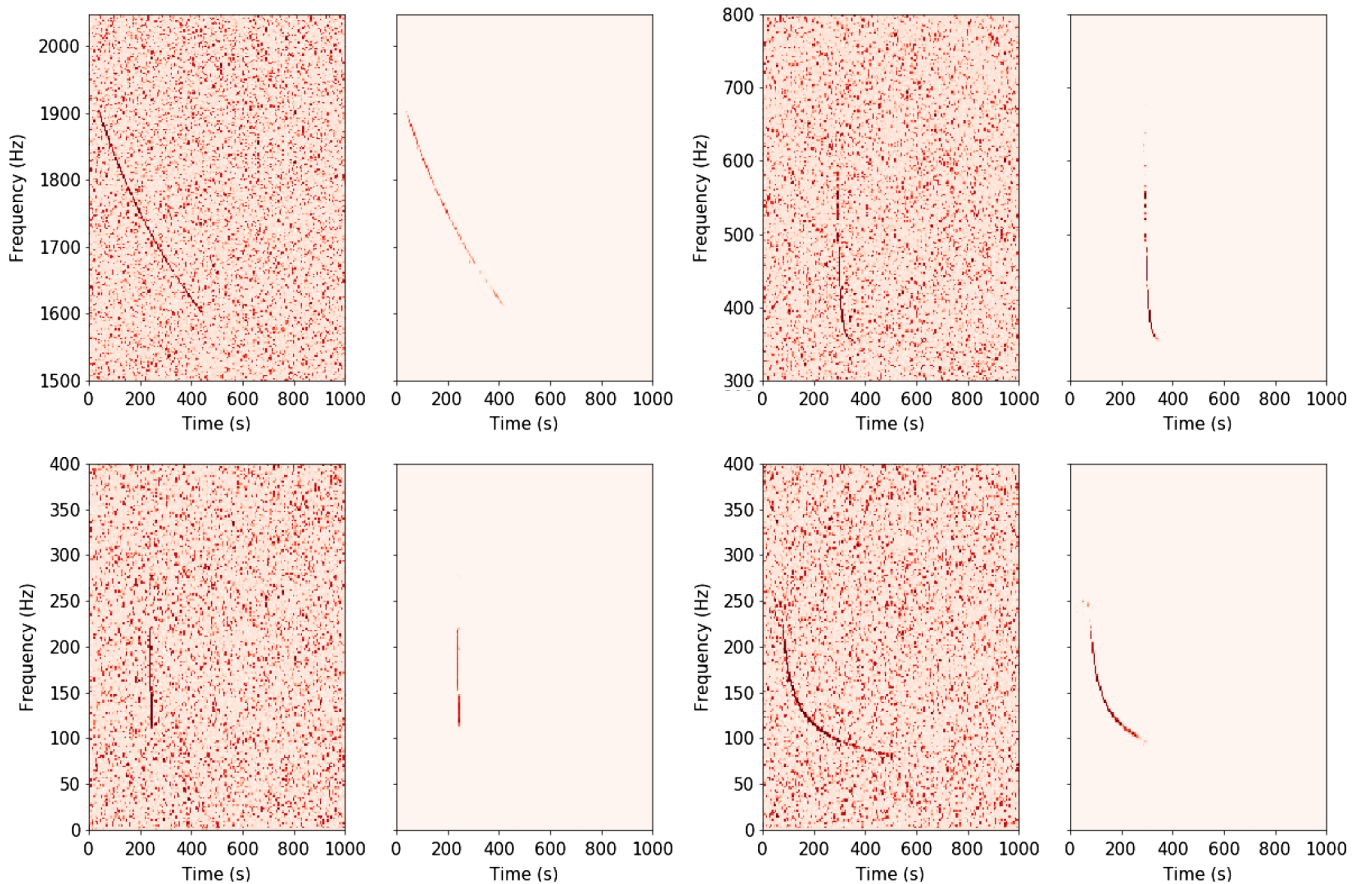
FIG. 10. Examples of detection performance on long-duration waveforms (top left: *Magnetar-D* [53], top right: *ISCOchirp-C* [36], bottom left: *ADI* [7] and bottom right: *GRBplateau* [28]). The left image of each panel is the red channel of the input image and the right panel shows the output of ALBUS.

from L1 data needs to fall in the same pixel as another glitch from H1 data to show up in the correlated TF map. That small number of cross-correlated glitches in the background can explain why ALBUS does not consider them as part of the background noise and actually detects

them. Figure 11 shows an example of a glitch and its localization by ALBUS.

The final point concerns harmonics of waveform models. It is reasonable to foresee that ALBUS will not detect them as they are not showing up in the training
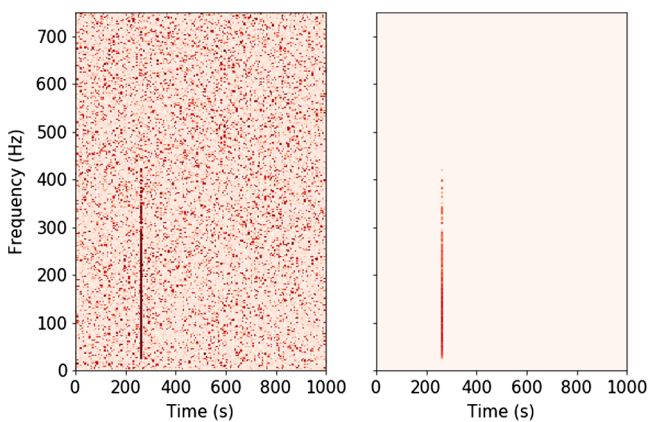

FIG. 11. Example of glitch detection. The left image shows the red channel of the input map while the right panel shows the output of ALBUS.
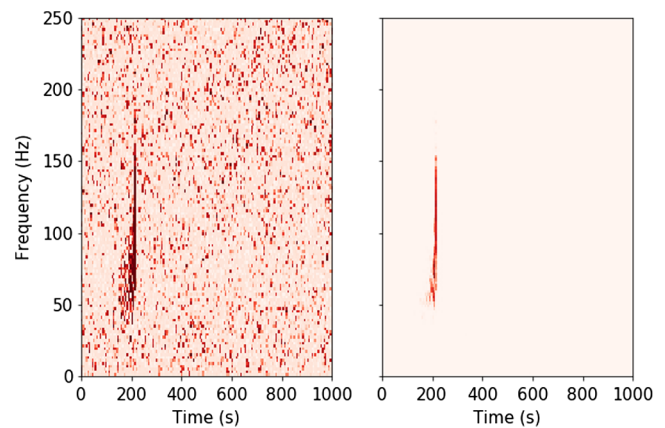

FIG. 12. Example of waveform showing harmonics (*ECBC-C*). The left image shows the red channel of the input map while the right panel shows the output of ALBUS.

dataset. Figure 12 displays this phenomenon to some extent. The harmonics that appear to the left of the rising chirp are not found in the output map except in the very beginning of the signal around 50 Hz. This effect is certainly due to the extrapolation capability mentioned above, for which ALBUS is "looking" for smooth connected signals. The choice of not incorporating harmonics of chirps in the data has minimal consequence since the waveform models showing these harmonics are still detected.

The detection performance as a function of the visibility levels can be seen in Figs. 13 and 14 for two different long-duration waveforms. The minimal visibility level to which ALBUS can identify most of the signal lies between 16 and 12, which is consistent with the perception of the human eye with no *a priori* knowledge of the injected signal.

C. Background analysis

Figure 15 compares the output of ALBUS for 3 different background images to its output when a long-duration waveform is injected. The output map shows correlation values smaller than 0.1. This trend is observed for all the processed spectrograms, with the exception of some isolated hot pixels that can show values up to 0.5. In any case, the highlighted pixels appear sparse and unconnected for background spectrograms, confirming
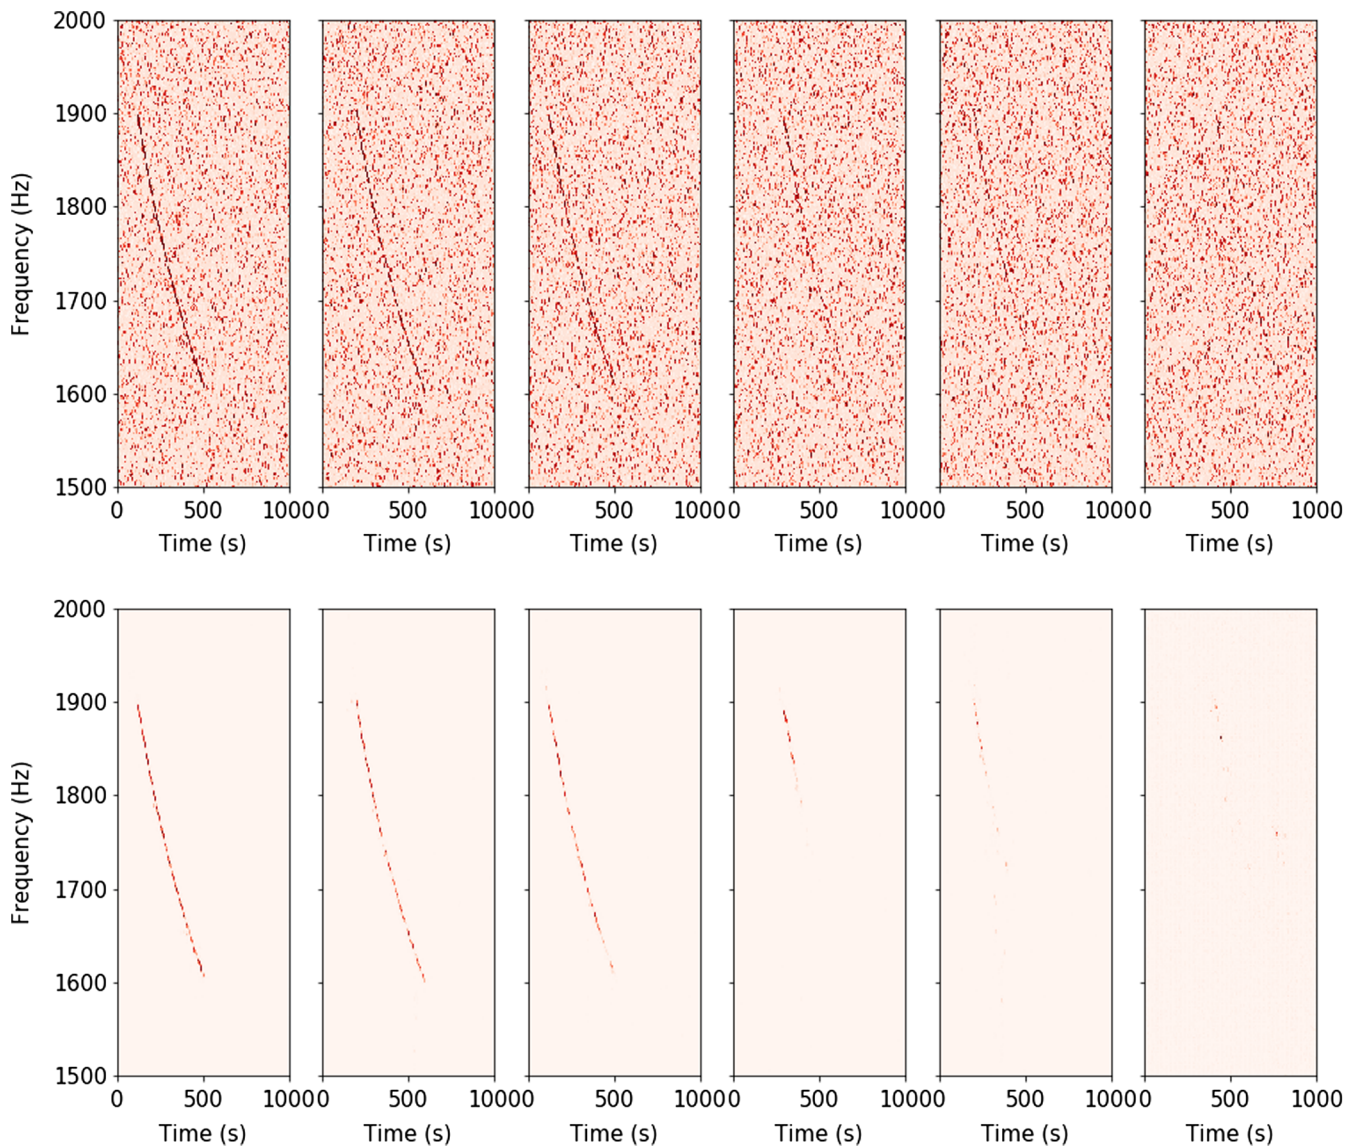

FIG. 13.   Detection performance of ALBUS on 6 different visibility levels (from left to right: 60, 40, 30, 20, 16, and 12) for the waveform model *Magnetar-D*. The top panel shows the input images and the lower panel shows the output of ALBUS.
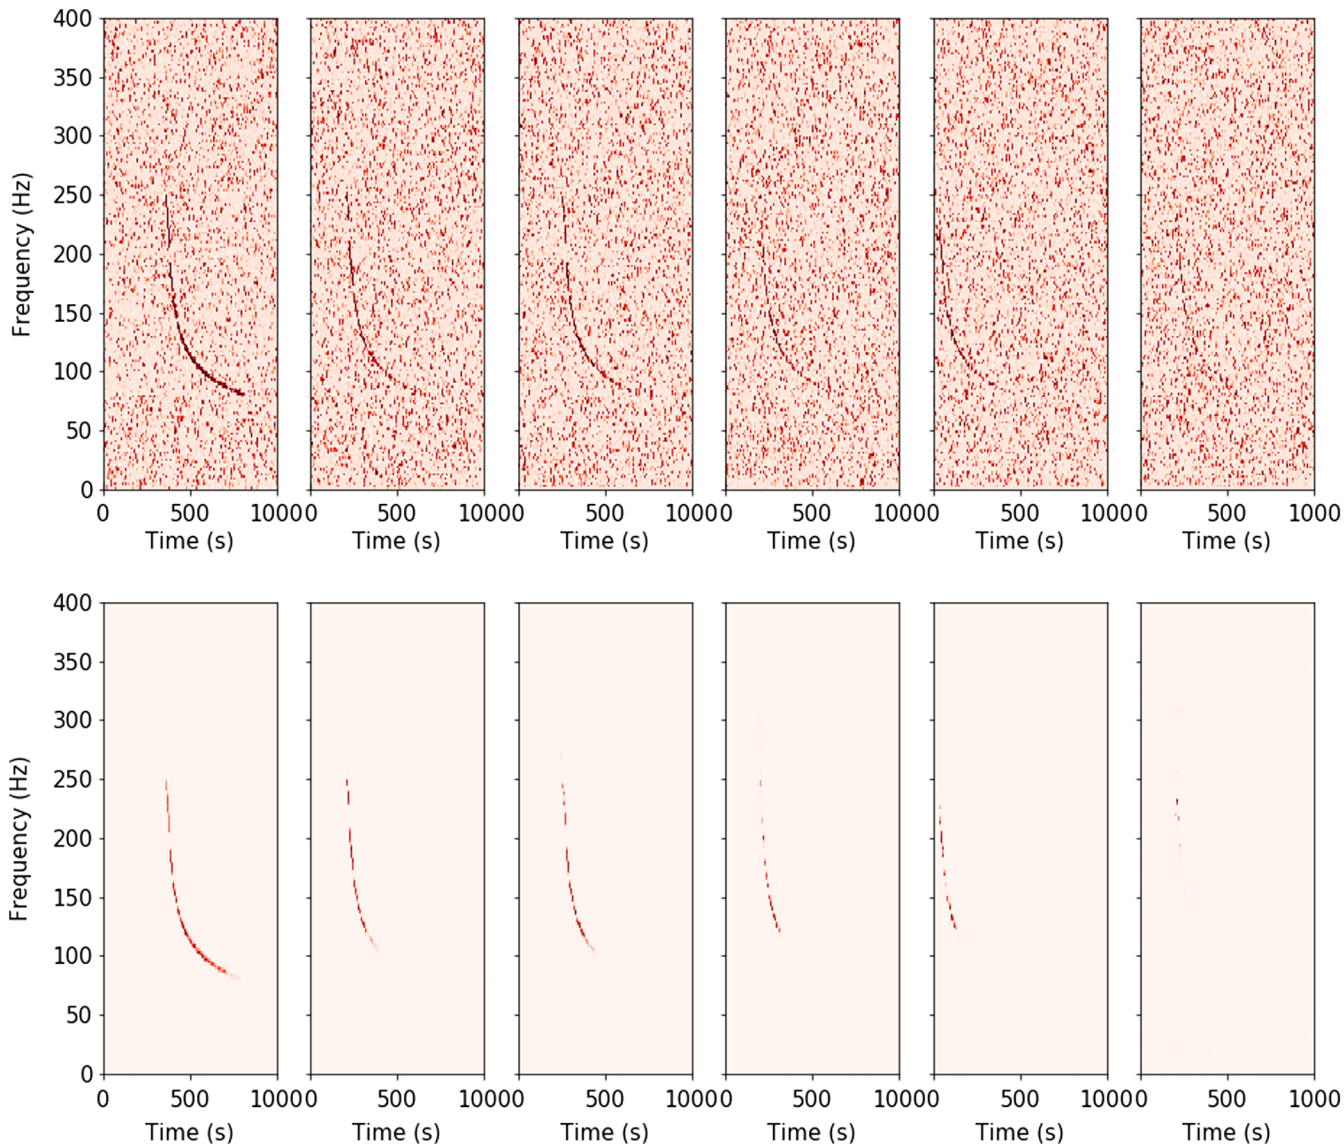
FIG. 14.   Detection performance of ALBUS on 6 different visibility levels (from left to right: 60, 40, 30, 20, 16, and 12) for the waveform model *GRBplateau*. The top panel shows the input images and the lower panel shows the output of ALBUS.

that our network is searching for connectivity among high value pixels.

V. DISCUSSION AND CONCLUSION

We have shown that convolutional neural networks can be applied to the search for minute-long gravitational wave transients in the time-frequency space of the cross-correlated LIGO noise. Our approach allows a fast and pixel-precise identification of the long-duration signals with no training on the latter. The speed of neural networks naturally allows to extend our algorithm to low-latency searches. With training on data accumulated during the first month of the observing run, after which the network would be active for the rest of the run. The whole process including the data acquisition,

the spectrogram generation and the forward pass to ALBUS can be carried out in just a few seconds. A low-latency implementation will then consist in repeating this process periodically.

In a low-latency implementation, coherent glitches appearing in both Hanford and Livingston interferometers will be detected. There is a need to remove these in order to avoid sending false alarms to other astronomers. Among the low-latency tools to identify and classify glitches, GravitySpy [46] and Omicron [56] can produce triggers in few minutes. We could make use of these triggers to discriminate cross-correlated glitches from short signals like ADI and ECBC (see Fig. 4), that have the potential to be misclassified as a glitch. An alternative
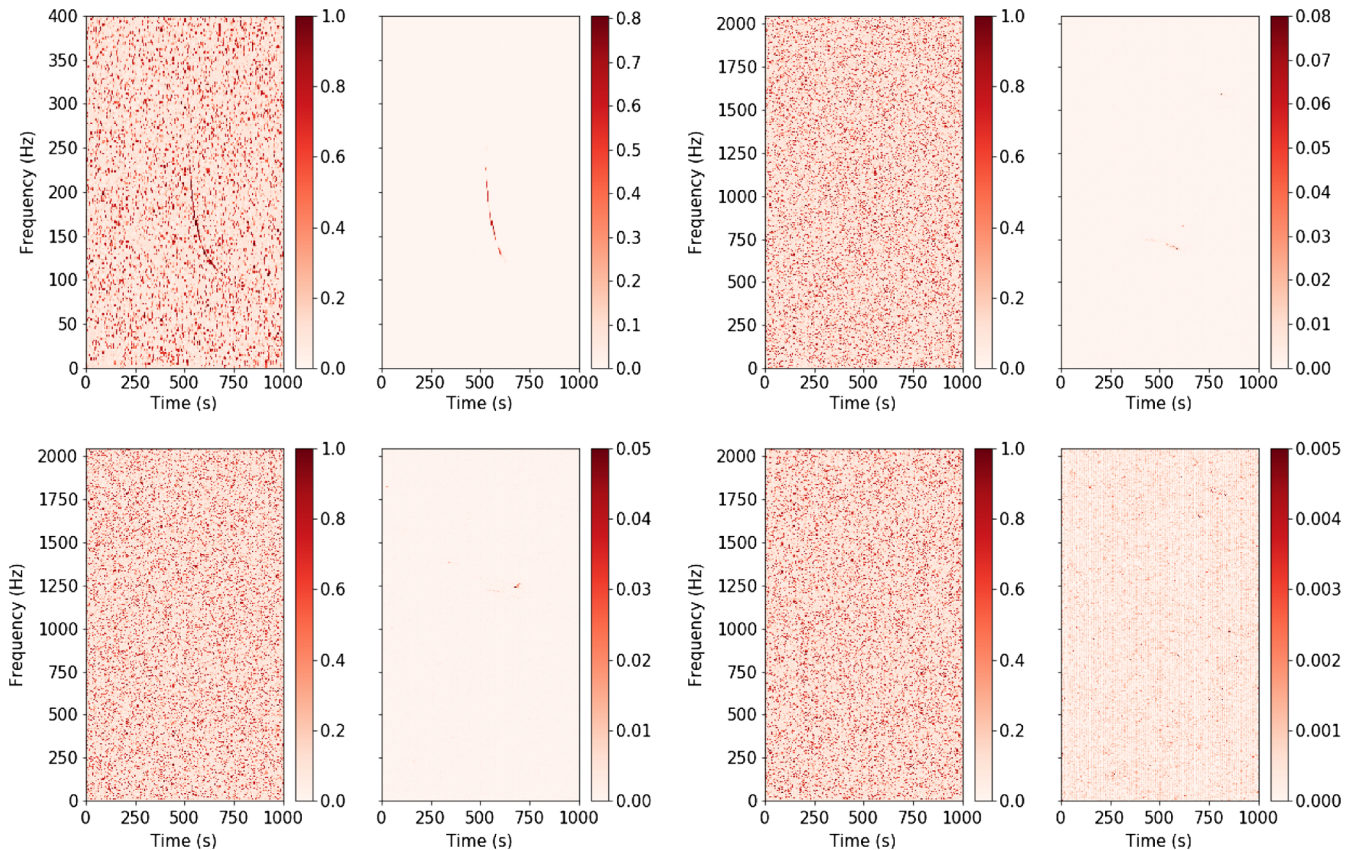
FIG. 15. Detection performances of ALBUS on background spectrograms. The top left panel displays the output map for a long-duration signal (*GRBplateau*) for comparison. The left image of each panel is the red channel of the input image and the right panel shows the output of ALBUS.

is to train a second network to recognize only coherent glitches, appearing as vertical lines, and run it in parallel to the current algorithm. This new network would then serve as a check for glitches.

The threshold for the detection of long-duration signals is determined by the highest background candidates, i.e., the background candidates that show the highest detection score as defined by a particular pipeline. Usually, the highest candidates are identified after analyzing at least 50 years of background, making more than 1 million spectrograms to process [35]. In order to rank these candidates and automate the detection, a detection statistics needs to be defined in follow-up works.

[1] J. Aasi *et al.*, Classical Quantum Gravity **32**, 115012 (2015).
[2] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116**, 061102 (2016).
[3] B. P. Abbott *et al.*, Phys. Rev. Lett. **119**, 161101 (2017).
[4] R. Abbott *et al.*, Astrophys. J. Lett. **915**, L5 (2021).
[5] R. Abbott *et al.* (The LIGO Scientific, the Virgo, the KAGRA Collaborations), arXiv:2111.03606.
[6] F. Acernese *et al.*, Classical Quantum Gravity **32**, 024001 (2015).
[7] M. H. P. M. van Putten, Phys. Rev. Lett. **87**, 091101 (2001).

[8] S. Dall'Osso, B. Giacomazzo, R. Perna, and L. Stella, Astrophys. J. **798,** 25 (2015).

[9] M. H. P. M. van Putten, A. Levinson, F. Frontera, C. Guidorzi, L. Amati, and M. Della Valle, Eur. Phys. J. Plus **134,** 537 (2019).

[10] A. Corsi and P. Meszaros, Astrophys. J. **702,** 1171 (2009).

[11] A. Piro and E. Thrane, Astrophys. J. **761,** 63 (2012).

[12] S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. F. Da Silva et al., Phys. Rev. D **93,** 042004 (2016).

[13] E. Thrane and M. Coughlin, Phys. Rev. Lett. **115,** 181102 (2015).

[14] T. Prestegard, University of Minnesota thesis, 2016, https://hdl.handle.net/11299/182183.

[15] A. Macquet, M. A. Bizouard, N. Christensen, and M. Coughlin, Phys. Rev. D **104,** 102005 (2021).

[16] R. Coyne, A. Corsi, and B. J. Owen, Phys. Rev. D **93,** 104059 (2016).

[17] M. Fays, Cardiff University thesis, 2017, http://orca.cardiff.ac.uk/id/eprint/110245.

[18] E. E. Flanagan and S. A. Hughes, New J. Phys. **7,** 204 (2005).

[19] P. Welch, IEEE Trans. Audio Electroacoust. **15,** 70 (1967).

[20] M. Was, M. A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, N. Leroy, F. Robinet, and M. Vavoulidis, Classical Quantum Gravity **27,** 015005 (2010).

[21] D. Davis et al., Classical Quantum Gravity **38,** 135014 (2021).

[22] Y. Bengio, in Neural Networks: Tricks of the Trade (Springer, New York, 2013).

[23] P. Barrett, J. Hunter, J. Miller, J.-C. Hsu, and P. Greenfield, in Astronomical Data Analysis Software and Systems XIV, edited by P. Shopbell, M. Britton, and R. Ebert, Astronomical Society of the Pacific Conference Series Vol. 347 (2005), p. 91.

[24] C. R. Harris et al., Nature (London) **585,** 357 (2020).

[25] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (IEEE Computer Society, Los Alamitos, CA, 2016), pp. 779–788.

[26] M. Tan and Q. V. Le, in Proceedings of the 36th International Conference on Machine Learning, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of Machine Learning Research Vol. 97 (PMLR, 2019), pp. 6105–6114.

[27] P. Astone, P. Cerda-Duran, I. Di Palma, M. Drago, F. Muciaccia, C. Palomba, and F. Ricci, Phys. Rev. D **98,** 122002 (2018).

[28] A. Corsi and P. Mészáros, Astrophys. J. **702,** 1171 (2009).

[29] X. Zhu and X. Wu, Artif. Intell. Rev. **22,** 177 (2004).

[30] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, in Proceedings of the 26th Annual International Conference on Machine Learning (Association for Computing Machinery, New York, 2009), ICML '09, pp. 41–48, ISBN 9781605585161.

[31] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in Proceedings of the 25th International Conference on Neural Information Processing Systems, Vol. 1 (Curran Associates Inc., Red Hook, 2012), NIPS'12, pp. 1097–1105.

[32] C. Szegedy, L. Wei, J. Yangqing, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, in 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015), pp. 1–9.

[33] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright et al., Nat. Methods **17,** 261 (2020).

[34] J. F. Kaiser, in Systems Analysis by Digital Computer, edited by F. Kuo and J. Kaiser (John Wiley and Sons, New York, 1966), pp. 218–285.

[35] R. Abbott et al. (The LIGO Scientific, the Virgo, the KAGRA Collaborations), Phys. Rev. D **104,** 102001 (2021).

[36] M. H. P. M. Van Putten, Astrophys. J. **819,** 169 (2016).

[37] E. A. Huerta, P. Kumar, B. Agarwal, D. George, H.-Y. Schive, H. P. Pfeiffer, R. Haas, W. Ren, T. Chu, M. Boyle et al., Phys. Rev. D **95,** 024038 (2017).

[38] M. H. P. M. Van Putten, Astrophys. J. **575,** L71 (2002).

[39] E. Cuoco et al., Mach. Learn. **2,** 011002 (2020).

[40] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning (MIT Press, Cambridge, 2016).

[41] A. Menendez-Vazquez, M. Kolstein, M. Martinez, and L. M. Mir, Phys. Rev. D **103,** 062004 (2021).

[42] M. López, I. Di Palma, M. Drago, P. Cerda-Duran, and F. Ricci, Phys. Rev. D **103,** 063011 (2021).

[43] G. Baltus, J. Janquart, M. Lopez, A. Reza, S. Caudill, and J. R. Cudell, Phys. Rev. D **103,** 102003 (2021).

[44] P. Krastev, Phys. Lett. B **803,** 135330 (2020).

[45] P. G. K. Krastev, K. Gill, V. A. Villar, and E. Berger, Phys. Lett. B **815,** 136161 (2021).

[46] M. Zevin et al., Classical Quantum Gravity **34,** 064003 (2017).

[47] V. Skliris, M. Norman, and P. Sutton, arXiv:2009.14611.

[48] I. J. Goodfellow et al., in Proceedings of the 27th International Conference on Neural Information Processing Systems, Vol. 2 (MIT Press, Cambridge, 2014), NIPS'14, pp. 2672–2680.

[49] J. McGinn, C. Messenger, M. J. Williams, and I. S. Heng, Classical Quantum Gravity **38,** 155005 (2021).

[50] F. Xing, Y. Xie, X. Shi, P. Chen, Z. Zhang, and L. Yang, BMC Bioinf. **20,** 472 (2019).

[51] A. Paszke et al., in Advances in Neural Information Processing Systems 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2019), pp. 8024–8035.

[52] D. Kingma and J. Ba, in 3rd International Conference on Learning Representations, edited by Y. Bengio and Y. LeCun, Conference Track Proceedings (ICLR, San Diego, CA, 2015).

[53] S. Dall'Osso, B. Giacomazzo, R. Perna, and L. Stella, Astrophys. J. **798,** 25 (2015).

[54] B. P. Abbott et al., Classical Quantum Gravity **33,** 134001 (2016).

[55] J. Aasi et al., Classical Quantum Gravity **32,** 115012 (2015).

[56] F. Robinet, N. Arnaud, N. Leroy, A. Lundgren, D. Macleod, and J. McIver, SoftwareX **12,** 100620 (2020).