

Deep learning detection of transients

Iftach Sadeh*

DESY, Platanenallee 6, 15738 Zeuthen, Germany

E-mail: iftach.sadeh@desy.de

The next generation of observatories will facilitate the discovery of new types of astrophysical transients. The detection of such phenomena, whose characteristics are presently poorly constrained, will hinge on the ability to perform blind searches. We present a new algorithm for this purpose, based on deep learning. We incorporate two approaches, utilising anomaly detection and classification techniques. The first is model-independent, avoiding the use of background modelling and instrument simulations. The second method enables targeted searches, relying on generic spectral and temporal patterns as input. We compare our methodology with the existing approach to serendipitous detection of gamma-ray transients. We use our framework to derive the detection prospects of low-luminosity gamma-ray bursts with the upcoming Cherenkov Telescope Array. Our method is an unbiased, data-driven approach for multiwavelength and multi-messenger transient detection.

*36th International Cosmic Ray Conference -ICRC2019-
July 24th - August 1st, 2019
Madison, WI, U.S.A.*

*Speaker.

1. Introduction

Transient astrophysical events at high energies enable the study of a broad range of fundamental phenomena. We present a new method, intended for the detection of such events. Our algorithm is based on deep learning, using recurrent neural networks. It employs two complementary approaches, optimised for both model-independent and targeted searches, respectively denoted as *anomaly detection*, and *classification*. Our approach is by design generic. The methodology is not restricted to a specific energy regime, type of input, or time scale. In particular, it is well suited for multiwavelength and multi-messenger searches, where different observables are combined. It can therefore easily be adapted for many types of transient searches. We illustrate the method for the case of serendipitous detection of low-luminosity gamma ray (γ -ray) bursts (GRBs; LL-GRBs) with CTA, the upcoming Cherenkov Telescope Array.¹

2. New detection methods

Machine learning is widely used in astronomy [1–6]. In the current study, we utilise a recurrent neural network (RNN), made up of long short-term memory (LSTM) units. RNNs are a type of artificial neural network, which is well suited for time series analysis [7, 8]. For a review of deep learning and RNNs, see [9].

We utilize the open-source software, *tensorflow*, for our implementation [10]. The architecture of the RNN is illustrated in Figure 1. The network accepts an input which corresponds to 25 time steps, each representing a 1 s interval of γ -ray data. The different steps are implemented as RNN cells. A cell is composed of a pair of LSTM layers, respectively comprising 128 and 64 hidden units in the current implementation. The network may be decomposed into an *encoder* and a *decoder*. The encoder receives 20 time steps as input. A potential transient signal event is then searched for within the final 5 time steps.

Anomaly detection represents a model-independent approach, where transient events are identified based on their divergence from the expected background. This simple methodology is completely data-driven, and is able to adapt to real-time evolution of the background. Anomaly detection is decoupled from instrument simulations. It is insensitive to uncertainties on modelling of the background, the atmosphere, or other observational artefacts which do not evolve strongly in time. In the current example, Poissonian statistics are assumed for both the background and the signal models.² The background in this case is derived in-situ, using data exclusively from within the ROI for the source. The network is trained using background-only events for all 25 time steps. The test statistic for detecting a transient is based on the 5 decoder time steps. It encapsulates the difference between the counts predicated by the RNN, and actual observations.

Our second method employs classification. In this case, we train an estimator to identify transient patterns. The classification approach increases the sensitivity of specific searches, such as for LL-GRBs. For instance, the time structure of transient events is naturally incorporated as a part of the training process, avoiding the need for explicit modelling. Further more, simple training

¹ CTA: <https://www.cta-observatory.org/>.

² For brevity, we refer to the background-only hypothesis as the *background model*, and to the background+signal hypothesis as the *signal model*.

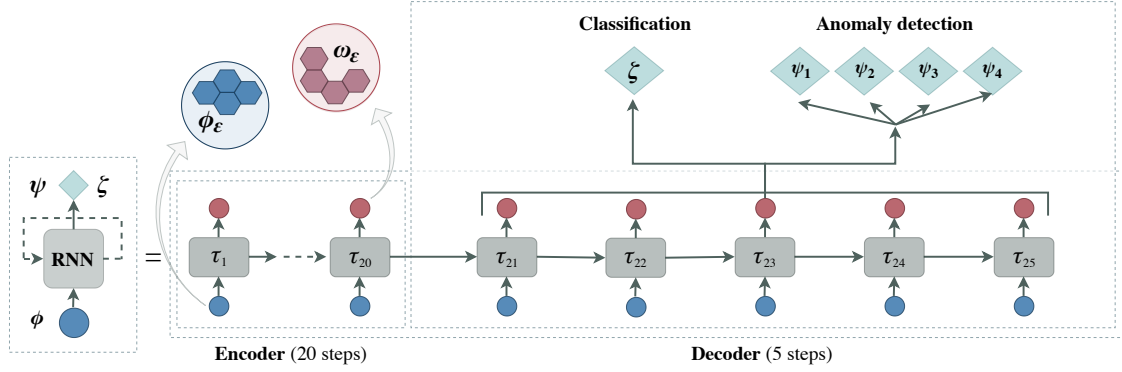


Figure 1: Schematic design of the RNN used in this study. The network may be decomposed into an encoder and a decoder, respectively representing 20 and 5 time steps, τ , of LSTM units (rectangles). The input data, ϕ , (blue circles) make up 4 numbers for each time step (blue hexagons), corresponding to γ -ray counts in different energy bins, ϵ . The direct output of the LSTMs, $\omega_{\epsilon, \tau}$, (red hexagons) are predictions for event counts for each step and energy bin. The output of the RNN (diamonds) is different for each implementation. For anomaly detection, ψ_{1-4} , provide the predicted background counts per energy bin, integrated over the decoder time steps. The output of the classification method, ζ , is used to derive a probability density function.

examples are sufficient for subsequent detection of more complicated intrinsic spectra, as discussed below. The network is trained using labelled examples of background and signal events.

The output of the RNN, ζ , is the inferred classification metric for a given event, (see Figure 2(a) below). The test statistic for identifying a signal event is derived from the distributions of classified signal and background events, respectively denoted as ζ_{sig} and ζ_{bck} . It is defined as,

$$\text{TS} = -2 \log \left(\frac{\zeta_{\text{bck}}}{\zeta_{\text{sig}}} \right), \quad (2.1)$$

following the prescription of Ref. [11].

For additional details on these methods, see Ref. [12].

3. Transient simulations

One of the interesting source populations that might be fully unveiled in the near future is that of low-luminosity γ -ray bursts [13, 14]. LL-GRBs are distinguished by low isotropic equivalent luminosities, generally, $10^{46} < L_{\gamma, \text{iso}} < 10^{49} \text{ erg s}^{-1}$. A sub-class of the population of long γ -ray bursts (GRBs), they have been connected to mildly relativistic supernovae [15], and are potential sources of ultra high-energy cosmic rays and neutrinos [16–18]. There are indications that the observable rate of LL-GRBs in the local Universe (redshift, $z < 0.1$) is high, of the order of $200 \text{ Gpc}^{-3} \text{ yr}^{-1}$ [19]. They are thus appealing targets for blind real-time searches. Observation of LL-GRBs is challenging with the current generation of ground- and space-based observatories³ [19,

³ E.g., *Fermi*: <https://fermi.gsfc.nasa.gov/>; *H.E.S.S.*: <https://www.mpi-hd.mpg.de/hfm/HESS/HESS.shtml>; *MAGIC*: <https://www.magic.mpp.mpg.de/>; *VERITAS*: <https://veritas.sao.arizona.edu/>.

20]. However, the upcoming CTA will significantly enhance their detection prospects, allowing γ -ray measurements down to ~ 20 GeV, within a large field-of-view (FoV).

We use the following as reference events for simulating the possible γ -ray signatures of LL-GRBs: GRBs 080916C, 090323, 090510, 090902B, and 110731A. These are bright high-luminosity GRBs, which have been detected at high energies with *Fermi*-LAT (see Ref. [21] and references therein). We assume that the GeV emission of these bursts is a simple power law (PL) extension of a Band-like model to high energies [22]. We only consider those bursts which exhibit durations of the order of tens of seconds. We thus exclude the population of ultra-long GRBs, which might involve unique emission mechanisms, such as shock breakouts [23, 24].

We randomly shift the reference GRBs in redshift and luminosity to the expected ranges for LL-GRBs, scaling the observed flux accordingly [25]. In order to simulate the signals at GeV energies, we nominally assume a simple spectral/temporal PL model,

$$M_{\text{PL}}(E, t) = k_0 \left(\frac{E}{E_0} \right)^{-\Gamma} t^{-\tau}. \quad (3.1)$$

The *prefactor* and *pivot energy*, k_0 and E_0 , are derived directly from the flux of the GRB. The spectral index, Γ , and temporal decay index, τ , are randomly selected for each event, uniformly distributed within $1.9 < \Gamma < 2.7$ and $0.8 < \tau < 2$. These properties generally correspond to the expectations for the low-luminosity population.

We generate CTA events using the open-source software, *ctools*, [26] simulating the Northern array [27], using the publicly available IRFs (version *prod3b-v1*). We exclusively use IRFs optimised for 30 min observations at zenith angles of $\sim 20^\circ$. The region of interest (RoI) for the simulation is chosen as a circular region with a radius of 0.25° . The circle is centred at the position of the source, which is displaced by 0.5° from the centre of the FoV.

The inputs to the RNN corresponding to a given time step are event counts in 4 logarithmically-spaced energy bins within $30 < E_\gamma < 200$ GeV. The inputs to the encoder are assumed to correspond to background-only counts in all cases. The input to the decoder and the output of the network depend on the type of inference being used. We conduct searches over the 5 s intervals that coincide with the beginning of bursts in the signal sample. We assume a conservative correction for trials [28], accounting for 100 h of observations at 1 s search intervals.

We explore the effects of the extragalactic background light (EBL) on the observed spectra. Comparing different models [29–31], we find that the EBL has little effect on the LL-GRBs in our simulation, due primarily to their low redshift.

4. Results

Using the trained RNN in the classification mode, we derive the distributions of the classification metric, ζ , for the background and signal samples, as shown in Figure 2(a). The corresponding relation between the classification test statistic, TS_{clas} , and ζ is presented in Figure 2(b). We evaluate our results using the *detectability* metric, $p_{\text{det}} = \langle \rho_{\text{det}} \rangle$, defined event-by-event for

$$\rho_{\text{det}}(t) = \begin{cases} 0, & t < \text{TS}_{5\sigma} \\ 1, & t \geq \text{TS}_{5\sigma} \end{cases}. \quad (4.1)$$

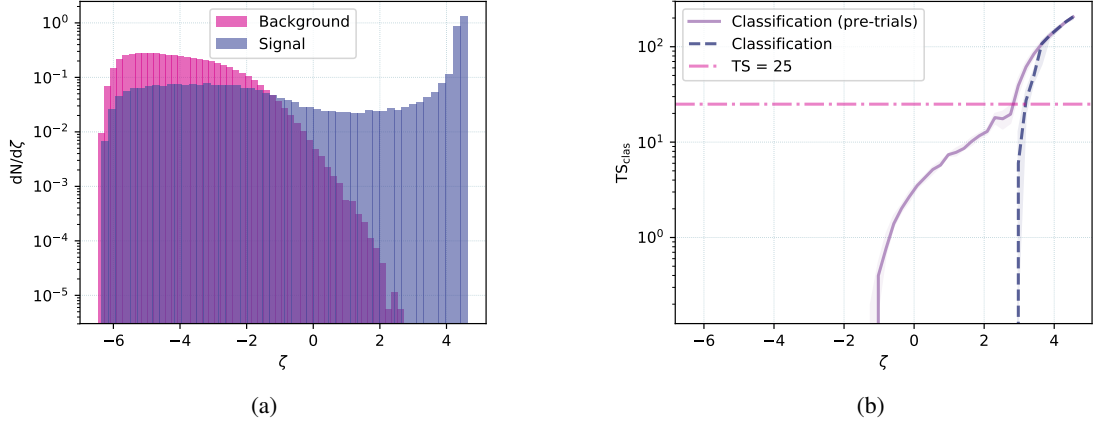


Figure 2: Parametrisation of the performance of the trained classification method. (a) Distributions of the classification metric, ζ , for the signal and background samples, as indicated. (b) The parametrised classification test statistic, TS_{clas} , as a function of ζ , before and after the correction for trials. The dashed-dotted horizontal line highlights the value, $TS = 25$.

Here t represents the test statistic derived for a given detection method; $TS_{5\sigma}$ is the corresponding threshold for a 5σ detection, where e.g., for a model with a single degree of freedom, $TS_{5\sigma} = 25$ [32].

Figure 3 shows f_{det} , the fraction of events with a TS value larger than a given threshold, as a function of this threshold. We find that *ctools* and the anomaly detection achieve comparable significance distributions. The two methods enable detection of a similar fraction of the events, with slightly better performance by *ctools*. The performance of the classification approach is better than that of *ctools*, with a relative improvement in detectability of $\sim 10\%$ on average.

It is important to verify that the new detection methods do not produce spurious detections, and that the corresponding test statistics are properly mapped to significance. We therefore evaluate the different algorithms on the background sample, and compare them to the reference *ctools* distribution. As shown in Figure 3(b), the anomaly detection and classification methods produce comparable or better (lower) rates of fake detections. For the given sample of 10^6 background simulations, none of the methods exceed a pre-trials TS value of 20, or a post-trials value of 1.

The dependence of p_{det} on the temporal and spectral indices of the simulated LL-GRBs is shown in Figure 4. One may observe that our new methods match or improve upon the performance of *ctools*. As expected, longer-lasting and harder spectra are more likely to be detected by all algorithms.

5. Summary

In this study, we present a new approach for source detection. Our algorithm is based on deep learning, utilising recurrent neural networks, which are ideally suited for time series analyses. The model can be used to evaluate observation sequences of second time scales with insignificant latency. The choice of technology is therefore particularly fitting for real-time searches.

We have developed two methods, based on anomaly detection and classification techniques. Anomaly detection represents a model-independent approach, where transient events are identified




Figure 3: Dependence of f_{det} , the fraction of events with a TS value larger than a given threshold, on the value of the threshold. The different detections methods are compared, derived for the signal (a) and background (b) samples, as indicated. The full lines in either figure correspond to the pre-trials test statistic. The dashed lines in (a) represent f_{det} after accounting for trials, where in (b) we found $f_{\text{det}}(TS > 1) = 0$ post-trials in all cases. For clarity, the relations are truncated to the range, $1 < TS < 100$, where the pre-trials background distributions in (b) do not extend beyond $TS \approx 20$. The dashed-dotted vertical line in (a) highlights the value, $TS = 25$.

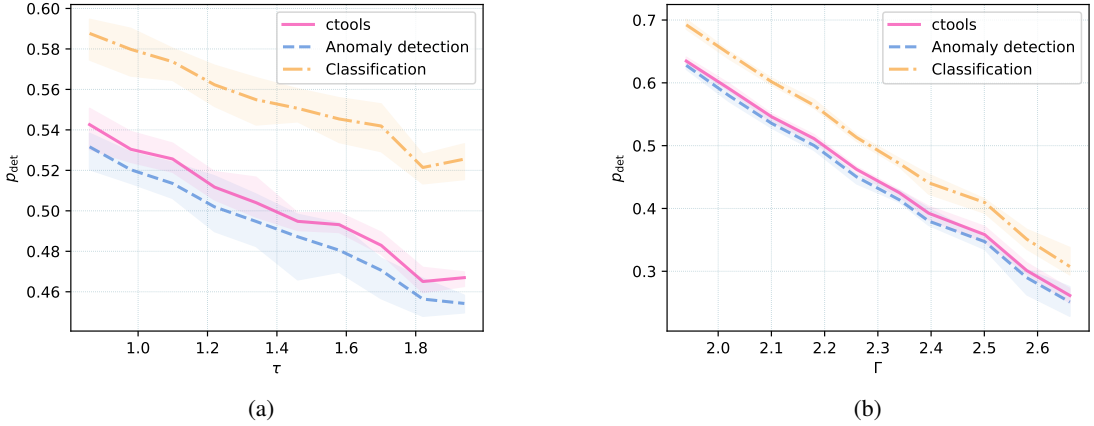


Figure 4: Dependence of p_{det} on the temporal (a) and spectral (b) indices of simulated LL-GRBs, after accounting for trials. The shaded regions correspond to 1σ uncertainties on the values of p_{det} , derived using the bootstrap method. The different detection methods are compared, as indicated.

based on their divergence from the expected background. The method is data-driven. We thus avoid the need for background modelling, as well as for explicit characterisation of the state of the instrument. The classification method allows one to perform targeted searches. In this case, the RNN is trained to identify generic transient patterns. The estimator provides high detection rates while maintaining low fake rates.

Acknowledgments

We would like to thank the following people for numerous useful discussions: D. Biehl, D. Boncioli, Z. Bosnjak, O. Gueta, T. Hassan, M. Krause, G. Maier, M. Nieves Rosillo, I. Oya, A. Palladino, E. Pueschel, R. R. Prado, L. Rauch, and W. Winter.

This research made use of *ctools* [26], a community-developed analysis package for Imaging Air Cherenkov Telescope data. *ctools* is based on *GammaLib*, a community-developed toolbox for the high-level analysis of astronomical gamma-ray data. This research has also made use of the CTA IRFs (version *prod3b-v1*), provided by the CTA Consortium and Observatory; see <http://www.cta-observatory.org/science/cta-performance/> for more details.

References

- [1] Iftach Sadeh, Filipe B. Abdalla, and Ofer Lahav. In: (2016). [Publ. Astron. Soc. Pac.128,104502(2016)]. arXiv: 1507.00490.
- [2] Maria Krause, Elisa Pueschel, and Gernot Maier. In: *Astroparticle Physics* 89 (Mar. 2017), pp. 1–9. arXiv: 1701.06928 [astro-ph.IM].
- [3] Itamar Reis et al. In: *MNRAS* 476 (May 2018), pp. 2117–2136. arXiv: 1711.00022 [astro-ph.IM].
- [4] Bryan Kim et al. In: *APS April Meeting Abstracts*. Vol. 2018. Jan. 2018, p. L01.031.
- [5] H. Domínguez Sánchez et al. In: *MNRAS* 484 (Mar. 2019), pp. 93–100. arXiv: 1807.00807 [astro-ph.GA].
- [6] Nima Sedaghat and Ashish Mahabal. In: *MNRAS* 476 (June 2018), pp. 5365–5376. arXiv: 1710.01422 [astro-ph.IM].
- [7] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. In: *NIPS’14*. Montreal, Canada: MIT Press, 2014, pp. 3104–3112.
- [8] Hongyu Shen et al. In: (2017). arXiv: 1711.09919 [gr-qc].
- [9] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. In: *Nature* 521 (May 2015), 436 EP–.
- [10] Martín Abadi et al. 2015.
- [11] Kyle Cranmer, Juan Pavez, and Gilles Louppe. In: (2015). arXiv: 1506.02169.
- [12] Iftach Sadeh. In: (2019). arXiv: 1902.03620 [astro-ph.HE].
- [13] Enwei Liang, Bing Zhang, and Z. G. Dai. In: *Astrophys. J.* 662 (2007), pp. 1111–1118. arXiv: astro-ph/0605200.
- [14] Francisco Virgili, Enwei Liang, and Bing Zhang. In: *Mon. Not. Roy. Astron. Soc.* 392 (2009), p. 91. arXiv: 0801.4751 [astro-ph].
- [15] Zach Cano et al. In: *Adv. Astron.* 2017 (2017), p. 8929054. arXiv: 1604.03549 [astro-ph.HE].
- [16] K. Murase et al. In: *Phys. Rev. D* 78.2, 023005 (July 2008), p. 023005. arXiv: 0801.2861.
- [17] B. T. Zhang et al. In: *Phys. Rev. D* 97.8, 083010 (Apr. 2018), p. 083010. arXiv: 1712.09984.

- [18] D. Boncioli, D. Biehl, and W. Winter. In: *ArXiv e-prints* (Aug. 2018). arXiv: 1808.07481.
- [19] Hui Sun, Bing Zhang, and Zhuo Li. In: *Astrophys. J.* 812.1 (2015), p. 33. arXiv: 1509.01592.
- [20] S. Abdollahi et al. In: *Astrophys. J.* 846.1 (2017), p. 34. arXiv: 1612.03165.
- [21] M. Ajello et al. In: *Astrophys. J. Suppl.* 209.1 (2013), p. 11. arXiv: 1303.2908.
- [22] D. Band et al. In: *ApJ* 413 (Aug. 1993), p. 281.
- [23] Eli Waxman, P. Meszaros, and S. Campana. In: *Astrophys. J.* 667 (2007), pp. 351–357. arXiv: astro-ph/0702450 [astro-ph].
- [24] A. J. Levan et al. In: *Astrophys. J.* 781 (2013), p. 13. arXiv: 1302.2352 [astro-ph.HE].
- [25] Susumu Inoue et al. In: *Astropart. Phys.* 43 (2013), pp. 252–275. arXiv: 1301.3014 [astro-ph.HE].
- [26] J. Knodlseder et al. In: *Astron. Astrophys.* 593 (2016), A1. arXiv: 1606.00393.
- [27] B. S. Acharya et al. In: (2017). arXiv: 1709.07997 [astro-ph.IM].
- [28] S.D. Biller. In: *Astroparticle Physics* 4.3 (1996), pp. 285–291.
- [29] Alberto Franceschini, Giulia Rodighiero, and Mattia Vaccari. In: *Astron. Astrophys.* 487 (2008), p. 837. arXiv: 0805.1841.
- [30] A. Dominguez et al. In: *MNRAS* 410.4 (2011), pp. 2556–2578.
- [31] Rudy C. Gilmore et al. In: *MNRAS* 422.4 (2012), pp. 3189–3207.
- [32] S. S. Wilks. In: *Annals Math. Statist.* 9.1 (1938), pp. 60–62.