The Astrophysical Journal, 945:106 (12pp), 2023 March 10 https://doi.org/10.3847/1538-4357/acba0a

A Deep-learning Anomaly-detection Method to Identify Gamma-Ray Bursts in the Ratemeters of the AGILE Anticoincidence System

N. Parmiggiani[1], A. Bulgarelli[1], A. Ursi[2], A. Macaluso[3], A. Di Piano[1,4], V. Fioretti[1], A. Aboudan[1,5], L. Baroncelli[1], A. Addis[1], M. Tavani[2,6], and C. Pittori[7,8]
[1] INAF/OAS Bologna, Via P. Gobetti 93/3, I-40129 Bologna, Italy; nicolo.parmiggiani@inaf.it
[2] INAF/IAPS Roma, via del Fosso del Cavaliere 100, I-00133 Roma, Italy
[3] German Research Center for Artificial Intelligence (DFKI), D-66123 Saarbruecken, Germany
[4] Università degli Studi di Modena e Reggio Emilia, DIEF, Via Pietro Vivarelli 10, I-41125 Modena, Italy
[5] CISAS G. Colombo University of Padova, via Venezia 15, I-35131 Padova, Italy
[6] Università degli Studi di Roma 'Tor Vergata', via della Ricerca Scientifica 1, I-00133 Roma, Italy
[7] INAF/OAR Roma, Via di Frascati 33, I-00078 Monte Porzio Catone (RM), Italy
[8] ASI Space Science Data Center (SSDC), Via del Politecnico snc, I-00133 Roma, Italy
Received 2022 July 10; revised 2023 January 31; accepted 2023 February 6; published 2023 March 13

Abstract

Astro-rivelatore Gamma a Immagini Leggero (AGILE) is a space mission launched in 2007 to study X-ray and gamma-ray astronomy. The AGILE team developed real-time analysis pipelines to detect transient phenomena such as gamma-ray bursts (GRBs) and react to external science alerts received by other facilities. The AGILE anticoincidence system (ACS) comprises five panels surrounding the AGILE detectors to reject background-charged particles. It can also detect hard X-ray photons in the energy range 50–200 keV. The ACS data acquisition produces a time series for each panel. The time series are merged into a single multivariate time series (MTS). We present a new deep-learning model for the detection of GRBs in the ACS data using an anomaly detection technique. The model is implemented with a convolutional neural network autoencoder architecture trained in an unsupervised manner, using a data set of MTSs randomly extracted from the AGILE ACS data. The reconstruction error of the autoencoder is used as the anomaly score to classify the MTS. We calculated the associated *p*-value distribution, using more than 10^7 background-only MTSs, to define the statistical significance of the detections. We evaluate the trained model with a list of GRBs reported by the GRBWeb catalog. The results confirm the model's capabilities to detect GRBs in the ACS data. We will implement this method in the AGILE real-time analysis pipeline.

Unified Astronomy Thesaurus concepts: Time series analysis (1916); Gamma-ray bursts (629); Gamma-ray astronomy (628); Gamma-ray detectors (630); Convolutional neural networks (1938)

1. Introduction

Astro-rivelatore Gamma a Immagini Leggero (AGILE; i.e., Light Imager for Gamma-Ray Astrophysics) is a space mission of the Italian Space Agency (ASI) devoted to high-energy astrophysics, launched in 2007 and still operational (Tavani et al. 2008, 2009). The AGILE payload consists of the Silicon Tracker (ST), the SuperAGILE X-ray detector (SA), the CsI (Tl) Mini-Calorimeter (MCAL), and an AntiCoincidence System (ACS; Perotti et al. 2006). The combination of ST, MCAL, and ACS composes the Gamma-Ray Imaging Detector (GRID). Due to a reaction wheel failure occurred in 2009, the AGILE satellite is continuously spinning around its Sun-pointing axis, with an angular velocity of $\sim$7 minute^{-1}. For this work, we only use data acquired during the spinning mode period.

The ACS is composed of five independent plastic scintillation panels (four lateral and one on top) surrounding the AGILE detectors. The primary role of the ACS is to reject charged background particles. It can also detect hard X-ray photons in the energy range of 50–200 keV. The ACS continuously records each panel count rate in telemetry as ratemeters (RM) data, with 1.024 s resolution. Each ACS panel RM count rate constitutes a time series.

The AGILE team developed real-time analysis (RTA) pipelines (Bulgarelli 2019; Parmiggiani et al. 2021) to detect transient phenomena such as gamma-ray bursts (GRBs) and promptly react to external science alerts sent by other facilities (such as LIGO–Virgo, IceCube, or other space missions). When the pipeline detects a GRB through an RTA pipeline, it sends an automatic notice to the General Coordinates Network[9] (GCN) as well as a notification to the AGILE team. This automated software system implements different algorithms to analyze data and is based on the RTApipe framework (Parmiggiani et al. 2022). It is integrated with the real-time data processing at the AGILE Data Center at ASI/SSDC (Pittori & The AGILE-SSDC Team 2019).

This work aims to develop a new detection method based on deep learning (DL) to identify GRBs inside the ACS data. Although the primary role of the ACS is not the detection of high-energy transients of astrophysical interest, we aim to use the ACS data to improve the overall AGILE capability of identifying GRBs.

The AGILE RTA system implements an algorithm that searches count peaks in the data of the ACS top panel (only one panel) by calculating the signal-to-noise ratio (S/N) of each bin

[9] https://gcn.nasa.gov/

over the background level of the time series near the bin under evaluation. This method is designed to highlight the bins over a predefined S/N threshold in the plot of the ACS data that the AGILE team can visualize from the results of the AGILE RTA system. It does not calculate the statistical significance and has a high level of false positives. For this reason, it does not send automated notifications. In fact, its goal is to put in evidence all possible signals over a threshold to the AGILE team members that manually evaluate the highlighted signals, so even if this method generates false positives, they are discarded by the AGILE Team. We developed a new method based on DL techniques to overcome the limits of this algorithm now implemented inside the AGILE RTA. We cannot compare the results of the two methods because they are developed with different purposes, analyze different data (only one panel for S/N), and generate different outputs (the S/N analysis does not calculate the statistical significance).

The new method based on DL improves the GRBs detection capabilities using the ACS data for several reasons:

1. The DL model uses the data of four AC panels, analyzing them as a multivariate time series.
2. The DL model is trained with the real ACS data, while the S/N method is a generic analysis tool for the identification of signals in a time series. In fact, convolutional neural network (CNN) learns the time-series behavior using a data set containing one year of data, taking all the background rate modulations experienced by the satellite along its orbit and due to the spinning into consideration. This training is task-specific and depends on the data acquired by the instrument.
3. The DL model detects anomalies (GRBs) considering the global patterns of a GRB that differs from the background-only data, while the S/N method implemented in the AGILE RTA is only based on the evaluation of the count rate.
4. The S/N analysis is independently applied to every single step of the time series without considering the historical behavior of the time series (e.g., how many times a certain number of counts occurred). On the contrary, the DL method uses the p-value analysis (Section 4) to consider the occurrence of background data that can be classified as GRB resulting in a false-positive result.
5. The method based on the DL allows us to calculate the statistical significance of the transient event using the p-value analysis and can be used to send automated notifications to the AGILE team or to the scientific community (implementing an automated algorithm to classify the anomalies).

The AGILE researchers have already used DL techniques to detect GRBs, such as the DL model developed to classify the data acquired by the GRID detector (Parmiggiani et al. 2021).

The DL models (Lecun et al. 2015) are part of the machine-learning (ML) model, which uses automated training algorithms to learn how to predict the correct output without human intervention. These models are applied to several problems (e.g., classification and regression). The training process uses a training data set prepared for this purpose. Usually, this data set is a subset of the possible inputs that the model will face during the production phase. The ML techniques are limited because they cannot directly analyze the raw data without an expensive feature-extraction step carried on by data experts. The ML model uses the extracted features as inputs. The DL methods, on the contrary, do not require this feature-extraction process because they extract the features during the training procedure. The deep neural networks used to implement the DL model contain several layers to extract features from the input data at different levels of abstraction during the training. The number of layers can vary and must be optimized on the specific data set. In recent years, the DL popularity has grown thanks to the improvement of computational hardware (e.g., graphical processing unit) and the availability of large amounts of data. These two factors, with the availability of high-level frameworks to implement DL models, lead to the extensive use of DL analysis techniques.

The rest of the manuscript is organized as follows. Section 2 describes the method used to identify the background signal of the ACS panels. In Section 3 we present the DL model chosen for this work and the procedures that we use to train and evaluate it. Section 4 describes the method used to calculate the p-value distribution and the statistical significance of the detections. In Section 5 we report the results obtained with the trained model by analyzing the AGILE data archive. Finally, Section 6 presents the conclusion and possible future applications of the method described in this manuscript.

2. Data Evaluation

We need to evaluate the ACS data acquired during the AGILE mission lifetime by comparing different years and verifying that the background levels are similar. Before the analysis, we applied a list of filters and preprocessing procedures to the ACS RM. This work considers AGILE data acquired from 2010 January to 2020 December.

We extracted the ACS RM data from the archive and imported them into an InfluxDB[10] database as time series. We decided to use InfluxDB because it is an open-source database developed to manage time series, and it has a graphical user interface that visualizes the stored data through several diagrams. In addition, it is possible to query this database from custom software applications written with Python. We opted to import the data into a database as it is faster to make queries to this database than open a long file list.

In this work, we discard the time series acquired by the fourth panel due to its orientation toward the Sun and thus its sensitivity to solar flares, which interferes with GRBs detection. We will analyze the data acquired by this panel in a dedicated work to design a detection method that can be implemented in the automated pipeline developed by the AGILE team for solar-flare monitoring. The time series of the four ACS panels analyzed in this work constitutes a single multivariate time series (MTS) due to time alignment. We expect that a transient event is detected by more than one panel simultaneously.

Figure 1 shows the raw data acquired by an ACS panel during a short period (less than one day) in 2020. All the AGILE detectors are usually put into an idle mode when the satellite passes into the South Atlantic Anomaly (SAA), due to the high concentration of trapped charged particles present in this geographic region, which affects the gamma-ray background and prevents reliable data acquisitions. However, in some exceptional configurations operated throughout its

[10] https://www.influxdata.com/products/influxdb-overview/

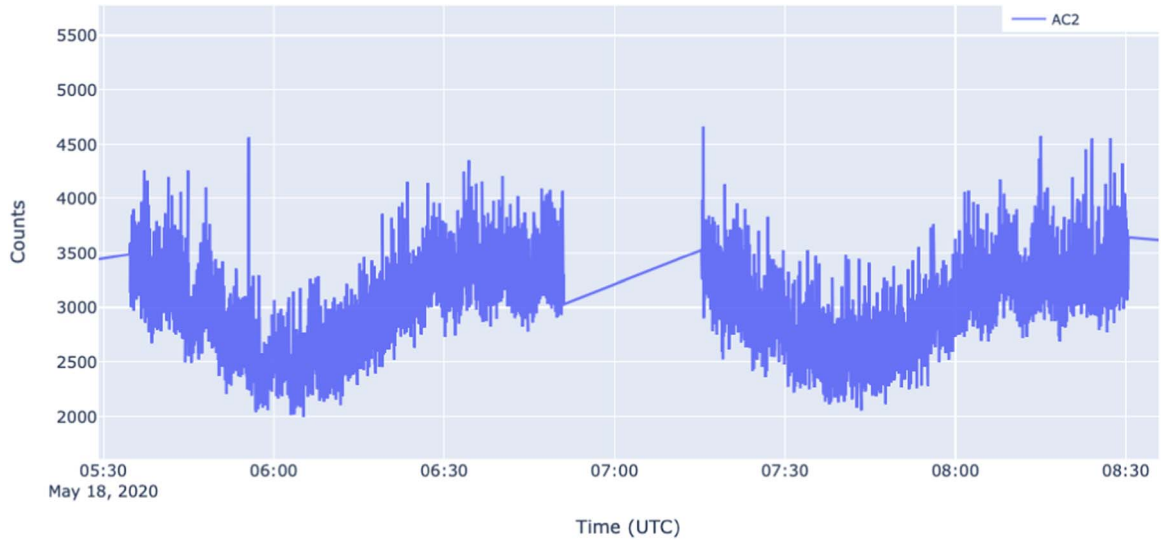

Figure 1. Raw data extracted from the ACS data, corresponding to a 3 hr time interval, acquired on 2020 May 18th with a 1.024 s time resolution. The data interruptions correspond to the satellite passages into the SAA, where all detectors are put in idle mode.

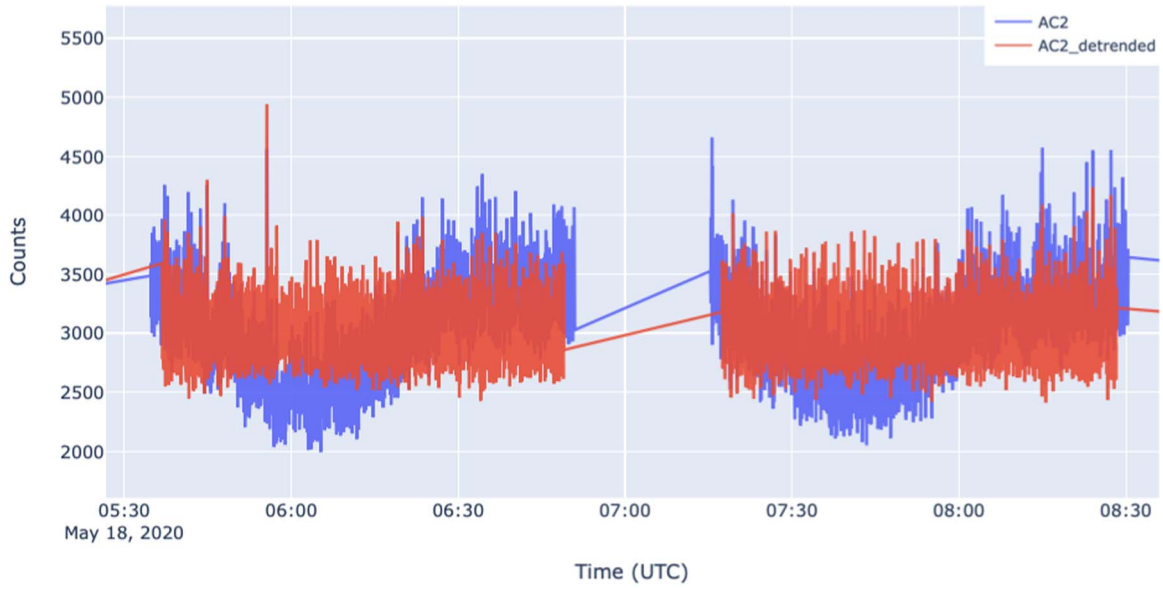

Figure 2. Comparison between detrended (purple) and raw data (blue) extracted from the ACS data for the same time interval reported in Figure 1.

mission lifetime, AGILE has sporadically acquired data even during passages into the SAA, revealing huge increases in the count rates of all detectors. As a detection algorithm could wrongly classify such increases as potential cosmic transients, we removed all the time windows related to SAA passages from the ACS data with a margin of five minutes before and after the SAA.

The X- and gamma-ray backgrounds collected by all the AGILE detectors are affected by periodic oscillations and modulations. The spacecraft spinning over its Sun-pointing axis generates most of these variations on timescales of $\sim$7 minutes (i.e., $\nu \sim 2.4 \cdot 10^{-3}$ Hz), while the different concentrations of high-energy particles along the orbital path cause variations with timescales of $\sim$96 minutes (i.e., $\nu \sim 1.7 \cdot 10^{-4}$ Hz). We implemented a detrending algorithm based on fast Fourier transforms (FFTs), which removes the effects of these background oscillations, independently from each ACS panel data. Figure 2 shows examples of the results obtained by applying the detrending algorithm to data acquired in 2020. This algorithm analyzes all time windows containing a time series with no interruption in the RM for at least 30 minutes. The algorithm skips shorter time series because it cannot efficiently detrend them. By applying this policy, we decided to drop some time windows to increase the data quality. We removed 120 s from the head and tail of each MTS to discard the MTS edges. We use the detrended data for the analyses described in the next sections.

Figure 2 shows that the detrending algorithm can remove the main trends from the raw signal. The intervals lacking data are due to the SAA passages.

The analysis method described in this work aims to detect GRBs inside the ACS RM when the AGILE RTA receives external science alerts. We defined a research interval of 140 bins of 1.024 s, including the burst trigger time, by analyzing the third Swift Burst Alert Telescope GRB catalog (Lien et al. 2016). The T_{50} and T_{90} parameters correspond to the time

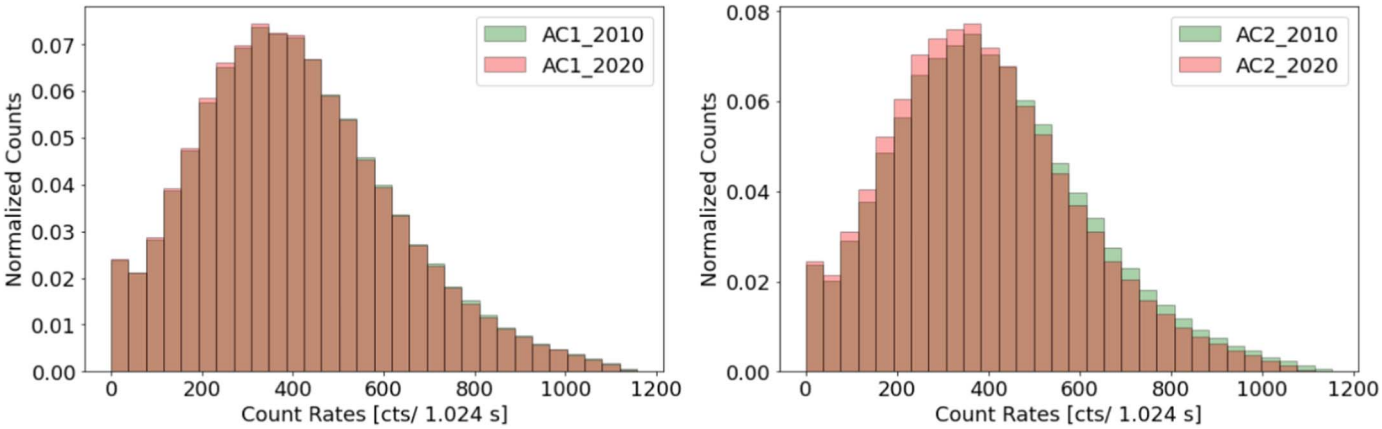

Figure 3. Normalized distributions of ACS RM background count rates acquired during two different years (2020 and 2010) by the ACS panel 1 and 2. It can be seen that the two distributions exhibit very similar profiles, pointing out that the AGILE ACS background rate remained quite stable throughout the years and that our analysis method can be applied to different periods of the mission lifetime.

intervals over which the central 50% and 90% of the GRB cumulative counts above the background are detected, respectively (Kouveliotou et al. 1993). We found an optimal duration using the 95 percentile of the T_{50} value distribution, obtaining 110 s. As the ACS might not be able to precisely reveal the initial and final stages of the detected bursts, we adopted the T_{50} as a reference parameter, as it mainly focuses on the most central part of the event, to characterize the GRB time profile. We added 30 s to the interval to include background before and after the supposed GRB trigger time notified by the external science alert.

Following this procedure, we obtained a time-series data set with a length of 140 bins each. The next step was to subtract the minimum value of each time series from their counts to have the minimum value equal to 0 in all time series.

We compared the distribution of the ACS data over the years to verify that the background levels were similar during the years. Thus we can analyze the data acquired during different periods using the same method. Figure 3 shows the comparison between the data, normalized following the procedure described in Section 3.1, acquired during 2020 and 2010. We should note that the distributions have a similar behavior, and this is true also for other years.

3. Deep-learning Model

We developed a DL model based on an anomaly detection technique to identify MTSs significantly deviating from the background-only data used for the model training. Among the DL architectures designed to execute this type of analysis (Pang et al. 2021), we decided to implement our model with a CNN autoencoder (Goodfellow et al. 2016) using several 1D CNN layers, designed to work with time series having one or more channels. CNNs are well-known in several fields for the successes obtained with image processing (e.g., object detection and object segmentation), but they also achieved important results in the time-series domain (Zhao et al. 2017; Munir et al. 2019).

The autoencoders are neural networks designed to encode the input data in a representation with reduced dimension and then decode the compressed information to the original data, minimizing the reconstruction error (the differences between the original input data and the reconstructed ones).

This kind of neural network is usually used for anomaly detection because, when a trained autoencoder receives as input an object different from the training data set, it cannot efficiently reconstruct it, resulting in a significant reconstruction error. We evaluated 36 different model architectures by varying the hyperparameters (parameters that define the structure of the network) to find the best model configuration for our use case. The best model has two 1D CNN layers for the encoding (with 250 and 500 filters, respectively) and two 1D deconvolutional layers for the decoding (with 500 and 250 filters, respectively). The last layer is a 1D deconvolutional layer with four filters used to reconstruct the four time series with the same shape of the input. The 1D CNN layers use the ReLu activation function followed by a dropout layer with a 20% value. The kernel size used for the convolution is 5. We implemented the network using two open-source frameworks: Keras[11] running on top of Tensorflow.[12]

The main reasons that we considered when we decided to use this DL architecture are:

1. The autoencoder architecture is designed to compress the size of the input, trying to maintain the important information required to correctly reconstruct the input. This behavior forces the model to learn the temporal patterns, which is useful for classifying the time series (Thill et al. 2021).
2. Using this DL architecture, we can analyze the MTS containing the data of all panels. With this kind of analysis, it is possible to improve the detection capability of GRBs when the burst is present in the time series of more than one ACS panel (Zheng et al. 2016).
3. The unsupervised DL, which does not require a labeled data set, fits well with our problem as we have a large amount of unlabeled data that we cannot simulate nor efficiently label. In addition, unsupervised training is widely used in anomaly detection problems (Munir et al. 2019) where the data sets are highly imbalanced because the number of anomalies is much lower than normal data.
4. The autoencoder applied to MTS also implements recurrent neural networks (RNN) such as the long short-term memory (LSTM). However, these networks

[11] https://keras.io
[12] https://www.tensorflow.org

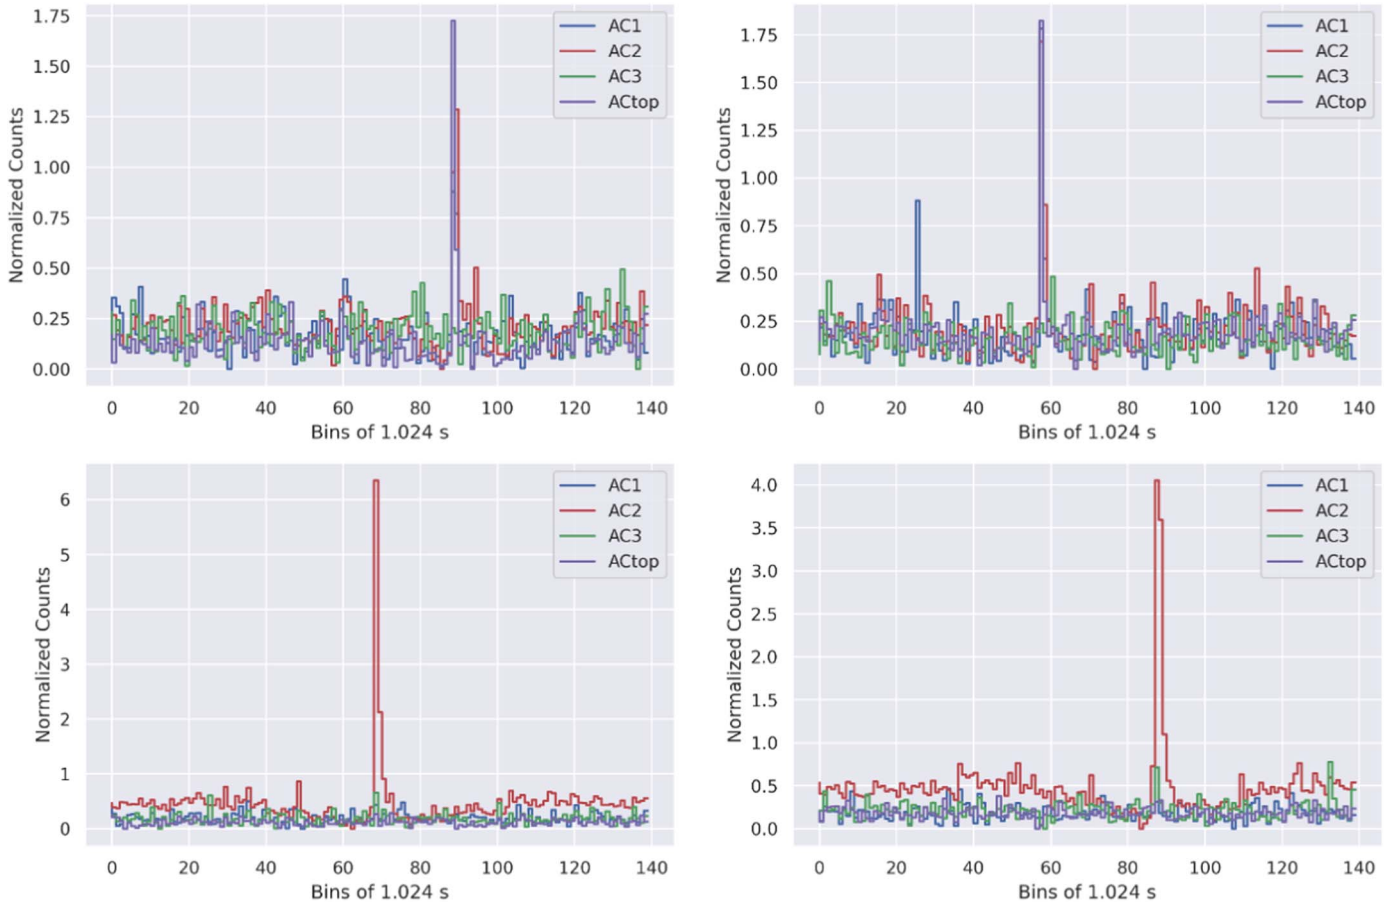

Figure 4. Examples of four different light curves that the filter described in Section 3.1 removed from the background-only data set used for the model training. These light curves represents anomalies that cannot be considered background signals.

require more computing time to be trained and are used to extract long-term features (Zhang et al. 2021). In our case, the GRB features are located in a short time window and are independent of what happened before. We verified that we obtained a higher detection performance with less computational effort with CNN.

3.1. Data Set Creation

We prepared a data set of 4572 MTSs from the AGILE ACS RM data acquired during 2020 by sampling time windows of 140 bins of 1.024 s. To ensure that the data set contains background-only MTSs, we excluded all time windows where the AGILE detectors were in idle mode, where the satellite was flying into the SAA, and where known GRBs (present in the GRBWeb catalog described in Section 5) were taking place. However, the filtered data can still contain anomalies not part of the background that can be associated with transient events not already detected from other instruments and not present in the catalog. We have to remove these outliers inside the background-only data set. We excluded the time series with values greater than the 99.993 percentile of the distribution of ACS time-series values considering each ACS panel independently. Figure 4 shows four examples of outliers that we excluded from the data set.

We divided the 4572 MTSs in two data sets of different sizes for the training (3765 MTSs) and test (807 MTSs) phases. The MTSs of all data sets are detrended and reduced with the minimum value to zero following the procedure described in Section 2. Then, the MTSs values are scaled between 0 and 1 using the minimum and maximum values of the entire training data set. These values are stored to scale the MTSs used as input of the model during the following phases of the work. When the model is used to classify new data, the data are normalized using the same values calculated with the training data set to ensure that the input follows the same values distribution.

The procedures described in Section 2 and the scaling are part of the overall normalization procedure applied to the MTSs before using them for the analyses. Figure 5 shows the MTSs of 140 bins taken from the training data set of two ACS panels before and after the normalization procedure.

3.2. Model Training

We trained the autoencoder with an unsupervised procedure to avoid data labeling. For this reason, the model cannot explicitly classify the anomalies detected (e.g., GRBs or instrument anomalies), and a domain expert or an automated algorithm is necessary to validate the results.

We executed the model's training with a batch size of 100 MTSs using the training data set containing 3765 MTSs. We used a validation split factor of 0.1, so 10% of the training data set is not used for the training but for the validation procedure. The algorithm automatically stopped the training after 45 epochs because the results had not improved for five epochs. This means that the loss is not the same but can increase by a

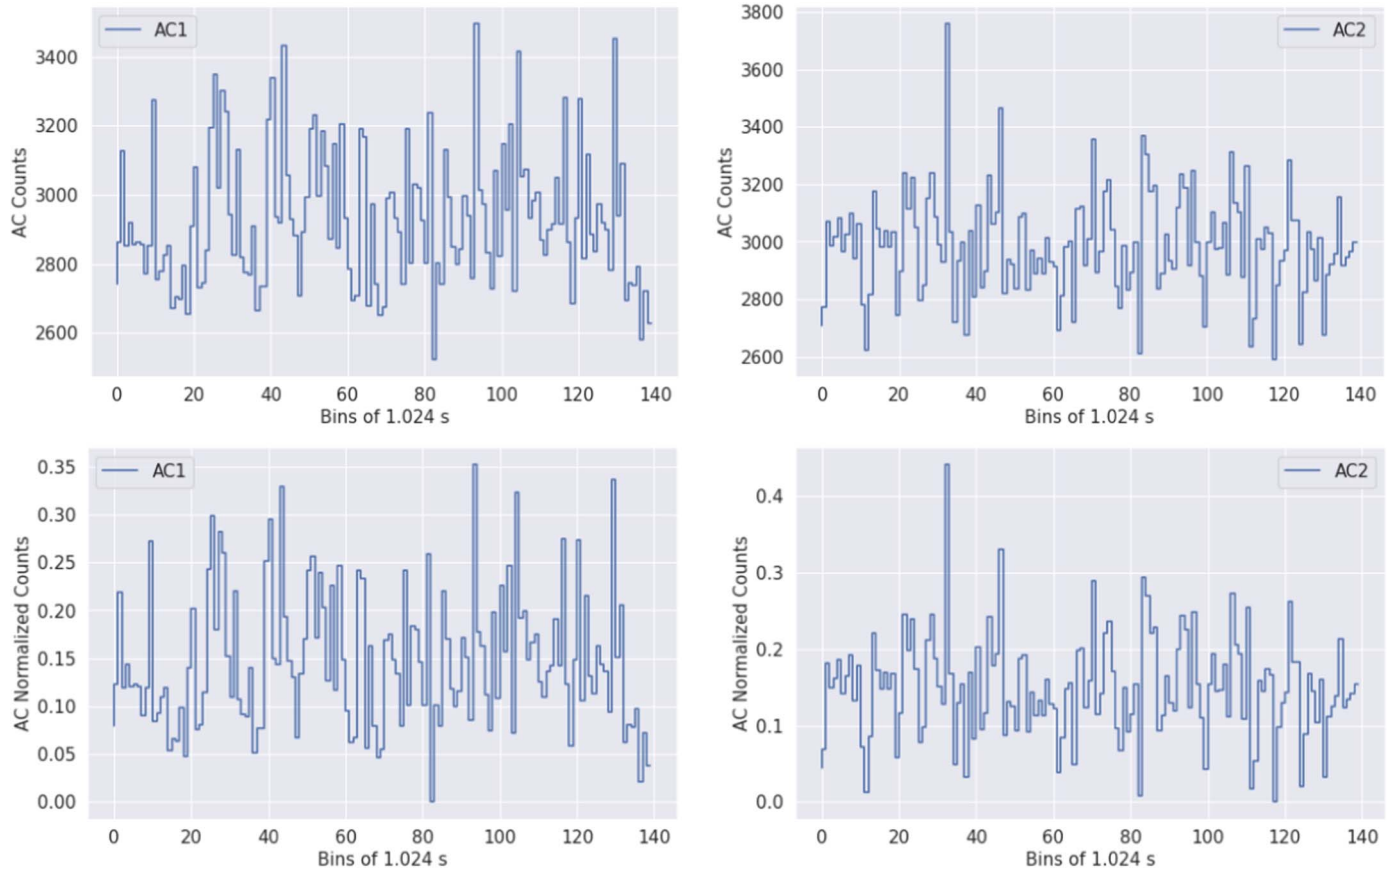

Figure 5. Example of a MTS extracted from the training data set showing the time series of the AC1 (left) and AC2 (right) before (upper plots) and after (bottom plots) the normalization procedure.

small amount that does not produce overfitting, as shown in Figure 6 where the validation loss has the same trend that the training loss. We configured this behavior to avoid the model overfitting on the training data set and to allow the model to search for additional optimum results instead of stopping it with the first optimum found. The optimization algorithm is Adam (Kingma & Ba 2014), configured with a learning rate of 0.001, and the reconstruction error is calculated as the absolute value of the difference between the original and reconstructed data.

3.3. Anomaly Score

To use the DL model for the anomaly detection task, the method presented here calculates an anomaly score of each input analyzed with the model. This anomaly score correlates with the reconstruction error obtained with the autoencoder network architecture.

The reconstruction error is calculated as the absolute value of the difference between the original and reconstructed data for each ACS panel and each MTS. The reconstruction errors in the first and last 5 s of the MTS are set to zero to avoid artificial anomalies at the edges. Figure 7 represents an example of this procedure. The array of error values is filtered to select elements over the 90 percentile of the values' distribution. This procedure is repeated on the reconstruction errors of each ACS panel. Finally, the error values of the four ACS panels are summed, bin by bin, to obtain the cumulative reconstruction error of any given MTS. The anomaly score is calculated by summing all values (already summed for the four panels) that

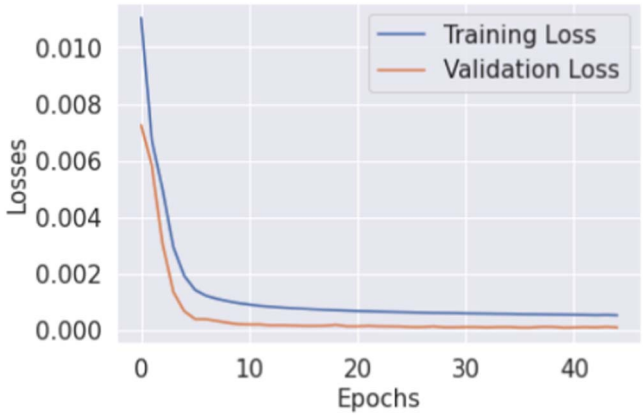

Figure 6. Results of the training process. The orange line shows the loss of the model when it evaluates the validation data, while the blue line represents the loss of the model when it evaluates the training data. When the two losses do not decrease for five epochs, the training algorithm is automatically stopped to prevent overfitting the model on the training data.

exceed the 90 percentile of the values' distribution. This procedure filters the values to keep the largest summed reconstruction error values, where a candidate GRB should be detected. A display of the result of this procedure is shown in Figure 8 where the bottom plot indicates the summed reconstruction error of the four ACS panels (green) and the filtered values (violet).

We use the anomaly score to classify MTSs and detect anomalies. If the anomaly score is higher than predefined thresholds (depending on the statistical significance level that

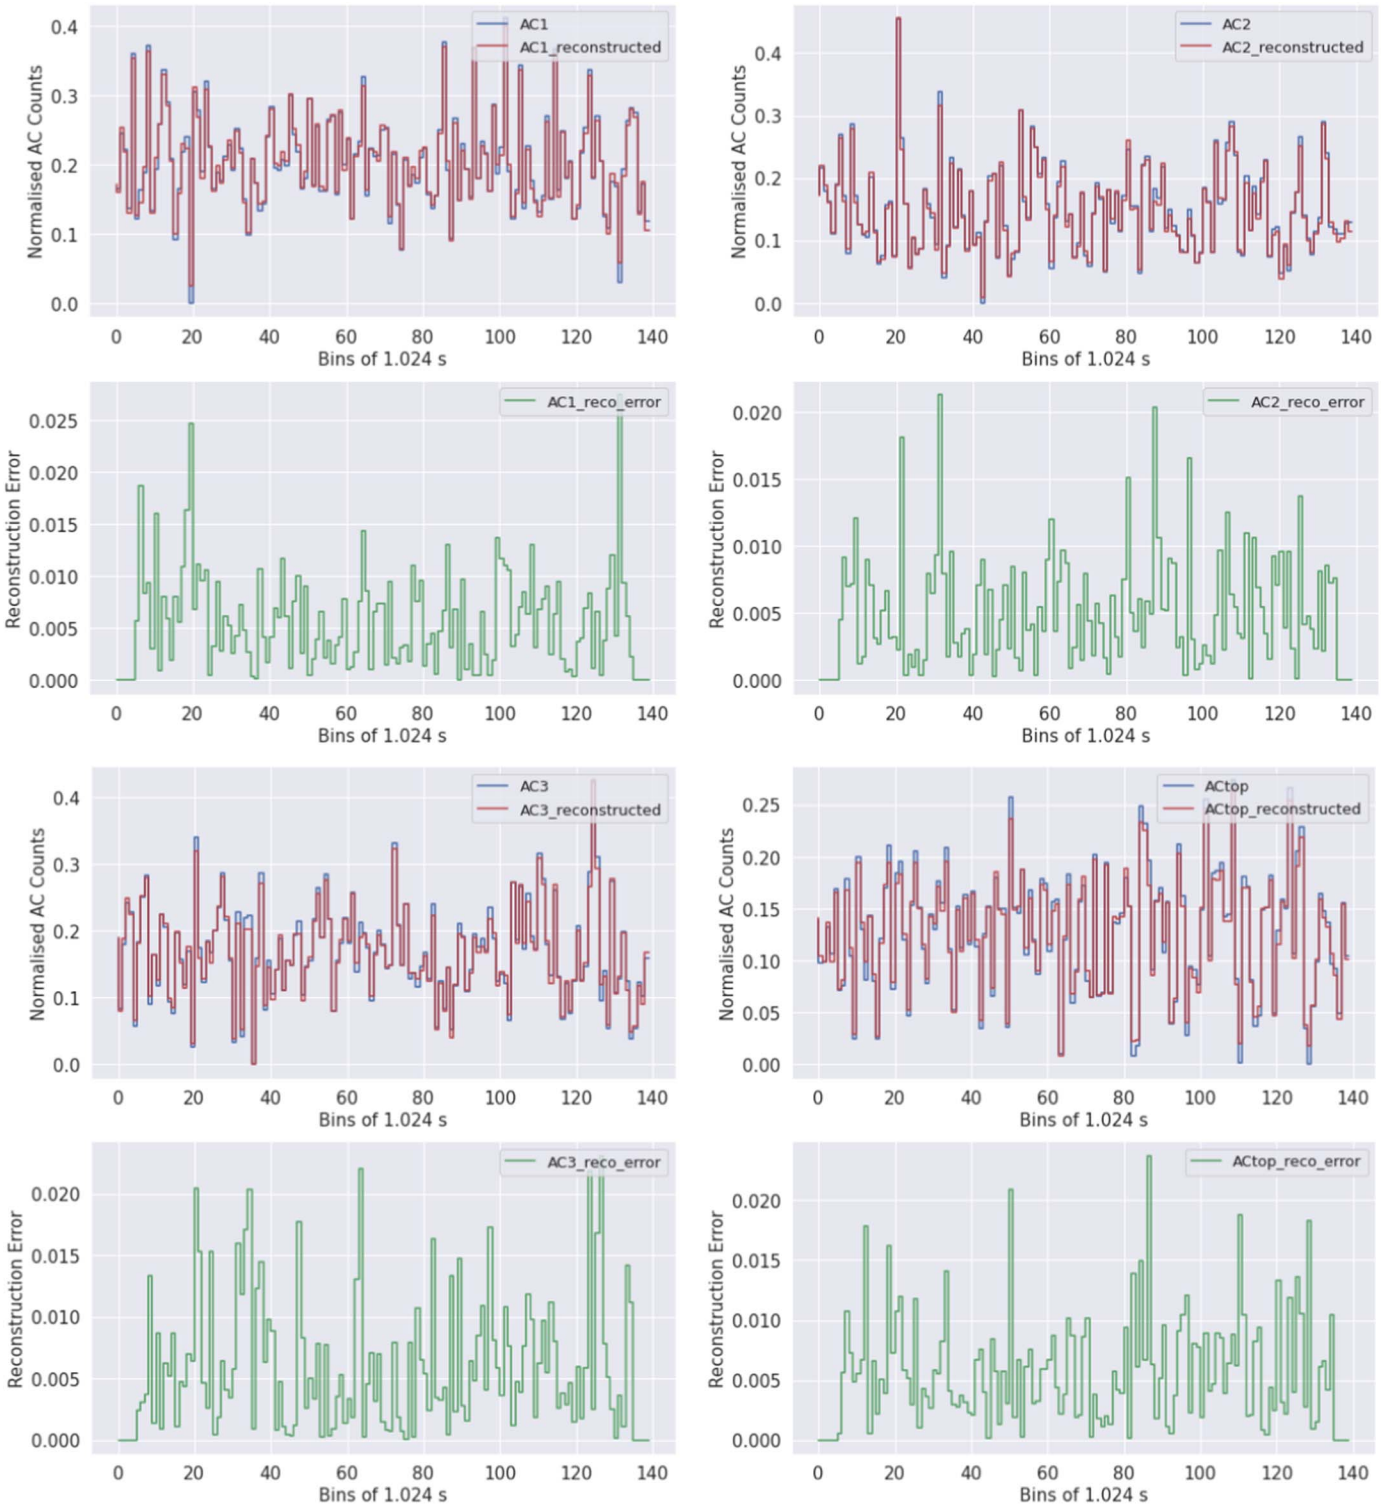

Figure 7. Background-only MTSs for a research interval of 140 bins of 1.024 s, acquired by all the four ACS panels extracted from the test data set. Each panel also shows the comparison between original data (blue) and reconstructed data (red), and the difference between them (green).

we want to reach), the time window is flagged as an anomaly (i.e., a possible GRB detection). The threshold are calculated using the *p*-value method described in the following section.

4. CNN *p*-value Evaluation

We developed this neural network to obtain a new method for GRB detection and implement it as part of the AGILE automated RTA system that follows up external science alerts received from other facilities through the GCN network. When the model detects a GRB with a certain statistical significance, the AGILE team can communicate the results to the community. To obtain the statistical significance of a detection we need to perform a statistical characterization of the background-only MTSs and define new thresholds on the anomaly score.

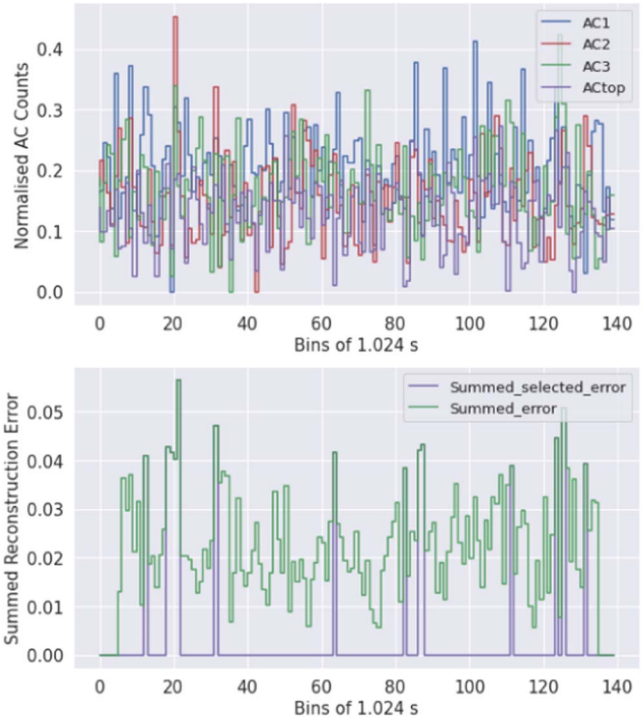

Figure 8. Example of background-only MTSs extracted from the test data set and plotted in the same diagram. The bottom figure represents the summed differences of all the four panels (green) and the values selected (greater than the 90 percentile of the values distribution) to calculate the anomaly score (violet).

We evaluated with the trained model a data set of background-only MTS to determine the p-value distribution of the results, similarly to Parmiggiani et al. (2021). With the p-value distribution, we obtain the thresholds that we can apply to the anomaly scores to reject the null hypothesis and classify the ACS time series analyzed by the model as GRBs with certain significance levels. In this work, we exclude detection with a statistical significance lower than 3σ.

We defined the distribution Φ of the anomaly score (AS) values as resulting from the analysis of the background-only data set. The probability that the result of a trial in an empty field has $\mathrm{AS} \geqslant h$ (that is the complement of the cumulative distribution function) is:

$$P(\mathrm{AS} \geqslant h) = \int_{h}^{+\infty} \Phi(x)dx, \quad (1)$$

which is also called p-value $p = P(\mathrm{AS} \geqslant h)$ and defines the probability of obtaining that AS value or greater when the null hypothesis is true.

The p-value analysis requires more than 10^7 time series to reach a significance of 5σ. Computing resources limit this value. We cannot obtain this number of samples from the AGILE acquired data, and the development of simulation software would require a detailed analysis of the ACS ratemeters during the spinning and orbital phases, but this is out of the scope of this work. The lack of sufficient training data is a common issue when using DL models. We decided to solve this issue by applying the data augmentation method to increase the data set volume with artificial data with a resampling technique. We extracted a data set of 74027 MTSs from the time series of each ACS panel and generated more than 1.5×10^7 synthetic time series from each of these MTS, randomly combining values and avoiding duplications. We avoided merging segments of different MTSs because the resulting synthetic MTS can have different background levels, generating artificial anomalies. Using this method, we generated more than 1.5×10^7 MTSs that constitute our background-only data set. The data set used for this data generation is extracted from data acquired by the ACS during 2019. We used a different time window to extract the training data set and this data set for the p-value analysis. We excluded the known GRBs and the outliers from the starting data set, following the procedure described in Section 3.1, before the data augmentation to avoid the replication of anomalies that are not part of the background data.

Figure 9 shows the normalized distribution of the anomaly scores and the associated p-value distribution with a red dotted line that represents the threshold of 3σ on the AS value that is the minimum value required to classify a time series as an anomaly.

5. Agile Data Archive Analysis

Once we defined the thresholds at different sigma levels (Table 1), we evaluated the trained model using a list of 1586 GRBs that we obtained by applying filters to the GRB catalog reported by the GRBWeb[13] platform by P. Coppin, a catalog of GRBs that combines data from different sources and several detectors. We selected the GRBs in the time windows used for this work and removed those with trigger time falling within the AGILE passages into the SAA or idle periods. We extracted the MTSs of 140 bins of 1.024 s from the AGILE ACS RM, starting each MTS 25 s before the external trigger time reported by the GRBWeb platform. Then, we analyzed these MTSs with the trained model, detecting 95 anomalies with $\sigma \geqslant 3$.

5.1. Performance Analysis

We manually checked these anomalies to validate the results, and we removed from the list two detections because they are due to data anomalies not associated with GRBs and six anomalies because we do not see a clear GRB in the data. In addition, we decided to remove from the validated results 14 anomalies that consists of a single peak in only one of the four time series. These anomalies that consist of a single peak can be due to a charged particle or a short GRB that for some reason is not seen by other ACS panels. However we prefer to be conservative and discard these kind of anomalies. For a data expert (or an automated algorithm) it is trivial to discard this kind of anomalies during the validation of the CNN results. Finally, we discarded another GRB because GRB100724A and GRB100724B are both reported by the GRBWeb platform with just 14 s of time distance. Our model detects both GRBs in the ACS RM but cannot distinguish between them as it is not intended for localization and the time series have the same values. We can discard GRB100724A as the fluences reported for the two transient events indicate that, with a high probability, the ACS system detected GRB100724B. At the end of the validation process, we have a total number of detected bursts equal to 72 (Table 2). Figure 10 shows an example of a GRB detected during this analysis.

[13] https://user-web.icecube.wisc.edu/~grbweb_public/index.html

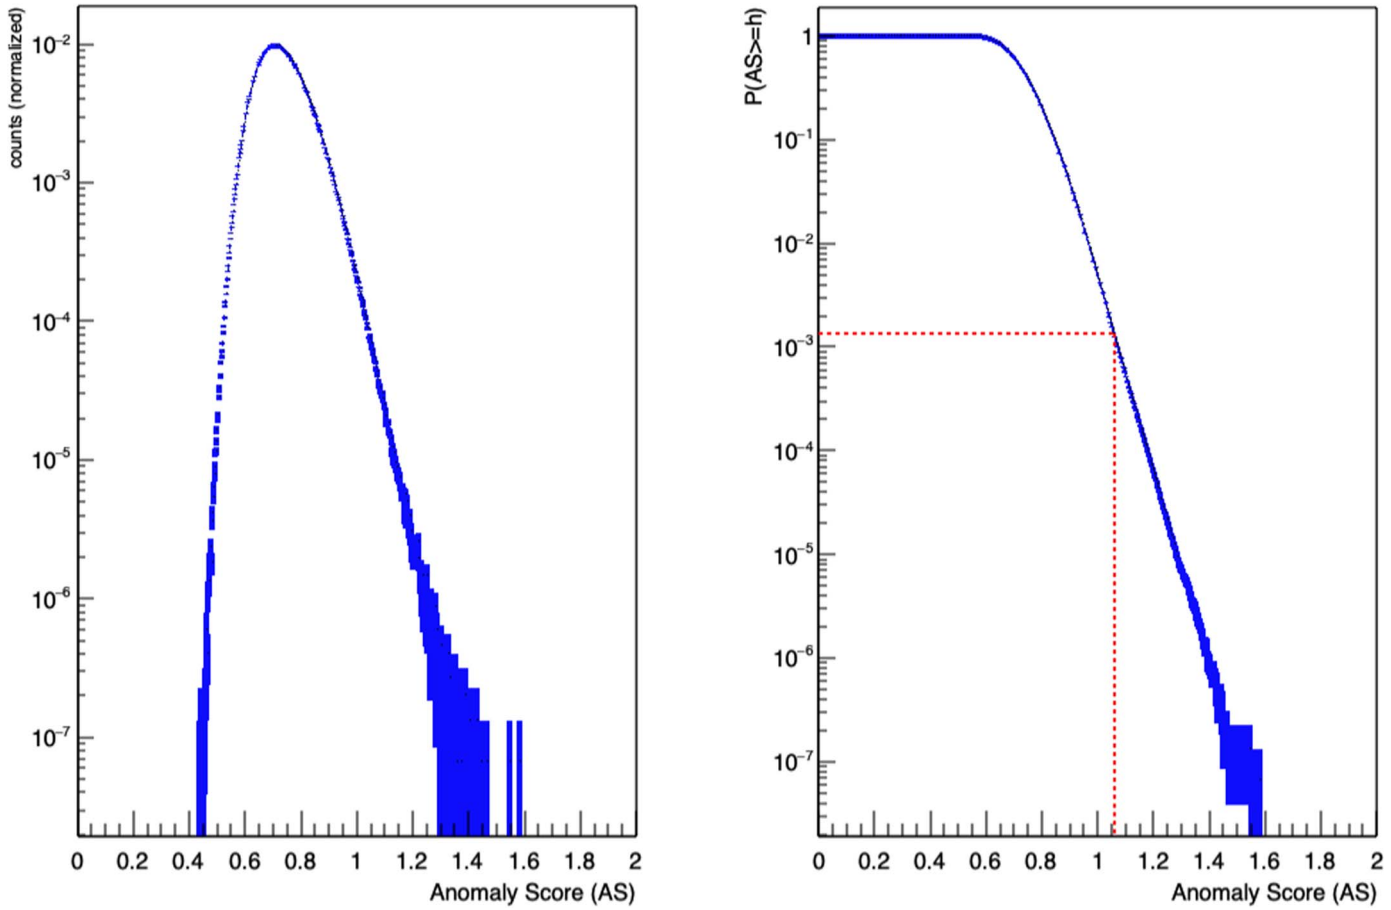

Figure 9. Distribution of the anomaly scores obtained on the background-only data set generated for the *p*-value analysis (left panel). *P*-value distribution used to determine the thresholds on the anomaly scores corresponding to different sigma levels (right panel). The red dotted line represents the threshold for 3σ; time series obtained from the DL model with AS values greater than this threshold are classified as anomalies.

Table 1
Relation between σ and the Threshold on Anomaly Scores

Nσ	*p*-value	Threshold
3	1.35×10^{-3}	1.06
3.5	2.32×10^{-4}	1.14
4	3.17×10^{-5}	1.23
4.5	3.40×10^{-6}	1.34
5	2.86×10^{-7}	1.45

Note. Different detection thresholds for the anomaly score corresponding to predefined statistical significance levels.

The AGILE ACS cannot detect all GRBs reported in the GRBWeb platform for several reasons. The ACS is developed for a different purpose than the detection of transient events and has a coarse time resolution that is not designed for GRB detection. In addition, the ACS data have several interruptions due to the SAA passages and time windows with the instruments in an idle state. Despite these factors, the results obtained prove the capability of this DL model to detect GRBs inside the ACS data.

In future works, we plan to develop a classifier able to separate GRB anomaly detections from other sources (e.g., instrumental anomalies during the acquisition) in order to improve the AGILE pipeline detection performance and avoid human intervention.

5.2. Comparison with the AGILE-MCAL Detector

We compared the GRBs detected with this method with the second AGILE-MCAL GRBs catalog,[14] described in Ursi et al. (2022).

The ACS and the MCAL detector operate in different energy ranges, with different sensitivities and time capabilities, and their results cannot be directly compared. However, for each detected event, the ACS data acquired in the 50–200 keV energy range can provide further information that can be used to integrate the MCAL data in the 0.4–100 MeV energy range. MCAL works as a triggered detector, issuing limited data acquisitions whenever a transient fulfills certain logic conditions: at the moment, the MCAL onboard logic is mainly configured to reveal short-duration, hard-spectrum GRBs (Ursi et al. 2019), often preventing the full detection of longer-duration, softer-spectrum bursts (Ursi et al. 2022). In this perspective, the ACS is more sensitive to long GRBs, as it does not operate on a trigger mode and acquires data continuously and in the X-ray energy range. The AC data can be therefore used to investigate the overall temporal profile of the detected bursts, integrate the limited MCAL triggered acquisitions, and perform preliminary multifrequency characterizations of the events. The synergy of MCAL and ACS helps provide a more detailed picture of each burst, increasing the overall information about AGILE-detected GRBs.

[14] https://www.ssdc.asi.it/mcal2grbcat/

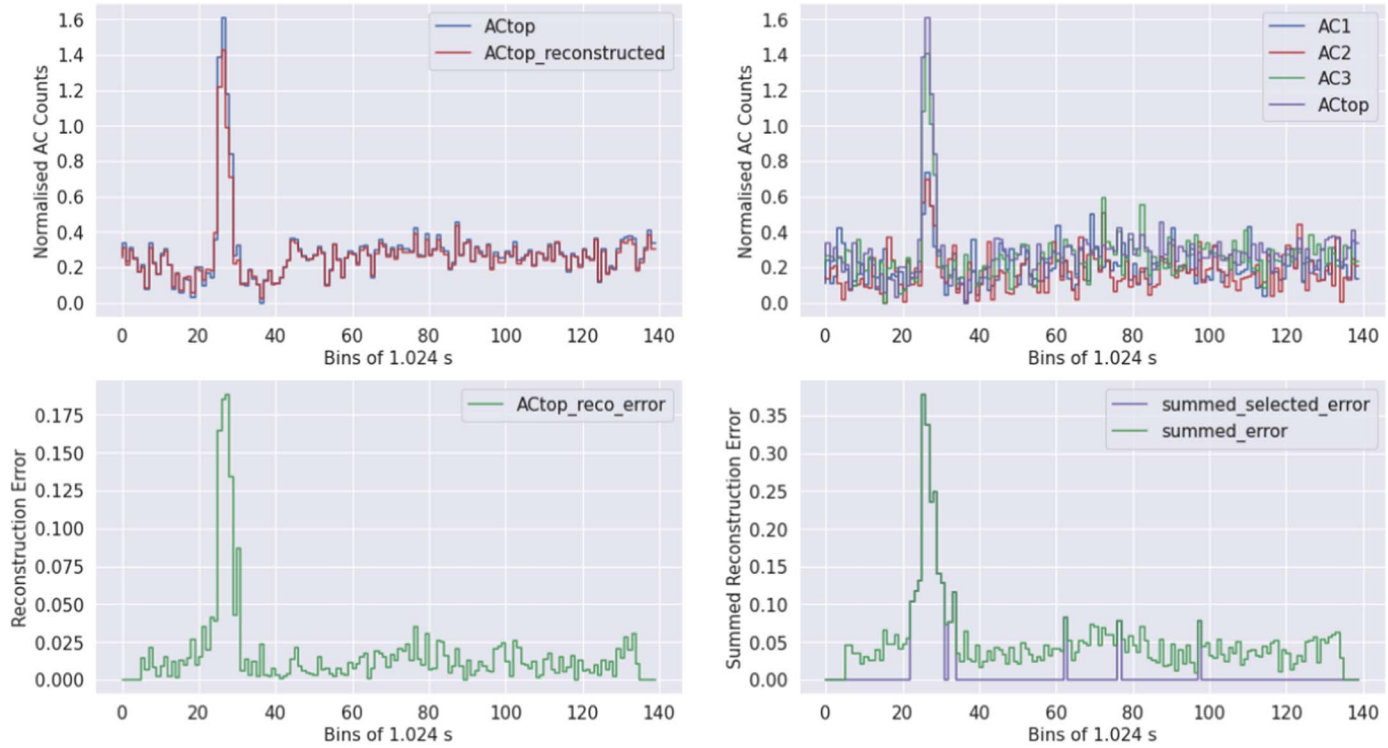

Figure 10. Example of a GRB (GRB100719D) detected in the ACS data. Plot (a) represents one of the four ACS panels (AC top). The original signal is blue, and in red is the reconstructed one. Plot (b) shows the GRB signals present in the original signals of all the ACS panels. The bottom plots represent the difference between the original and the reconstructed signals in green, in the single AC top panel (c), and in all ACS panels (d). The values used to calculate the anomaly score are highlighted in violet.

Table 2
GRB Detected with $\sigma \geqslant 3$

GRB	$N\sigma$	GRB	$N\sigma$	GRB	$N\sigma$	GRB	$N\sigma$
GRB100511A	3.5	GRB130606B	5	GRB160113A	4	GRB180914B	5
GRB100719D	5	GRB130821A	3.5	GRB160131A	5	GRB181121A	5
GRB100724B	5	GRB131028A	5	GRB160530B	5	GRB181227A	5
GRB101014A	5	GRB131108A	4	GRB161010A	4	GRB190103A	5
GRB101023A	5	GRB140416A	5	GRB161218B	3.5	GRB190114C	5
GRB101123A	5	GRB140508A	5	GRB170115B	5	GRB190329A	5
GRB110625A	5	GRB140509B	5	GRB170127C	3.5	GRB190501A	5
GRB110715A	4	GRB140821A	4.5	GRB170311B	5	GRB190611B	4
GRB110825A	3.5	GRB141022B	5	GRB170522B	5	GRB190706C	4
GRB111211A	4.5	GRB141028A	4.5	GRB170607B	5	GRB190727B	5
GRB120129A	5	GRB141104A	4.5	GRB170626B	4	GRB190731A	5
GRB120426A	5	GRB141207A	3	GRB170808B	5	GRB190928A	5
GRB120707A	4.5	GRB150330A	4	GRB171010A	5	GRB191221B	5
GRB120711A	5	GRB150403A	3	GRB171011B	5	GRB200131A	3
GRB120911B	5	GRB150424A	5	GRB171119A	3.5	GRB200829A	5
GRB121118B	5	GRB150510A	5	GRB171227A	5	GRB200903E	5
GRB130320B	5	GRB150523A	3	GRB180720B	5	GRB200919C	5
GRB130408A	3	GRB160106A	5	GRB180806A	3	GRB201020B	5

Note. List of 73 GRBs detected with a significance $\geqslant 3\sigma$ by using our DL model to analyze a list of the 1586 MTSs generated from ACS data, starting from the GRBWeb catalog of external GRBs.

An example of this synergy is shown in Figure 11 that represents the GRB180914B with light curves from MCAL and ACS data. The MCAL bottom panel light curve (black) has a better time resolution (with different timescales) to describe the GRBs in detail. On the other hand, the MCAL light curve is fragmented, while the ACS top-panel light curve (green) provides a complete behavior of the GRBs in a longer timescale. In addition, the MCAL instrument can be IDLE under certain payload configurations while the ACS ratemeters are still acquired. For this reason, using the ACS data to detect GRBs increases the probability of having available data that cover the GRB time window.

MCAL detected 57 GRBs out of the 72 detected by the ACS using our DL method. The remaining 15 GRBs are not detected

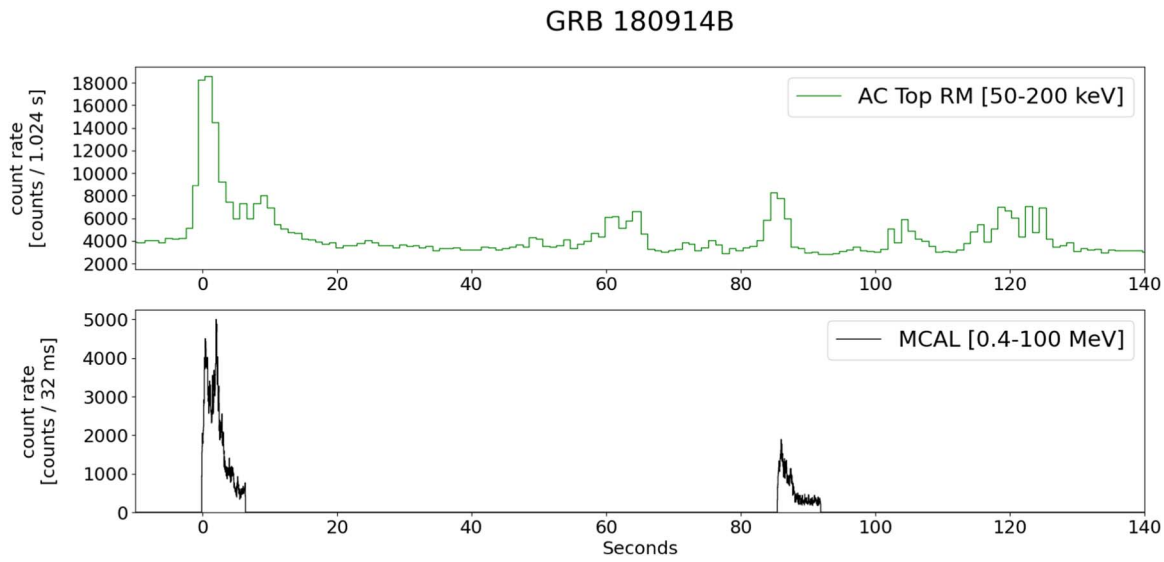

Figure 11. Comparison between the light curve obtained with the MCAL detector and the ACS top panel of the GRB180914B. The ACS top-panel light curve (green) shows complete GRB behavior in a time window of 140 s, while the MCAL light curve (black) shows more details with a better time resolution, but the light curve is fragmented.

by MCAL but only by the ACS data using this DL method. These GRBs are: GRB120426A, GRB120711A, GRB130320B, GRB140509B, GRB141022B, GRB150330A, GRB160113A, GRB161010A, GRB161218B, GRB170626B, GRB170808B, GRB171010A, GRB180720B, GRB181227A, and GRB201020B. This result confirms that the DL method presented in this work improves the detection capabilities of the AGILE real-time analysis pipeline to GRBs, providing a useful tool for automatically analyzing data in the hard X-ray energy range acquired by the ACS. This method can be used during the follow-up of external science alerts and provides additional information to the results obtained by analyzing the data acquired by other AGILE detectors, such as the MCAL.

6. Conclusions

We developed a new DL model based on the CNN autoencoder architecture to detect GRBs inside the AGILE ACS data. This model is trained with a data set of background-only MTSs randomly extracted from the ACS data. We calculated the p-value distributions using more than 1.5×10^7 background-only MTSs generated with a data augmentation technique starting from a data set of real MTSs. This p-value distribution defines the thresholds on the anomaly scores to calculate the statistical significance of a GRB detection up to 5σ. We used a list of 1586 GRBs obtained from the GRBWeb catalog to test the trained model. From this list, the DL model detected 72 GRBs with a significance $\geqslant 3\sigma$, 15 of which are not present in the second MCAL GRB catalog. Although the ACS cannot detect all GRBs reported in the GRBWeb catalog due to visibility issues, different energy ranges, and lower sensitivity compared to the instruments designed for GRB detection, these results confirm the detection capabilities of our DL model and the opportunity to provide supplementary information for the AGILE team. These results allow us to implement this detection method in the AGILE real-time analysis pipeline to follow up external science alerts, adding a new tool to perform automated analyses on data acquired by the ACS. After detecting a candidate GRB inside the ACS data, the pipeline will send an automated communication to the AGILE team to verify and share the detection with the community.

In the future, we plan to develop an algorithm to classify the anomalies detected by our model to automatically identify GRBs and avoid human intervention in validating the candidate GRB detections. In addition, we plan to use the method presented in this manuscript to analyze the RM of other AGILE instruments, such as MCAL or SA, by analyzing the full archive of AGILE RM data to perform a blind search analysis. Finally, we think that other gamma-ray facilities can use the method described in this work to develop DL models to detect transient events inside the MTS produced by their detectors.

The AGILE Mission is funded by the Italian Space Agency (ASI) with scientific and programmatic participation by the Italian National Institute for Astrophysics (INAF) and the Italian National Institute for Nuclear Physics (INFN). The investigation is supported by the ASI grant I/028/12/6 and I/028/12.7-2022. We thank the ASI management for unfailing support during AGILE operations. We acknowledge the effort of ASI and industry personnel at the ASI ground station in Malindi (Kenya), at the Telespazio Mission Control Center at Fucino, and the data processing done at the ASI/SSDC in Rome: the success of AGILE scientific operations depends on the effectiveness of the data flow from Kenya to SSDC and the data analysis and software management.

ORCID iDs

N. Parmiggiani https://orcid.org/0000-0002-4535-5329
A. Bulgarelli https://orcid.org/0000-0001-6347-0649
A. Ursi https://orcid.org/0000-0002-7253-9721
A. Macaluso https://orcid.org/0000-0002-1348-250X
A. Di Piano https://orcid.org/0000-0002-9894-7491
V. Fioretti https://orcid.org/0000-0002-6082-5384
A. Aboudan https://orcid.org/0000-0002-8290-2184
L. Baroncelli https://orcid.org/0000-0002-9215-4992
A. Addis https://orcid.org/0000-0002-0886-8045
M. Tavani https://orcid.org/0000-0003-2893-1459
C. Pittori https://orcid.org/0000-0001-6661-9779

References

Bulgarelli, A. 2019, ExA, 48, 199
Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep Learning (Cambridge, MA: MIT Press)
Kingma, D. P., & Ba, J. 2015, arXiv:1412.6980
Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, ApJL, 413, L101
Lecun, Y., Bengio, Y., & Hinton, G. 2015, Natur, 521, 436
Lien, A., Sakamoto, T., Barthelmy, S. D., et al. 2016, ApJ, 829, 7
Munir, M., Siddiqui, S. A., Dengel, A., & Ahmed, S. 2019, IEEE Access, 7, 1991
Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. 2021, ACM Comput. Surv., 54, 1
Parmiggiani, N., Bulgarelli, A., Beneventano, D., et al. 2022, A&C, 39, 100570
Parmiggiani, N., Bulgarelli, A., Fioretti, V., et al. 2021, ApJ, 914, 67
Parmiggiani, N., Bulgarelli, A., Ursi, A., et al. 2021, Proc. ICRC (Berlin), 37, 933
Perotti, F., Fiorini, M., Incorvaia, S., Mattaini, E., & Sant'Ambrogio, E. 2006, NIMPA, 556, 228
Pittori, C. & The AGILE-SSDC Team 2019, RLSFN, 30, 217
Tavani, M., Barbiellini, G., Argan, A., et al. 2008, NIMPA, 588, 52
Tavani, M., Barbiellini, G., Argan, A., et al. 2009, A&A, 502, 995
Thill, M., Konen, W., Wang, H., & Bäck, T. 2021, Appl. Soft Comput., 112, 107751
Ursi, A., Romani, M., Verrecchia, F., et al. 2022, ApJ, 925, 152
Ursi, A., Tavani, M., Verrecchia, F., et al. 2019, ApJ, 871, 27
Zhang, A., Zhao, X., & Wang, L. 2021, in 2021 IEEE 5th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC) (Piscataway, NJ: IEEE), 571
Zhao, B., Lu, H., Chen, S., Liu, J., & Wu, D. 2017, J. Syst. Eng. Electron., 28, 162
Zheng, Y., Liu, Q., Chen, E., Ge, Y., & Zhao, J. L. 2016, Front. Comput. Sci., 10, 96