



Full length article

flashcurve: A machine-learning approach for the simple and fast generation of adaptive-binning light curves with Fermi-LAT data

T. Glauch^{a,b,c,d}, K. Tchiorniy^a *^a Technische Universität München, Physik-Department, James-Frank-Str. 1, Garching bei München, D-85748, Germany^b Institute for Advanced Study, Technische Universität München, Lichtenbergstrasse 2a, Garching bei München, D-85748, Germany^c Heidelberg University, Institute of Environmental Physics, Im Neuenheimer Feld 229, Heidelberg, D-69120, Germany^d Deutsches Zentrum für Luft- und Raumfahrt, Institut für Physik der Atmosphäre, Oberpfaffenhofen, Germany

ARTICLE INFO

Keywords:

Machine learning

Gamma-rays

Light curves

Active galaxies

BL lacertae objects

Data analysis

Astronomical databases

ABSTRACT

Gamma rays measured by the Large Area Telescope (LAT) on board the *Fermi Gamma-ray Space Telescope* tell us a lot about the processes taking place in high-energetic astrophysical objects. The fluxes coming from these objects are, however, extremely variable. Hence, gamma-ray light curves optimally use adaptive bin sizes in order to retrieve most information about the source dynamics and to combine gamma-ray observations in a multi-messenger perspective. However, standard adaptive binning approaches are slow, expensive and inaccurate in highly populated regions. Here, we present a novel, powerful, deep-learning-based approach to estimate the necessary time windows for adaptive binning light curves in *Fermi*-LAT data using raw photon data. The approach is shown to be fast and accurate. It can also be seen as a prototype to train machine-learning models for adaptive binning light curves for other astrophysical messengers.

1. Introduction

Gamma rays are photons at the high-energy end of the electromagnetic spectrum, with energies starting from a few hundred keV and going up to ultra-high energies in the TeV and PeV ranges. They are produced in various processes like nuclear decays and the collision of high energy particles (Gaisser et al., 2016). Hence, they are a unique way to study the high-energy phenomena in our Universe.

The Large Area Telescope (LAT) on board the *Fermi Gamma-ray Space Telescope* is a satellite that studies gamma-rays from Galactic and extragalactic environments in the energy range between 100 MeV up to 1 TeV. Several thousand gamma-ray point sources have been detected by *Fermi*-LAT in the extragalactic sky (Ballet et al., 2024). The sources at high Galactic latitudes are mostly blazars – active galactic nuclei with a jet of high energy particles pointing towards the observer – but also starburst-, radio galaxies and narrow-line Seyfert 1 galaxies (Ajello et al., 2015; D’Ammando, 2019). It has been shown that the gamma-ray fluxes from those objects are, in most cases, highly variable (Sbarrato et al., 2011). In order to catch the dynamics of the processes in those sources, it is necessary to calculate time-dependent light curves. This gets even more important in the emerging field of multi-messenger astronomy, which tries to combine (time-dependent) measurements of different particles (e.g. photons, neutrinos, protons) at different

energies (for sub-eV to PeV) to get a comprehensive picture of the processes happening within our Universe. Many of those studies already use such approaches or could greatly profit from doing so (Padovani et al., 2018; Giommi et al., 2020).

Due to the high variability and varying strength of gamma-ray sources, it is challenging to choose an adequate time width for each light curve bin such that the integration time is long enough to catch the signal but short enough to catch relevant features. For a long time, the preferred method for choosing the (adaptive) length of the time bins was the so-called adaptive binning approach (Lott et al., 2012) that aims to produce light curves with bins having a constant significance or flux uncertainty Δ_0 . Numerically, this requires to solve the equation

$$\sigma_f(T_0, T_1, \bar{\gamma}, \bar{F}) = \Delta_0 \quad (1)$$

for the ending time of a time interval T_1 giving some starting time T_0 and some function σ_f . $\bar{\gamma}$ and \bar{F} are the average flux and photon spectral index over the interval $[T_0, T_1]$, respectively. However, the (sequential) search for the time bins fulfilling this equation is extremely expensive, so the computation of σ_f requires approximations instead of a full likelihood model. Nevertheless, with traditional methods, a single adaptive light curve often takes several hours, if not days, to compute. This paper presents a novel approach for the (fast) generation of

* Corresponding author.

E-mail address: kristiantcho@gmail.com (K. Tchiorniy).

adaptive binning for light curves that uses a machine-learning approach to approximate the *Fermi* likelihood function.

In recent years, machine-learning-based methods have become quite popular in astronomy as they are quick in the calculation and provide a high accuracy (Fluke and Jacobs, 2020; Sen et al., 2022; Glauch et al., 2022; Kronmueller and Glauch, 2019). As a result, our new tool,¹ *flashcurve*, allows large-scale generation of adaptive-binning for gamma-ray light curves as well as quick follow-up studies in astronomical real-time programs.

The paper is structured as follows. In Section 2, we present the dataset of light curves that we use to train the machine-learning estimator; in Section 3, we explain its structure and the training process; and in Section 4, we show how it can be used. We conclude with a summary in Section 5.

2. Data sample of γ -ray light curves

2.1. *Fermi*-LAT LCR & analysis

The *Fermi*-LAT Light Curve Repository (LCR) (Fermi-LAT, 2023) is a database of light curves including 1525 sources from the 10-year *Fermi*-LAT point source 4FGL-DR2 catalogue (Ballet et al., 2020). This database and the corresponding raw photon data downloaded from the *Fermi*-LAT data server were used to train the machine-learning estimator model presented here.

The sources included in the LCR database have a variability index greater than 21.67, indicating the average fractional flux variability (dF/F) measured on timescales of 1 year. Having a variability over this value for over 12 intervals means that these sources are estimated to have a <1% chance of being a steady source (Fermi-LAT, 2020). The resulting source selection comprises blazars, further classified as flat spectrum radio quasars, BL Lacs, and blazar candidates of unknown type.

Each source in the catalogue has light curves with three different fixed time-bin lengths: 3 days, 1 week, and 1 month. This was done for over 14 years of data, from 2008 August to 2023 December, which accumulated to over 3.7 million individual time bins.

The gamma-ray flux of each time bin was estimated using a maximum likelihood analysis as described in Fermi-LAT (2009). The signal hypothesis of the likelihood function includes a power-law point source flux ϕ on top of a diffuse gamma-ray background and the known gamma-ray sources in the region. Free parameters are the flux normalisation ϕ_0 , and the spectral index γ i.e.,

$$\phi(E_\gamma; \gamma, \phi_0) = \phi_0 \times (E_\gamma/E_0)^{-\gamma} \quad (2)$$

with the photon energy E_γ . The test statistic (TS) quantity is defined as

$$TS = -2 \ln (L_{\max,0}/L_{\max,1}) \quad (3)$$

where $L_{\max,0}$ is the maximum likelihood value for a model without an additional source (the null hypothesis) and $L_{\max,1}$ is the maximum likelihood value for the signal hypothesis.

In order to maximise the signal and background likelihood, the observed energy and direction of the photons are used with the LAT instrument response function (IRF). The exact location of the source is found through a grid search for the maximum test-statistic value. The analysis settings for the *Fermi*-LAT LCR analysis can be found in Table 1. Our machine-learning algorithm uses the same data settings regarding energy range, event selection, IRF and acceptance cone. For *Fermi* point sources analyses, the TS is distributed approximately as a χ^2 distribution (Fermi-LAT, 2009); thus, the square root of the TS is approximately equal to the detection significance for a given source (Wilks, 1938). This TS will be the target variable for our machine-learning algorithm using the raw photon data as input. The LCR analysis rejects the background hypothesis only when $TS \geq 4$.

Table 1

Analysis settings for the *Fermi*-LAT LCR analysis.

Time Bins	3 day, 7 day, 30 day
Energy range	100 MeV–100 GeV
Event selection	P8R3_SOURCE
IRF	P8R3_SOURCE_V2
Acceptance Cone (ROI)	12 deg (radius)
Zenith Angle Cut (zmax)	90 deg
Fit optimiser	MINUIT
Galactic interstellar emission model	gll_iem_v07.fits
Isotropic spectral template	iso_P8R3_SOURCE_V3_v1
4FGL-DR2 Catalog	gll_psc_v27.fit




Fig. 1. An example of an LCR 'sanity check' for time bins with analyses which may have failed. If a flux-to-flux uncertainty ratio within a time bin wildly deviates from an approximately linear behaviour over the square root of the test statistic, this time bin should not be used in high-level analysis.

Source: Plot taken from Fermi-LAT (2023).

To obtain the LCR light curve data conveniently, we used the Python API script *pyLCR*,² which was used to get the exact time bins and corresponding TS for each source. We include time bins with a minimum TS of 0. By including time bins with $TS < 4$, we aim to train the estimator so that it possibly recognises the background. For these time bins, however, only upper limits for fluxes are provided.

2.2. Data sample cleaning

The LCR data source comes with some caveats. One caveat was that some results may have come from non-convergent analyses. Thus, we removed these samples from the final data sample selection by checking for non-zero MINUIT (James and Roos, 1975) return codes. Some data points had negative TS values, resulting from non-convergent fits. Those data points were also removed.

Apart from this, the LCR usage notes³ recommend performing 'sanity checks' to check whether the flux uncertainty to flux to ratio, σ_F/F , behaves approximately linearly to the square root of the TS, as well as to check whether fluxes are deviating quite strongly from the total distribution of fluxes for a given source. Fig. 1 taken from Fermi-LAT (2023) shows an example of data which deviates from the linear approximation. However, we found no deviations in our selected sources when performing such a sanity check with the data.

Finally, we removed sources within 10 degrees of the Galactic plane to avoid confusion with the Galactic background, for a total of 163 sources. The LCR database also includes a few extended sources, which we excluded. All of the aforementioned filters amount to a total of 1362 sources and around 1.5 million time bins that could be used to train the estimator model. Fig. 2 shows the distribution of the TS for all the selected time bins.

² <https://github.com/dankocevski/pyLCR>, last accessed: 01.04.2024.

³ <https://fermi.gsfc.nasa.gov/ssc/data/access/lat/LightCurveRepository/about.html>, last accessed: 01.10.2024.

¹ Available at <https://github.com/kristiantcho/flashcurve>.




Fig. 2. Binned distribution of test statistic from the *Fermi* LCR time bins which were to be used for training the estimator model.

3. Machine-learning-based adaptive light curve

3.1. General idea

Artificial neural networks have made huge strides in the last few years, especially in computer vision. Many techniques and architectures have been developed to increase the prediction quality of image recognition, both in categorical and regression estimation (Chai et al., 2021; Zhao et al., 2024). Similarly, we use the spatial and energy information of the time-integrated *Fermi*-LAT observations to estimate the detection significance.

For training data, we use the previously discussed *Fermi*-LAT LCR dataset. Although this dataset only covers light curves with three different time-bin lengths, it covers an extensive range of TS values due to the sources' variability. Moreover, sources are distributed isotropically in the sky, avoiding directional bias in the training.

3.2. Data preparation

Overall, our training dataset contains 1.5 million time bins. We generated neural network input images for each time bin by binning the photon counts in right ascension (RA), declination (Dec), energy, and time.

Similar approaches have been used in Caron et al. (2018), which binned photon counts in right ascension, declination and several energy bins as image channels instead of a single (broad range) channel.

The generation of the input images for the neural network has two steps. Firstly, the photons are rotated into a spherical coordinate system in which the source is located at a pole. We do this to standardise the angular distances of each photon event from the source. Secondly, the photons are binned into four dimensions: time, energy, and the two tangential dimensions of the rotated spherical coordinates, which we denote as X and Y in the flattened image (see Fig. 4). In order to avoid creating an exceedingly large feature space, only six energy bin channels were used. Thus, similar to energy bins typically used in *Fermi*-LAT spectral energy distributions, energy was binned in half powers of 10 MeV, (10^2 MeV, $10^{2.5}$ MeV, 10^3 MeV, $10^{3.5}$ MeV, 10^4 MeV, $10^{4.5}$ MeV, and $3 \cdot 10^5$ MeV). Since the point spread function of *Fermi*-LAT decreases with energy (Fig. 3), we decrease the range of the X-Y histogram with increasing energy (Fermi-LAT, 2013) to 12 deg, 5 deg, 3 deg, 1.5 deg, 1 deg, and 0.6 deg following the 95% quantile of the *Fermi*-LAT point spread function. In each dimension, 56 bins are used.

Figs. 4 and 5 illustrate the neural network inputs generated with this procedure. The first image is from 4FGL J0319.8+4130, and the second




Fig. 3. 68% and 95% containment angles of the acceptance weighted PSF for both the front/back PSF event types for the LAT P8R3_SOURCE_V2 IRF. Plot taken from https://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass8r3v2/lat_Performance.htm, last accessed: 01.03.2024.




Fig. 4. Example neural network input image for 4FGL J0319.8+4130 used to train the estimator model. The total input has six (energy bin) channels with photon counts binned in angular distance for a tangential plane around the source. The colour bars indicate counts in each bin. Each panel has counts of photon events within a specified energy range (energy bin). The source depicted here has a TS of 569.12.

comes from the source 4FGL J2253.9+1609, which had the highest significance time bin of all samples. Both images depict time bins of 1 month with significant gamma emission (TS of ~ 570 and $\sim 1.3 \cdot 10^5$ corresponding to significances of $\sim 25\sigma$ and $\sim 360\sigma$ respectively).

Fig. 5 visually shows clear emission from the source in all energy channels, while in Fig. 4, it is less visible to the human eye. Especially at low energies, photons from a nearby source are dominating the image. The picture is clearer at higher energies, contributing most to the overall significance.

3.3. The artificial neural network - architecture & training

3.3.1. Convolutional & residual neural networks

Convolutional neural networks (CNN) are a class of deep neural networks (DNN) that are designed to be trained to make predictions using visual data or images (LeCun et al., 2015). A CNN generally takes an N-dimensional grid input and uses convolutional layers and non-linear activation functions to map the input to an M-dimensional output. The input can be, for example, a multichannel image with dimensions (w, h, c), w, h being the width and height of the image in terms of the number of pixels and then c the number of channels. In the case of a colour image, we have three channels: red, green, and blue. In our case, we have six energy bins. The output could be a set of predictions for different image class options or, as in our case, a single number (regression problem).

Additional improvements in network performance can be achieved using residual neural networks (ResNets) (He et al., 2015). Instead of




Fig. 5. Same as the figure before, but using data from the source 4FGL J2253.9+1609, which, in the given time window, had a TS of 129479.46, making it the most significant time bin in the data sample.




Fig. 6. The basic working principle of a residual block used in ResNets, which consists of an input x having transformations $F(\cdot)$ applied to it, whilst having an identity mapping added onto its output in the form of a ‘skip connection’: $x + F(x)$. Image taken from Wikimedia Commons <https://commons.wikimedia.org/wiki/File:ResBlock.png>, last accessed: 01.10.2024.

learning independent features in each layer, a residual neural layer only learns corrections to the previously learnt features. This can be implemented by including an identity connection between subsequent layers, as illustrated in Fig. 6.

A ResNet consists of a sequence of multiple residual blocks. This is because stacking more layers in a typical DNN could lead to a swift reduction in training accuracy, often referred to as the ‘degradation’ problem (He et al., 2015). If a DNN is able to reach a maximal training accuracy with the initial layers (or some other sub-combination of layers inside the architecture), then these skip connections would ensure that only the additional unnecessary layers are skipped and that only the identity mapping is kept. ResNets have been shown to increase training stability and, therefore, accuracy (Zhang et al., 2024).

3.3.2. Flashcurve architecture & training

Fig. 7 depicts the network architecture of *flashcurve*. The model consists mainly of the previously mentioned residual blocks, each with a 3×3 convolutional layer with a stride of 1 and a filter count of 32, 64, or 128. At the end of each layer, we use the *RELU* function as a non-linear activation layer.

The identity mapping is performed with a 1×1 convolutional layer, which ensures that the number of channels in both layers matches when added at the end of the residual block.

Some preprocessing steps are performed in the input layer: First, a rescaling layer normalises the counts in each energy channel using the maximum values found in the training data set. In addition, we used a random flip layer as well as a rotation layer, which rotates the images

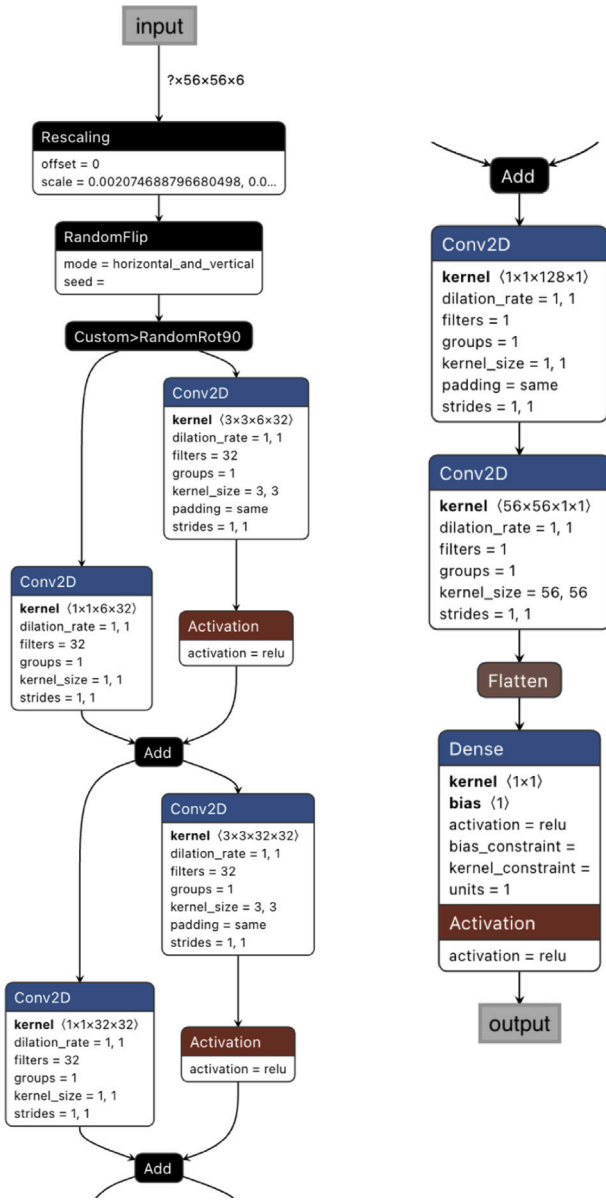


Fig. 7. Left: First layers of the *flashcurve* estimator model. All layers between two consecutive 'Add' layers comprise a residual block. Included here are also the 'Rescaling' preprocessing layer, which normalises the images, and the 'Random Flip' and 'Random Rotation' layers, which are only used during training to horizontally flip or rotate the images in 90-degree steps.

Right: Last layers after the residual blocks, including a 'Dense' layer with a *RELU* activation, which acts as a linear output for the TS estimation.

in random multiples of 90 degrees. Those data augmentations increase the size of the training dataset by a factor of 8, allowing a better model generalisation (Perez and Wang, 2017).

In the network's last layer, the network features are mapped to a single number, the test statistic, see Fig. 7. The final linear layer has another *RELU* activation to ensure that the TS value is a strictly positive number.

We trained the model with a batch size of 256 using the *Tensorflow* python package, the *Keras* API, the *ADAM* (Kingma and Ba, 2017) optimiser, and a mean-squared error (MSE) loss function. Table 2 summarises our model's total number of residual blocks. This combination of layers was chosen such that the number of parameters is one order of magnitude smaller than the total number of samples. In total, the network has ~400000 trainable parameters.

Table 2

Ordered list of residual blocks and the corresponding number of filters used in the convolutional layer within each block used for the final estimator model.

# of Res. Blocks	# of filters
15	32
5	64
1	128




Fig. 8. Learning curve displaying the output of the mean squared error loss function during the training of the best estimator model at different epochs for both the validation and training set.




Fig. 9. Predicted TS from an unseen test set of images vs. the true TS. To simplify the plot, mean, median, and the 10% to 90% quantile region of binned values are taken and plotted against the medians of the binned true TS.

3.4. Validation & performance

In order to test the model performance and its ability to generalise, we split the dataset into three sets: 80% was used as a training set, 10% as a validation set, and the remaining 10% as a test set, which would remain unseen until the model completed training. The validation set is used during training to test the model's performance after each epoch. It is also used to select the final model based on the epoch at which the model produced the best validation loss. Fig. 8 shows the output of the loss function over all the epochs for our best-performing model.

Fig. 9 shows the model's performance on the unseen test set, which achieved approximately constant relative uncertainty across the entire TS space.

To evaluate how the size of the training dataset impacts the model performance, we rerun the training procedure with different subsets of the training dataset, all while still keeping the same validation and test




Fig. 10. Learning curve displaying the output of the mean squared error (MSE) loss function during the training of an estimator model at different epochs for different subsets of the main training set.




Fig. 11. Learning curve displaying the output of the mean squared error (MSE) loss function during the training of an estimator model at different epochs, with different subsets of the main training set, and testing then on the main validation set.

set. The results are shown in Figs. 10 and 11. The losses on the training set show that more data leads to a more stable training; meanwhile, the validation losses indicate that lower losses are achieved with more data. This indicates that the amount of training data available currently limits our model.

3.5. Investigation of bad predictions

We find two types of bad predictions from the test dataset: the model predicts a low TS, whereas the true TS is orders of magnitude higher, or vice versa. The former case was investigated by performing the checks recommended by Fermi-LAT (2023) and then subsequently running *Fermipy* to cross-check the TS. One such case occurred with the source 4FGL J1234.0-5735, in which a TS of ~ 2 was predicted, but the *Fermi*-LCR stated a TS of ~ 300 . Figs. 12–14 show the recommended checks performed.

The badly predicted TS does not appear to be an outlier when looking at Fig. 12, which, according to the *Fermi*-LCR, would be the main reason to investigate further. However, looking at the distributions of flux and flux uncertainty, this data point exhibits a sudden burst of flux compared to the typical lower activity of the source. Fig. 15 depicts the respective image with which the model under-predicted the TS.




Fig. 12. Set of TS vs. their corresponding squared ratio of flux/flux uncertainty from the *Fermi*-LCR for a 3-day binned lightcurve of 4FGL J1234.0-5735. The red cross indicates a TS, which the estimator had vastly under-predicted. However, it does not appear very clearly here to be an outlier.




Fig. 13. Log-binned counts of flux data of 4FGL J1234.0-5735, see 12 for more details. The badly predicted TS (red cross) appears to be an outlier from the distribution.

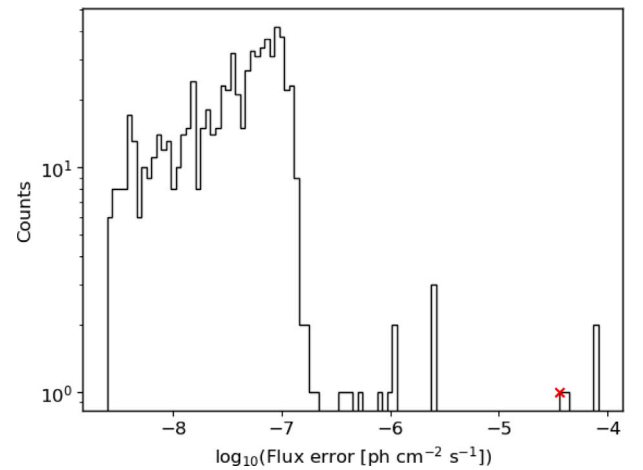


Fig. 14. Log-binned counts of flux uncertainty data of 4FGL J1234.0-5735, see 12 for more details. The badly predicted TS (red cross) also appears to be an outlier from the distribution.




Fig. 15. Image of binned photon event counts for 4FGL J1234.0-5735 in a 3-day period, corresponding to a TS ~ 300 , for which the estimator model under-predicted a TS of only ~ 2 . See Fig. 4 for more information.

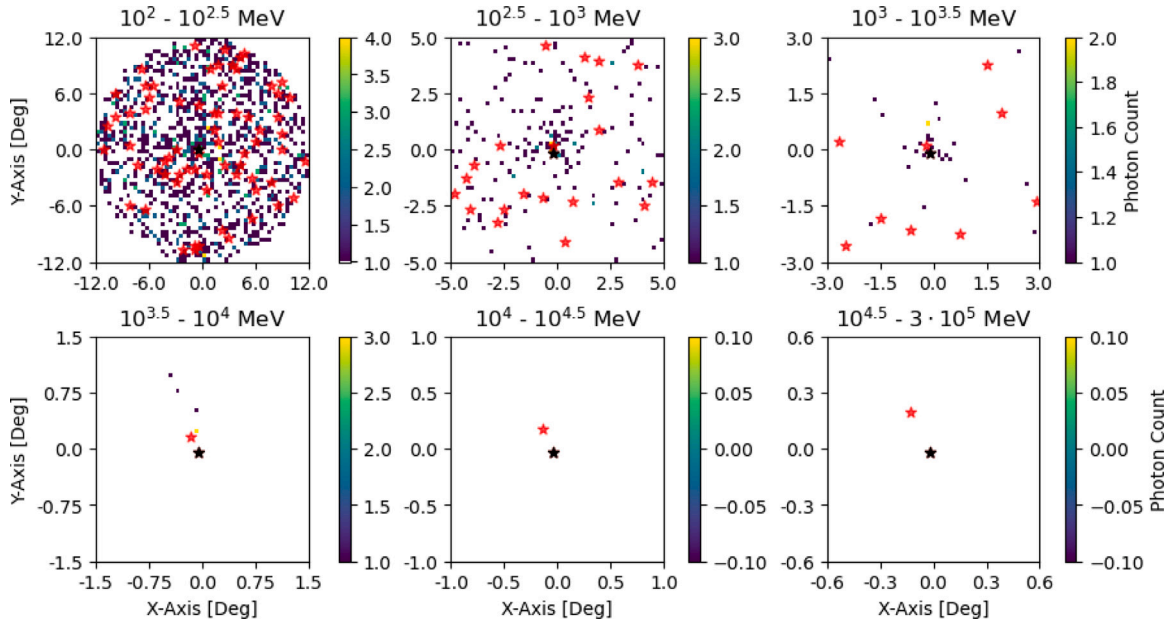


Fig. 16. Same as the figure above for 4FGL J1311.0+3233 and a time bin of 1 week. *Fermi*-LCR reports a TS of 3.63, whereas the estimator model predicted ~ 300 . Additionally, overlaid on this image as red stars are the (relative) positions of nearby sources; the black star is 4FGL J1311.0+3233. It is apparent in this image that there is a high amount of activity close to the central source.

Finally, the time bin (3-days centred at 673617601 MET) was rerun independently through *Fermipy* and along with the 3-day bin before and after the time bin in question, all of which came out with extremely low flux ($O(10^{-10})$) and extremely low TS ($O(10^{-5})$). Thus, we concluded that this was an outlier in the LCR database that was missed during the initial data cleaning.

However, judging by the flux and flux uncertainty distributions, one can see that the expected amount of outliers per source could be in the order of 10, which would mean that the total number of expected outliers in the entire data set would be in the order of 10^4 , two orders of magnitude less than the total number of data-points. Hence, the impact of these outliers on the model is expected to be small. For the future, we plan to create more stringent filters for the data cleansing, for example,

by filtering out data points whose flux or flux uncertainty sit outside 5% and 95% quantiles as in Figs. 13 and 14.

In the opposite case, there were numerous over-predictions from the model, one example of such being a prediction of ~ 300 for a true TS of ~ 3.6 . One of the sources in question was 4FGL J1311.0+3233 for a time bin centred at 242308801 MET with a duration of 1 week. Fig. 16 shows the image with photon counts along with the relative positions of nearby sources overlaid. We see that there is a considerable amount of gamma-ray sources close to 4FGL J1311.0+3233, even at relatively small angular distances, the most prominent being OP 313, with a source strength (for photon flux > 1 GeV) ~ 4 times larger than 4FGL J1311.0+3233. This causes source confusion, something which is also not handled well by Lott et al. (2012). On the other hand, the

complete *Fermi*-LAT analysis correctly accounts for the fluxes from all surrounding sources.

There are two types of over-predictions; one type comes from the statistical noise of the model itself. This is what is shown in Fig. 9. Here, you see a normal spread of over and under-predictions around the median. The second type of over-predictions comes from cases with strong source confusion, i.e. a strong neighbouring source. It is hard to quantify the number of over-predictions, but as a general estimate, the latest *Fermi* catalogues have around 6658 gamma-ray sources in the sky. This is equal to one source every 6.2 square degrees. Hence, it is rather rare to have sources as close as 1 degree.

To handle source confusion better, in the future, we plan to explore a model which additionally incorporates these nearby source positions and fluxes, perhaps as a seventh layer on top of the current image format or overlaid on each layer as in Fig. 16.

3.6. Time bin search algorithm

The machine-learning-based TS estimation is central to our new adaptive light curve approach. We also require an algorithm that partitions the entire data range into possible time windows. Different methodological and numerical approaches could be considered here, independent of the machine-learning estimator. Our choice is depicted in Fig. 17 and consists of the following steps:

1. A time window (red interval) is searched with sub-time intervals (blue intervals) based on the time stamps of selected photon events t_n . Each interval uses the time window's lower bound (t_{init}) as its own lower bound. To speed up the search, only photons with a relatively large impact on the TS are considered for calculating the intermediate's time stamps t_n . These are photons that have an energy and angular proximity to the source above and below chosen thresholds respectively.⁴ The TS is predicted using the machine-learning estimator for each resulting time interval.
2. The time stamps whose intervals produce TS within a TS target range (here [4,9]) are within the green interval. The **last** largest time window within the green interval is then used as the end time $t_{opt,1}$ for the light curve. If the TS are below the target range, the time window gets extended, and further events are processed.
3. The search for more time bins continues repeatedly with a new time window from each t_{opt} onwards until all photons are parsed through. The last bin will, therefore, not necessarily have a TS above the threshold.

Since the search is performed chronologically, if one takes the first event's timestamp within a time window that produces a TS within the target range as t_{opt} , exceedingly short time bins tend to occur. If one runs these time bins through the *Fermi*-LAT analysis, they will be evaluated as sudden strong fluxes with sudden hard spectra. For this reason, we take the last timestamp with a TS that is still within the target range instead of the first. If the last timestamp of a time window still produces a TS within the target range, then the time window is increased (while keeping the same initial t_{init}) until timestamps which produce TS above the target range are found. This is depicted in Fig. 18.

The final time bins can then be fed through the *Fermipy* python API to produce the final light curve with adapted time-binning.

⁴ The thresholds are set by default to a minimum energy of 1 GeV and a maximum angular proximity of 1 degree. Most photons contributing to the source signal fall below this threshold. This works well for most sources. If the algorithm produces exceedingly high TS, the thresholds should be loosened. On the other hand, when the algorithm gets too slow, the thresholds can be made stricter to speed up the search. We are currently working on an updated search algorithm that does not require this threshold. It will be added to the project GitHub.




Fig. 17. Visualisation of the steps of the time bin search algorithm with a TS target range of [4,9].




Fig. 18. Visualisation of the steps of the time bin search algorithm in the case of an unfinished range of event timestamps whose TS are within the target range (here [4,9]).

4. Example light curves

In this section, we provide a few example light curves that we produced using *flashcurve* and compare them with the light curves based on the method of Lott et al. (2012).⁵ Four independent test sources were studied:

- TXS 0506+056, an IBL/LBL blazar which was identified as a source of high-energy neutrinos (Padovani et al., 2018; IceCube-Collaboration, 2018)
- CTA 102, a flat spectrum radio quasar
- MKN 421, an LBL, which is one of the closest and brightest blazars
- MKN 501, a blazar with relatively lower gamma-ray flux.

For the generation of the *flashcurve* light curves, we use the entire energy range between 100 MeV and 300 GeV. The light curves produced with the classic approach, on the contrary, are made with an energy cut-off based on the optimal minimal energy E_{min} calculated with the Lott et al. (2012) methodology.

Note that the adaptive binning light curves made with the Lott et al. (2012) method were produced with a constant relative flux error, as

⁵ Preliminarily provided by Narek Sahakyan (Sahakyan et al., 2024).

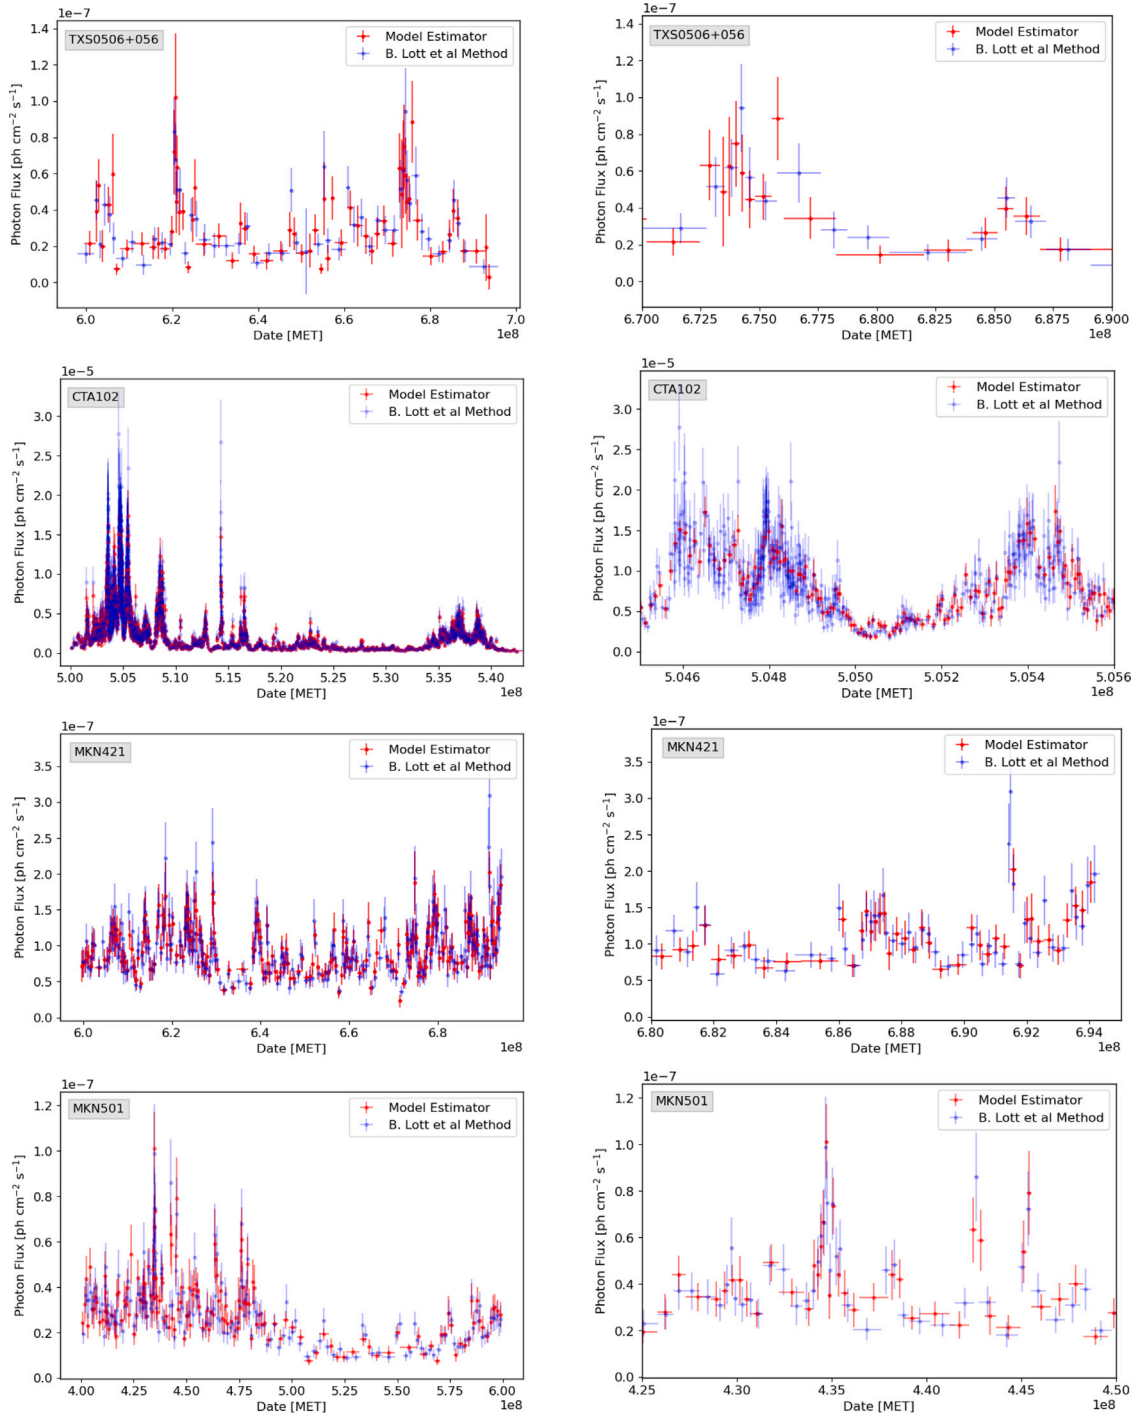


Fig. 19. Light curves of four test sources comparing *flashcurve* with the method of Lott et al. (2012). The four right plots are a zoomed-in section of interest of each source, with flux error and time bin length as vertical and horizontal errors, respectively.

the constant TS mode was not available at the time of this study.⁶ In order to achieve comparable results, the range of TS calculated from the constant relative flux error method was set as the thresholds for *flashcurve*. The details are listed in Table 3.

The results of our study can be seen in Figs. 19 to 22. At a general glance, *flashcurve* performs well, producing light curves with relatively constant significance time bins and without upper limits while also

Table 3

List of test sources used to make example light curves, along with the corresponding target range (Min. TS to Max. TS) for the time bin search algorithm.

Source	Min. TS	Max. TS
TXS 0506+056	25	50
CTA 102	50	75
MKN 421	115	170
MKN 501	75	125

⁶ https://www.slac.stanford.edu/~lott/tuto_adaptive_binning.pdf, last accessed: 01.07.2024.

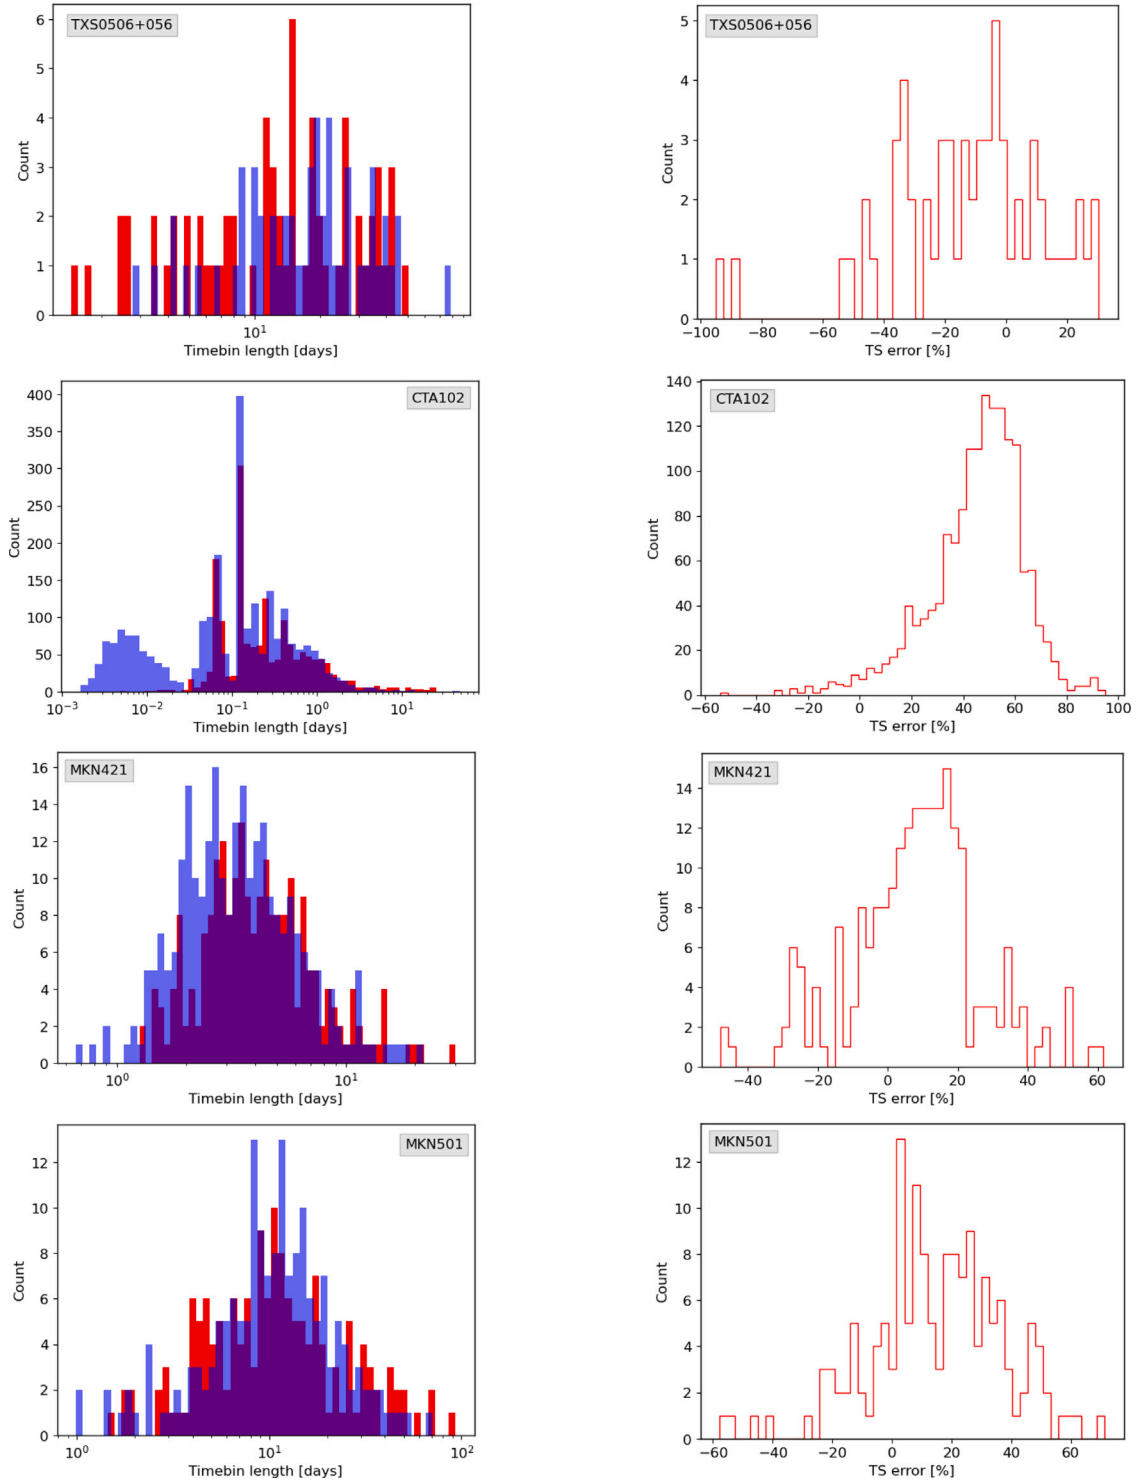


Fig. 20. *Left plots:* Histograms of the time bin lengths using *flashcurve* (red) and the method of Lott et al. (2012) (blue) for the four test sources. *Right plots:* Histograms of the relative TS deviation of the TS estimated using *flashcurve* from the TS calculated using the full *Fermi*-LAT likelihood analysis for the four test sources.

capturing the finer dynamics of the source flux. Moreover, we see reasonable agreement with the Lott et al. (2012) method.

Fig. 19 shows the light curves of each of the test sources. We observe a good agreement with the classic method for three of the selected sources. This is the case for TXS 0506, MKN 421, and MKN 501. Looking at the histograms of the time bin lengths in Fig. 20 and the spectral indices in Fig. 22 of these three sources, it is evident that the two methods produce similar results, except for a few short time bins producing overly hard spectra. This is due to *flashcurve*'s time bin

search algorithm, which only produces time bins that begin and end on photon event time stamps.

In the next step, we study the TS and relative flux errors of the light curves in Fig. 21 as well as the relative deviation of the estimated TS from the *Fermi*-LAT analysis in Fig. 20. Since *flashcurve*'s time bins for TXS 0506+056 were relatively shorter, this resulted in higher relative flux errors and lower TS than that of the Lott et al. (2012) method. Overall, however, we observe consistency between the two methods. We also note that *flashcurve* generally produces constant TS across

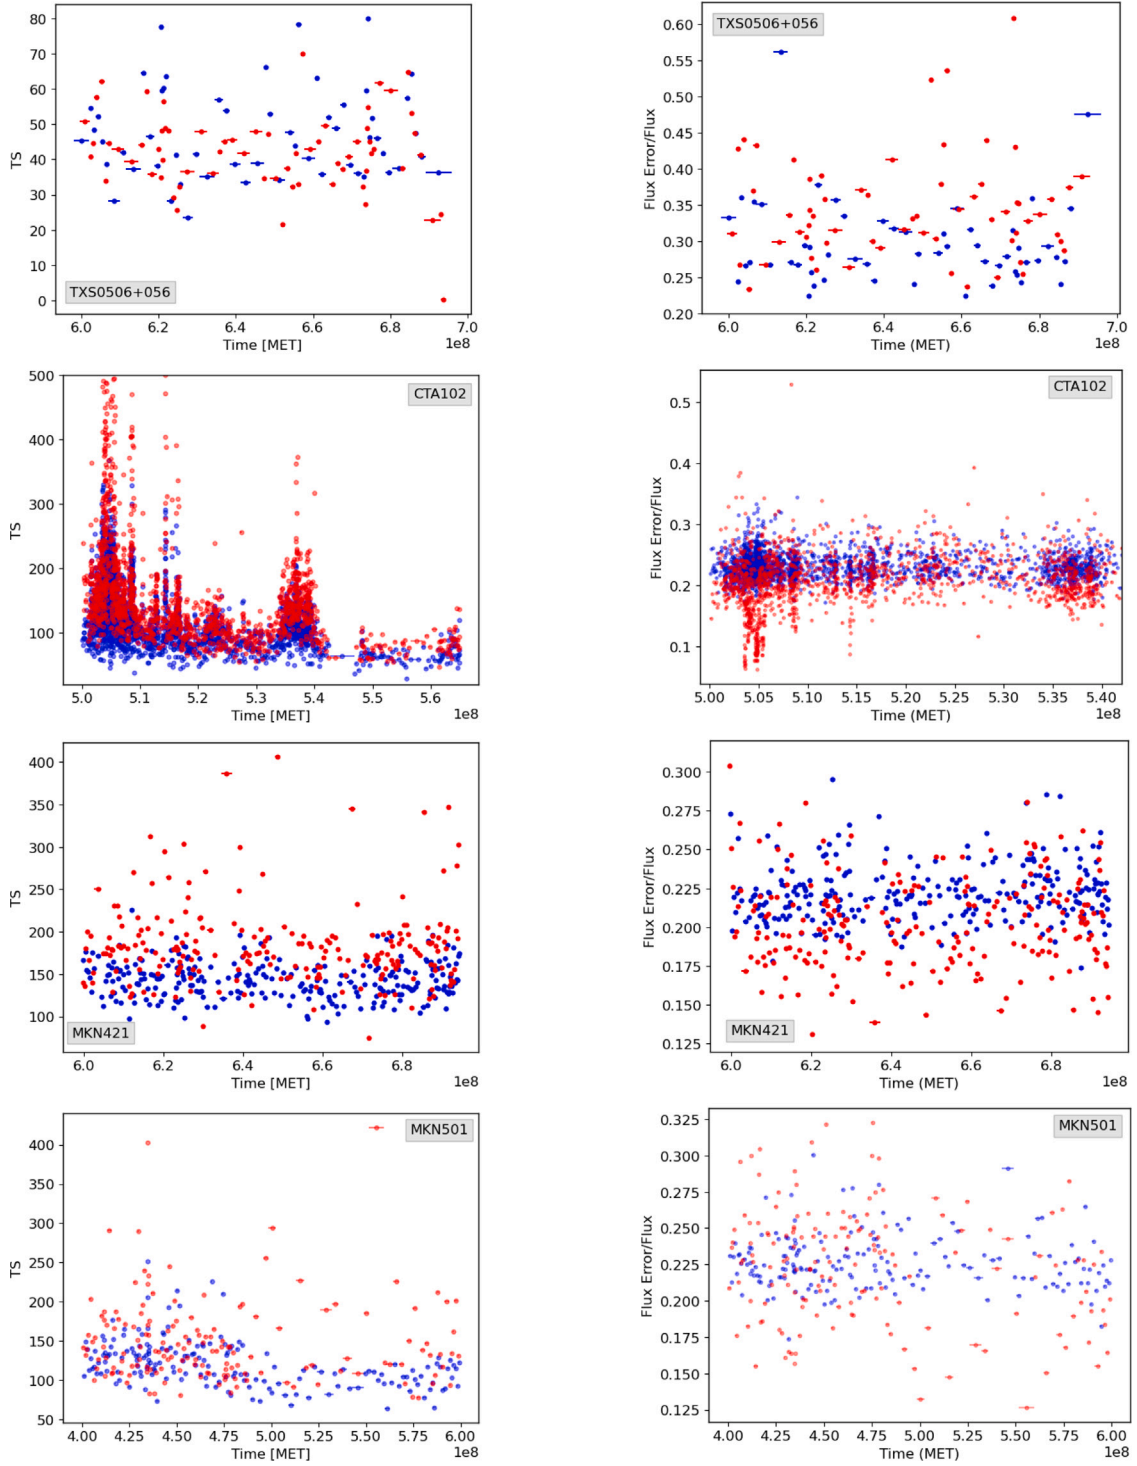


Fig. 21. Left plots: TS calculated using the *Fermi*-LAT analysis for the time bins found using our *flashcurve* (red) and the method of Lott et al. (2012) (blue) for the four test sources, with respective time bin lengths. In the case of CTA 102, a few extremely high TS are cut out for better visualisation. Right plots: Flux error versus flux ratios for the time bins found by the two methods.

all its time bins. Similarly, there is good consistency between the methods for MKN 421 and MKN 501. However, in this case, *flashcurve* tends to predict larger time bins, resulting in relatively larger TS with smaller flux error. This could be attributed to the energy and proximity thresholds of the time bin search algorithm, which filtered out intermediate photon events. Again, we see a relatively constant and decently accurate TS prediction overall.

The situation is different for CTA 102. Here, the two methods produce different results. CTA 102 was chosen due to its high activity

and hard fluxes (D'Ammando et al., 2019), which was meant to provide a challenging case for both light curve algorithms. Compared to the true TS values, it appears that *flashcurve* frequently under-predicts the TS values, as is evident in the TS error histogram in Fig. 20. This also results in small relative flux errors in the light curve, especially in its highly active flaring periods, since higher significance negatively correlates to lower relative flux error. In these periods, we also see larger variance in true TS; however, looking at relative TS error, we see a clear peak indicating that TS prediction, although collectively

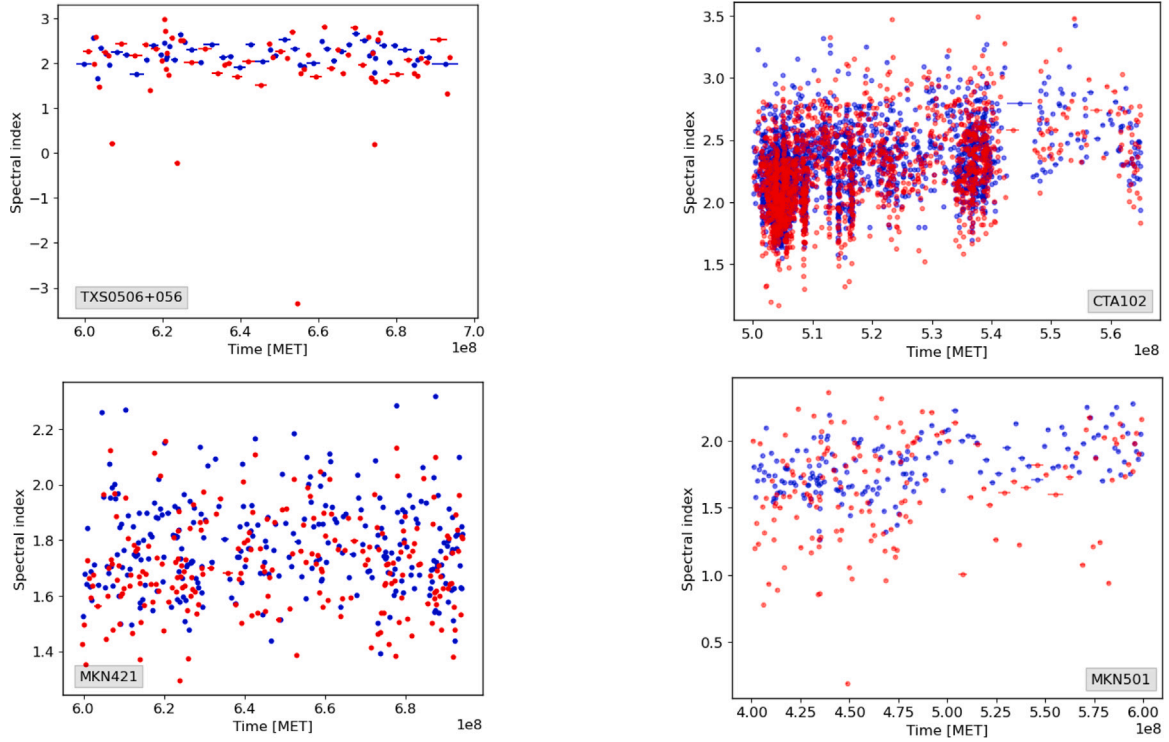


Fig. 22. Best-fit spectral indices for the time bins found using *flashcurve* (red) and the method of Lott et al. (2012) (blue) for the four test sources.

deviating from the target range, was consistent across all time bins. This, again, shows that the varied TS was more a fault of *flashcurve*'s time bin search algorithm rather than its TS prediction accuracy. Also, we see that the method of Lott et al. (2012) also displayed a relatively larger variance in TS and relative flux error during these periods, showing that these periods are generally difficult to process for either method.

Both *flashcurve* and the classic method have similar time bin length distributions. However, there is an additional population of extremely short time-bins, which was found by the Lott et al. (2012) method, as seen in CTA 102's time bin length distribution in Fig. 20. These shorter time bins do not necessarily indicate a better choice, as many of them come from the period of larger variance in flux error for the method of Lott et al. (2012). Using energy and proximity thresholds in *flashcurve*'s time bin search algorithm can cause it to predict fewer extremely short time bins.

In less active periods, we see more constant TS and similar flux dynamics relative to the classic method, as seen in the additional zoomed-in light curve of CTA 102 in Fig. 23.

We also see that, again, no upper limits were produced by the time bins predicted by *flashcurve*.

A final remark about the difference in the estimation of the time bins between the two methods is that it was clear that *flashcurve* is much faster than any previous method. The time scales for producing these time bins (not including the *Fermi*-LAT analysis after that) ranged from ~ 30 min to a few hours in the case of the very high activity sources. Using the method of Lott et al. (2012) can take up to several days, depending on the source strength. For an average source, it takes around 10 h to compute the time bins. This was mainly made possible using a neural network model that produces fast TS estimations. This can be sped up by multiprocessing, with around 100 threads used in these example cases. The main bottleneck of *flashcurve* is its time bin search algorithm, which can be replaced with a quicker and more sophisticated algorithm in the future. This might also improve the light curves by allowing the energy and proximity thresholds to be set lower in order to avoid missing shorter duration time bins, as in the case of CTA 102.

5. Conclusion and outlook

Using a machine-learning-based estimator model, we have presented a novel approach to adaptive binning of *Fermi*-LAT gamma-ray light curves. By leveraging neural networks, specifically convolutional neural networks, *flashcurve* estimates the test statistic for gamma-ray sources with respect to the pure background for gamma-ray events detected in a given time window. The method offers a significant improvement in computational speed over traditional methods like Lott et al. (2012). Light curves generated with *flashcurve* aim for a constant significance. This allows for fast variations in flux to be captured while at the same time avoiding upper limits in quiet periods. We validated *flashcurve* against multiple test gamma-ray sources. The results showed that for most sources, *flashcurve* produces high-quality light curves, which contain time bins with relatively constant TS, have no upper limits, and encompass smaller flux variations whenever appropriate. Results were also generally similar to that of the Lott et al. (2012) method, with only minor differences observed, particularly in highly active sources like CTA 102. Both methods, however, generally produced results with higher variance in relative flux error. In less active periods, however, *flashcurve*'s performance still demonstrated accurate results.

The training process, involving over 1.5 million time bins, utilised a residual block-based convolutional neural network architecture optimised with the ADAM optimiser. Despite challenges such as data imbalance towards lower TS values, the model for *flashcurve* generalised well on unseen test sets.

Validation against traditional methods showcased the potential of this approach in producing reliable light curves for gamma-ray sources at quick speeds, emphasising the importance of future improvements in model training and data handling techniques.

Future development of *flashcurve* will focus on several key areas. Firstly, expanding the dataset to include more sources, particularly those near the Galactic plane and extended sources, will help the generalisation capabilities of the estimator. Additional input layers could help to include information on the (known) sources in the region.

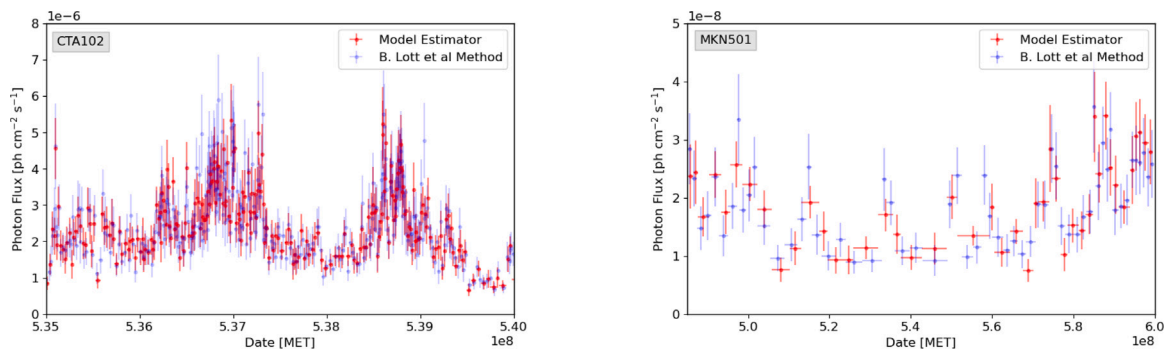


Fig. 23. Continuation of Fig. 19. Additional zoomed-in light curves.

Secondly, similar to the prediction of the TS values, the same approach can also be used to predict the flux uncertainty as an alternative method for generating adaptive binning light curves.

The model can be continuously improved by using *flashcurve*'s output as samples for the training. This way, the network could be trained in an infinite training loop for maximal accuracy, only limited by computational resources.

As for the time bin search algorithm, other numerical approaches, such as a bisection method (similar to the one described by Burden and Faires (1985)) could, in future, be implemented to help further speed up computation. This would also allow testing a more extensive range of time windows, reducing the variance in the light curve test statistics or flux uncertainties.

In general, using a machine-learning estimator to retrieve the time windows of adaptive-binning light curves is not limited to *Fermi*-LAT data. It could be similarly used to produce light curves in any other wavelength if sufficient training data is available.

CRedit authorship contribution statement

T. Glauch: Writing – review & editing, Writing – original draft, Supervision, Conceptualization. **K. Tchiorniy:** Writing – original draft, Software, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Martina Karl for proofreading the paper, Paolo Giommi and Narek Sahakyan for their consultation and provision of adaptive binning light curves, and Elisa Resconi for the green-lighting of this project. Funding for this research was provided through the SFB1258 grant. We also thank the unknown reviewers for reading the paper thoroughly and providing valuable comments.

Data availability

We have used data made public by the *Fermi*-LAT collaboration.

References

- Ajello, M., et al., 2015. The origin of the extragalactic Gamma-ray background and implications for dark-matter annihilation. *Astrophys. J. Lett.* 800 (2), L27. doi:10.1088/2041-8205/800/2/L27, arXiv:1501.05301.
- Ballet, J., Bruel, P., Burnett, T.H., Lott, B., collaboration, T.F.-L., 2024. Fermi large area telescope fourth source catalog data release 4 (4FGL-DR4). arXiv:2307.12546, URL: <https://arxiv.org/abs/2307.12546>.
- Ballet, J., Burnett, T.H., Digel, S.W., Lott, B., 2020. Fermi large area telescope fourth source catalog data release 2. arXiv:2005.11208, URL: <https://arxiv.org/abs/2005.11208>.
- Burden, R.L., Faires, D.J., 1985. *Numerical Analysis*, third ed. PWS Publishing Company, Boston.
- Caron, S., Gómez-Vargas, G.A., Hendriks, L., de Austri, R.R., 2018. Analyzing γ rays of the Galactic Center with deep learning. *J. Cosmol. Astropart. Phys.* 2018 (05), doi:10.1088/1475-7516/2018/05/058, 058–058.
- Chai, J., Zeng, H., Li, A., Ngai, E., 2021. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Mach. Learn. Appl.* 6, 100134. doi:10.1016/j.mlwa.2021.100134.
- D'Ammando, F., 2019. Relativistic jets in Gamma-ray-emitting narrow-line seyfert 1 galaxies. *Galaxies* 7 (4), 87. doi:10.3390/galaxies7040087, arXiv:1911.03500.
- D'Ammando, F., Raiteri, C.M., Villata, M., et al., 2019. Investigating the multiwavelength behaviour of the flat spectrum radio quasar CTA 102 during 2013–2017. *Mon. Not. R. Astron. Soc.* 490 (4), 5300–5316. doi:10.1093/mnras/stz2792.
- Fermi-LAT, 2009. Fermi/large area telescope bright gamma-ray source list. *Astrophys. J. Suppl. Ser.* 183 (1), 46–66. doi:10.1088/0067-0049/183/1/46.
- Fermi-LAT, 2013. Pass 8: Toward the full realization of the Fermi-LAT scientific potential. arXiv:1303.3514, URL: <https://arxiv.org/abs/1303.3514>.
- Fermi-LAT, 2020. Fermi large area telescope fourth source catalog. *Astrophys. J. Suppl. Ser.* 247 (1), 33. doi:10.3847/1538-4365/ab6bb6.
- Fermi-LAT, 2023. The Fermi-LAT lightcurve repository*. *Astrophys. J. Suppl. Ser.* 265 (2), 31. doi:10.3847/1538-4365/acbb6a.
- Fluke, C.J., Jacobs, C., 2020. Surveying the reach and maturity of machine learning and artificial intelligence in astronomy. *Wiley Interdiscip. Rev.: Data Min. Knowl. Discov.* 10 (2), e1349.
- Gaissner, T.K., Engel, R., Resconi, E., 2016. *Cosmic Rays and Particle Physics*, second ed. Cambridge University Press.
- Giommi, P., Glauch, T., Padovani, P., Resconi, E., Turcati, A., Chang, Y.L., 2020. Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection. *Mon. Not. R. Astron. Soc.* 497 (1), 865–878. doi:10.1093/mnras/staa2082.
- Glauch, T., Kersch, T., Giommi, P., 2022. BlaST - A machine-learning estimator for the synchrotron peak of blazars. *Astron. Comput.* 41, 100646. doi:10.1016/j.ascom.2022.100646, URL: <https://www.sciencedirect.com/science/article/pii/S2213133722000622>.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Deep residual learning for image recognition. arXiv:1512.03385, URL: <https://arxiv.org/abs/1512.03385>.
- IceCube-Collaboration, 2018. Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert. *Sci.* 361 (6398), 147–151. doi:10.1126/science.aat2890.
- James, F., Roos, M., 1975. Minuit: A system for function minimization and analysis of the parameter errors and correlations. *Comput. Phys. Comm.* 10, 343–367. doi:10.1016/0010-4655(75)90039-9.
- Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization. arXiv:1412.6980, URL: <https://arxiv.org/abs/1412.6980>.
- Kronmueller, M., Glauch, T., 2019. Application of deep neural networks to event type classification in IceCube. doi:10.48550/arXiv.1908.08763, arXiv e-prints, arXiv:1908.08763, arXiv:1908.08763.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nat.* 521 (7553), 436–444.
- Lott, B., Escande, L., Larsson, S., Ballet, J., 2012. An adaptive-binning method for generating constant-uncertainty/constant-significance light curves with Fermi-LAT data. *Astron. Astrophys.* 544, A6. doi:10.1051/0004-6361/201218873.

- Padovani, P., Giommi, P., Resconi, E., Glauch, T., Arsioli, B., Sahakyan, N., Huber, M., 2018. Dissecting the region around IceCube-170922A: the blazar TXS0506+056 as the first cosmic neutrino source. *Mon. Not. R. Astron. Soc.* 480 (1), 192–203. doi:10.1093/mnras/sty1852.
- Perez, L., Wang, J., 2017. The effectiveness of data augmentation in image classification using deep learning. *arXiv:1712.04621*, URL: <https://arxiv.org/abs/1712.04621>.
- Sahakyan, N., Vardanyan, V., Giommi, P., Bégué, D., Israyelyan, D., Harutyunyan, G., Manvelyan, M., Khachatryan, M., Dereli-Bégué, H., Gasparyan, S., 2024. Markarian Multiwavelength Data Center (MMDC): A tool for retrieving and modeling multi-temporal, multi-wavelength and multi-messenger data from blazar observations. *arXiv:2410.01207*, URL: <https://arxiv.org/abs/2410.01207>.
- Sbarrato, T., Foschini, L., Ghisellini, G., Tavecchio, F., 2011. Study of the variability of blazars gamma-ray emission. *Adv. Space Res.* 48 (6), 998–1003. doi:10.1016/j.asr.2011.05.019, URL: <https://www.sciencedirect.com/science/article/pii/S0273117711003565>.
- Sen, S., Agarwal, S., Chakraborty, P., Singh, K.P., 2022. Astronomical big data processing using machine learning: A comprehensive review. *Exp. Astron.* 53 (1), 1–43.
- Wilks, S.S., 1938. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Stat.* 9 (1), 60–62. doi:10.1214/aoms/1177732360.
- Zhang, X., Han, N., Zhang, J., 2024. Comparative analysis of VGG, ResNet, and GoogLeNet architectures evaluating performance, computational efficiency, and convergence rates. *Appl. Comput. Eng.* 44, 172–181. doi:10.54254/2755-2721/44/20230676.
- Zhao, X., Wang, L., Zhang, Y., Han, X., Deveci, M., Parmar, M., 2024. A review of convolutional neural networks in computer vision. *Artif. Intell. Rev.* 57 (4), 99. doi:10.1007/s10462-024-10721-6.