ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DEPARTMENT OF COMPUTER SCIENCE
AND ENGINEERING

ARTIFICIAL INTELLIGENCE

MASTER THESIS
in
Artificial Intelligence In Industry

AN INTERPRETABLE AND PROBABILISTIC APPROACH TO ACS COSI BACKGROUND MODELING

CANDIDATE	SUPERVISOR
Tommaso Perniola	Prof. Michele Lombardi
	CO-SUPERVISORS
	Dr. Nicolò Parmiggiani, Dr. Andrea Bulgarelli

Academic year 2024-2025
Session 5th

Abstract

Gamma-ray detectors aboard space missions operate in a background environment strongly influenced by spacecraft orientation and orbital position. The Anticoincidence System (ACS) of the Compton Spectrometer and Imager (COSI) satellite monitors the gamma-ray background and generates time series of detector counts with time bins of 50 ms. These measurements are used to identify transient events such as Gamma-Ray Bursts (GRBs) through onboard trigger algorithms. For this reason, accurate modeling of the background is essential in order to reliably distinguish transients from background fluctuations.

This thesis proposes a probabilistic and interpretable machine learning framework for modeling the ACS background signal. Instead of predicting a deterministic estimate, the approach models the full conditional distribution of detector counts given the spacecraft state parameters. Several probabilistic regression models were investigated, including Skew-Normal regression and Gaussian Mixture Regression, trained via maximum likelihood to capture asymmetric and regime-dependent variability in the detector counts.

To ensure interpretability, SHAP-based feature attribution methods were used to analyze the contribution of each input variable. The results indicate that background variations are primarily driven by spacecraft pointing directions and orbital geometry, consistent with physical effects such as Earth albedo and cosmic-ray induced background, and variations caused by spacecraft passage through the South Atlantic Anomaly (SAA). Kolmogorov–Arnold Networks were also explored as interpretable models capable of representing nonlinear relationships.

Finally, the probabilistic model was integrated into a transient detection strategy using adaptive significance thresholds derived from the distributions.

Experiments with simulated GRB signals demonstrate that the approach can identify significant excesses while accounting for background variability.

"And then whenever I stop and think—why did I want to do something? how did I work out how to do it?—I get a very strong desire just to stop thinking about it."

Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Gamma-ray astrophysics

High–energy astrophysics studies the most energetic phenomena in the Universe through the observation of X–ray and gamma–ray radiation. These photons are produced in extreme physical environments such as relativistic jets, compact objects, supernova explosions, and particle acceleration regions. Among these phenomena, *gamma-ray bursts* (GRBs) represent the most luminous transient events ever observed, releasing in a few seconds an amount of energy comparable to that emitted by the Sun over its entire lifetime.

GRBs were first discovered in the late 1960s by the *Vela* satellites [26], originally designed to detect nuclear explosions on Earth. The unexpected detection of intense flashes of gamma rays originating from outside the Solar System revealed the existence of a previously unknown astrophysical phenomenon. Since then, several missions such as BATSE [41], Swift [19], Fermi [40], and AGILE [45] have significantly improved our understanding of these events, enabling systematic observations of GRB temporal and spectral properties.

GRBs (Figure 1.1) are typically characterized by a short and intense *prompt emission* phase in gamma rays, lasting from milliseconds to several minutes, followed by a longer *afterglow* emission observable across multiple wavelengths, from X-rays to radio [15, 28].

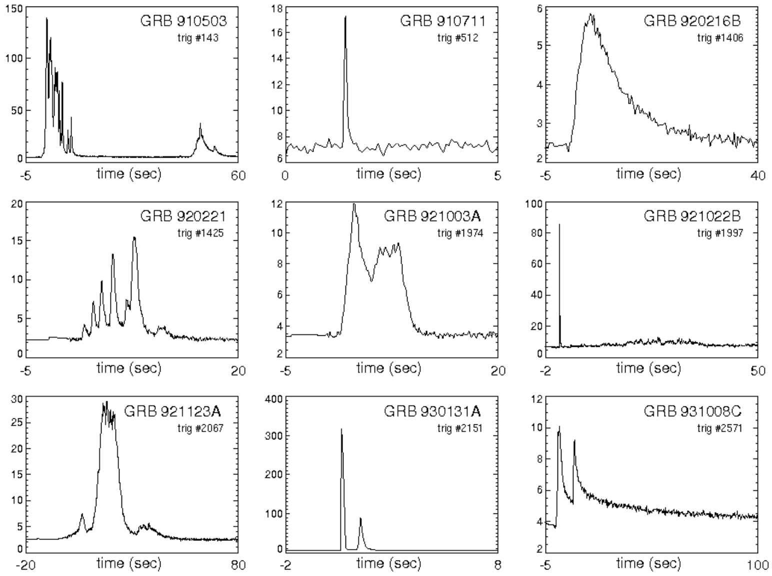

Figure 1.1: Examples of gamma-ray burst light curves detected by the BATSE instrument. The photon count rate is shown as a function of time for several bright GRBs, illustrating the large diversity in duration and temporal structure. Adapted from the public BATSE GRB catalog. Credit: Daniel Perley.

The detection of GRBs plays a crucial role in modern astrophysics. These events provide valuable information about relativistic particle acceleration, extreme magnetic fields, and the formation of compact objects.

Despite decades of observations, several aspects of GRBs remain poorly understood. In particular, the physical mechanisms responsible for the prompt emission and the detailed properties of the central engine are still under investigation. In addition, detecting faint or long–duration transients requires accurate background modeling of the photon counts recorded by space detectors. Background radiation in gamma-ray instruments originates from multiple sources, including cosmic rays, atmospheric albedo, and interactions of charged particles with the spacecraft structure. These contributions produce complex time-dependent count-rate variations that can mask faint astrophysical signals.

For this reason, accurate estimation of the detector background is a fundamental step for transient detection pipelines. Several recent works have proposed data-driven approaches for modeling the background count rates of space-based detectors using satellite orbital and attitude parameters. For instance, neural-network-based methods have been applied to GRB detection problems, including the analysis of gamma-ray intensity maps [35] and the identification of statistically significant deviations from the expected background signal [36].

In this context, developing robust background modeling techniques is essential for enabling reliable transient detection in modern gamma-ray missions.

1.2 The COSI Space Mission

The Compton Spectrometer and Imager (COSI; [46, 47]) is a NASA Small Explorer mission designed to observe the gamma-ray sky in the energy range between approximately 0.2 and 5 MeV. The mission aims to address several key scientific questions in high–energy astrophysics, including the origin of Galactic positrons, nucleosynthesis processes in the Galaxy, polarization of gamma-ray sources, and the detection of transient high-energy phenomena such as GRBs.

COSI operates as a compact Compton telescope based on an array of high-purity germanium detectors capable of reconstructing the interaction sequence of incoming gamma-ray photons through Compton scattering (Figure 1.2). This design allows the instrument to achieve high spectral resolution while maintaining a large field of view and broad sky coverage. In addition to the primary detector system, the spacecraft is equipped with an Anti-Coincidence Shield (ACS) designed to suppress charged particle background. The ACS is composed of Bismuth Germanate (BGO) scintillator detectors surrounding the main instrument, with panels located laterally and on the bottom side of

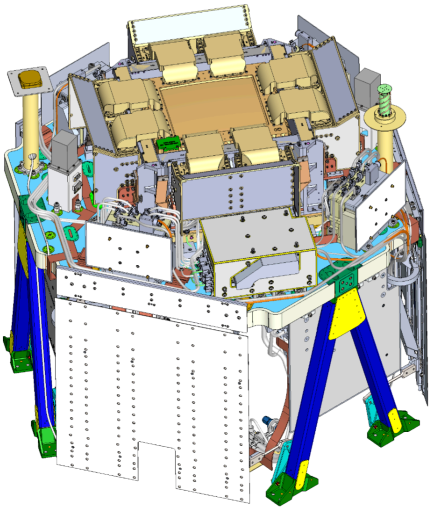

Figure 1.2: COSI spacecraft overview, adapted from [13].

the spacecraft, and continuously records count rates generated by particle interactions. While primarily designed for background rejection, such detectors are also sensitive to high-energy transient events and can therefore be used to detect sudden increases in photon flux associated with gamma-ray bursts [37].

The large field of view and continuous monitoring capabilities make instruments such as the COSI ACS particularly suitable for detecting transient gamma-ray phenomena. However, the background count rate recorded by these detectors strongly depends on the spacecraft orbital position, and orientation. Consequently, accurate modeling of the background signal is required in order to identify statistically significant deviations associated with astrophysical events [37].

1.2.1 Background physical drivers

The background radiation observed by gamma-ray detectors originates from several physical components [16]. One of the main contributions arises from atmospheric gamma-ray albedo, generated when cosmic rays interact with the Earth's atmosphere and produce secondary gamma-ray photons [33]. As a consequence, the detected count rate can increase when the detector field of view faces Earth.

Another important component is the flux of cosmic rays, which can interact with the spacecraft and detector materials, producing secondary particles and activation effects that contribute to the measured count rates. In addition, charged particles trapped in the Earth's radiation environment, particularly within the Van Allen radiation belts (Figure 1.3; [48]), can generate enhanced background levels along specific regions of the spacecraft orbit.

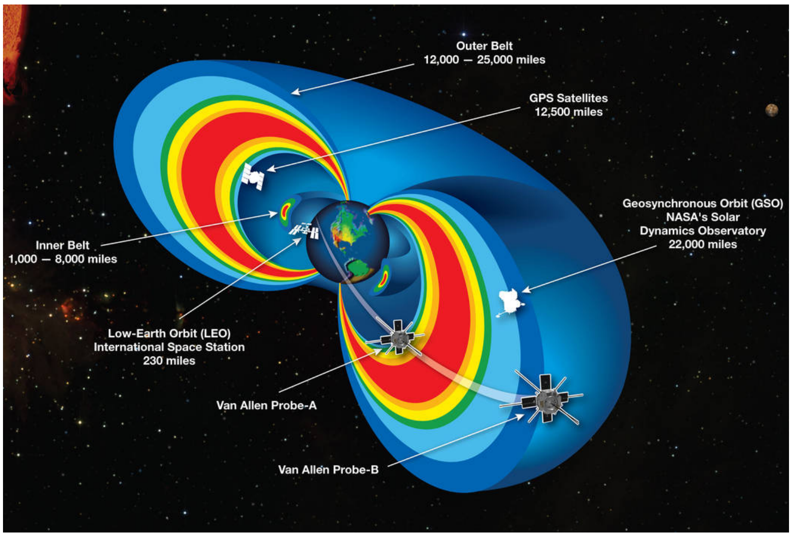

Figure 1.3: Schematic illustration of the Earth's Van Allen radiation belts, showing the two toroidal regions of trapped charged particles surrounding the planet. The approximate locations of several spacecraft in Low Earth Orbit, including the International Space Station, are indicated. Credit: NASA.

Because the intensity of these components varies with the spacecraft position and pointing direction, the observed detector count rates exhibit complex time-dependent variations. These dependencies motivate the use of spacecraft orbital and attitude parameters as inputs for data-driven background modeling approaches.

1.2.2 COSI ACS geometry

The COSI spacecraft reference frame is defined by three orthogonal axes ($+X$, $+Y$, $+Z$) that describe the orientation of the instrument in space. The $+Z$ axis

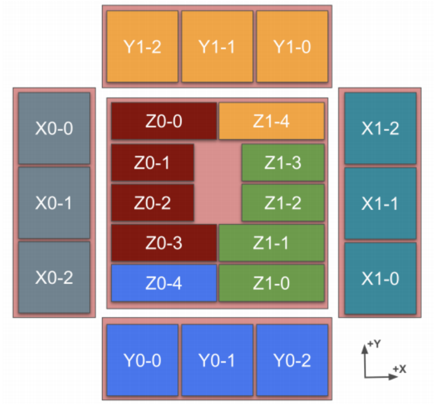

Figure 1.4: COSI ACS spacecraft geometry (a drawing from [13]).

corresponds to the main instrument boresight and defines the nominal pointing direction of the telescope. The spacecraft attitude is described through a Cartesian reference frame defined by the X, Y, and Z axes of the payload. This three-dimensional geometry determines how the ACS detector panels, oriented along different spacecraft axes, interact with the surrounding radiation environment. The orbital metadata provide the pointing direction of these axes, allowing the reconstruction of the spacecraft orientation and of the detector exposure to the background.

The geometry of the spacecraft plays a key role in shaping the detector response. Different detector elements are exposed to varying levels of radiation depending on their orientation relative to the Earth, the cosmic gamma-ray background, and charged particle fluxes. As a result, the measured background count rates exhibit strong correlations with the spacecraft attitude and orbital parameters. Understanding this geometrical configuration is therefore essential when modeling the detector background and interpreting the observed photon counts.

1.3 Gamma-ray Transient Event Detection

The detection of transient high-energy phenomena can be formulated as the problem of identifying statistically significant signals within noisy observational data. Gamma-ray detectors continuously measure photon counts that include contributions from both astrophysical sources and background radiation. A transient event such as a GRB manifests itself as a sudden excess of counts over the expected background level.

Traditional detection pipelines rely on threshold-based trigger algorithms applied directly to the observed count rates. These methods typically compare the measured counts with a background estimate obtained through simple time-window averaging. Classical statistical techniques widely used in gamma-ray astronomy include likelihood-based significance estimators such as the Li & Ma test [29] and segmentation approaches such as Bayesian Blocks [42, 13]. While widely used in gamma-ray astronomy, these methods typically rely on an estimate of the background obtained from the observed data itself.

Recent studies have shown that machine learning techniques can improve transient detection performance by exploiting patterns in detector measurements. For instance, convolutional neural networks have been applied to classify gamma-ray intensity maps and identify candidate GRB signals in AGILE data [45], achieving improved detection capability compared to traditional significance-based methods [35]. More generally, machine learning approaches can also be used to model background variability. In these frameworks, a predictive model is trained on background observations to estimate the expected detector response over time, enabling the identification of statistically significant upward fluctuations corresponding to candidate transient events [36].

Within this framework, GRB detection can be formulated as an anomaly detection problem in which observed count rates are compared with the predicted background distribution. Events that significantly deviate from the expected background are flagged as candidate transients and subsequently analyzed to determine their astrophysical origin.

Accurate background modeling therefore represents a crucial component of modern gamma-ray transient detection pipelines and provides the foundation for identifying anomalous astrophysical signals in space-based observations.

Chapter 2

Data and Background Characterization

2.1 COSI data and Preprocessing

2.1.1 Data Source

For this work we used the data simulated by the COSI team for the Data Challenge 3 [14]. The dataset includes simulated spacecraft orbital and attitude parameters over a continuous 24-hour period, along with the corresponding ACS shield background counts.

The background signal is generated using the spacecraft geometry and orbital configuration, allowing a controlled evaluation of probabilistic background models.

2.1.2 Data Inspection

Feature Explanation

Our input features directly describe the spacecraft geometry and orbital behavior. As shown in Figure 2.1, the time evolution of these variables reflects both the orbital motion and the spacecraft pointing strategy.

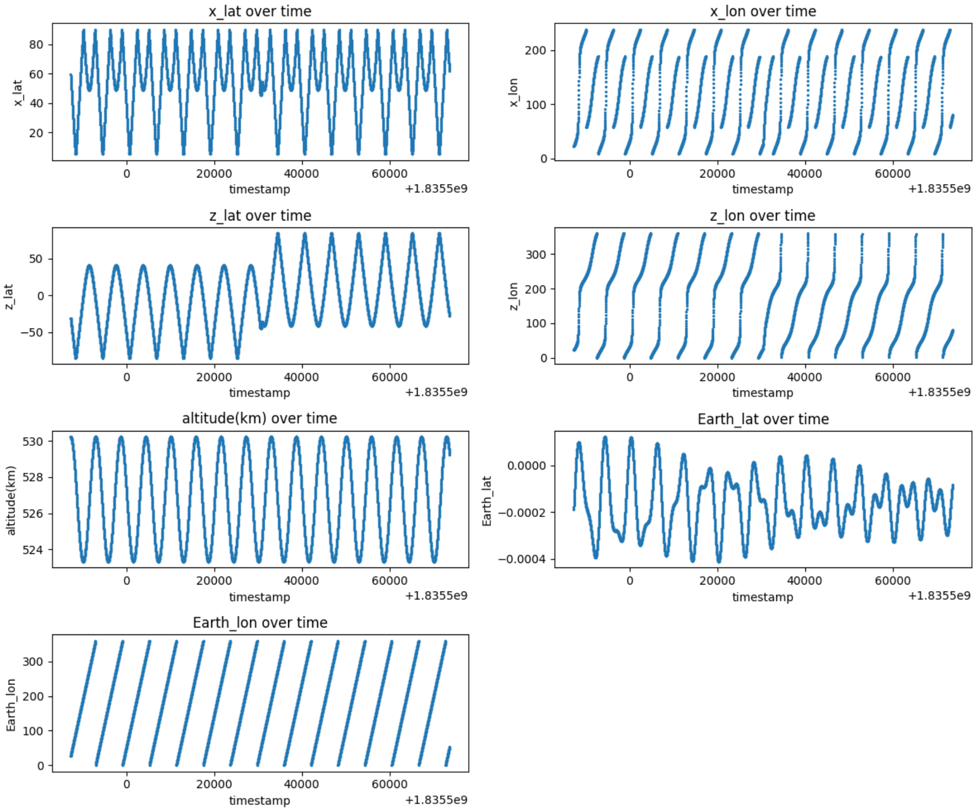

Figure 2.1: COSI ACS satellite dataframe input features over 24 hours of data.

In the spacecraft body-fixed reference frame (Figure 1.4), the attitude is defined by three orthogonal axes: $+X$, $+Y$, and $+Z$. The $+Z$ axis corresponds to the instrument boresight and therefore defines the nominal pointing direction of the spacecraft. Its orientation in the chosen galactic reference frame (a coordinate system in which longitude and latitude are measured with respect to the plane of the Milky Way and the direction of the Galactic center) is described by the angular coordinates $\mathrm{z}_{\mathrm{lon}}$ and $\mathrm{z}_{\mathrm{lat}}$.

The $+X$ axis provides complementary information about the spacecraft orientation, capturing rotations around the boresight direction. Its orientation is encoded through $\mathrm{x}_{\mathrm{lon}}$ and $\mathrm{x}_{\mathrm{lat}}$.

In this work, the orientations of the $+Z$ and $+X$ axes are used as inputs, since two orthogonal axes are sufficient to determine the spacecraft attitude.

In addition to pointing information, the feature set includes orbital parameters. The variable $\mathrm{Earth}_{\mathrm{lon}}$ describes the geographic longitude of the spacecraft position projected onto the Earth's surface, effectively encoding the orbital phase of the spacecraft. $\mathrm{Earth}_{\mathrm{lat}}$ represents the corresponding latitude, which in our data remains close to zero due to the nearly equatorial orbit. The altitude variable captures small variations in orbital height (approximately 523–530 km).

As visible in Figure 2.1, all angular variables exhibit clear periodic patterns over the 24-hour window shown in the figure, reflecting the underlying spacecraft orbital motion and the evolution of its pointing direction. These geometric and orbital descriptors constitute the physical drivers of the detector-dependent background modeled in this work.

2.1.3 Interpolation

The spacecraft orientation metadata are originally provided with a fixed time resolution of 15 s, while the ACS detector counts are available at higher temporal resolution, down to 50 ms. To apply the method at these finer time scales, the satellite metadata are linearly interpolated to match the desired resolutions, including 1 s and 50 ms bins.

2.1.4 Feature Transformation

From Figure 2.1 it is evident that the longitudinal and latitudinal variables exhibit a clear periodic behavior over time. This periodicity is directly related to the orbital motion of the spacecraft and to its pointing evolution.

Looking at the numerical ranges makes this even clearer:

- $\mathrm{Earth}_{\mathrm{lon}}$ and $\mathrm{z}_{\mathrm{lon}}$ span the interval $[0°, 360°]$.
- $\mathrm{x}_{\mathrm{lon}}$ is defined on the same circular domain, although only a sub-range is observed in the available data.

- $\mathrm{z_{lat}}$ lies approximately in $[-85°, 85°]$.
- $\mathrm{x_{lat}}$ lies approximately in $[5°, 90°]$.

Although these quantities are angular by nature, their raw numerical representation is linear. As a consequence, directions that are physically identical (e.g. $0°$ and $360°$) appear numerically far apart. This creates artificial discontinuities at the angular boundaries and may force the model to learn spurious edge effects that have no physical meaning.

To address this issue, angular variables were encoded using sine and cosine transformations. Longitudes were mapped as

$$\theta \to (\sin\theta, \cos\theta),$$

while latitudes were encoded using $\sin(\theta)$ only, since they are bounded within $[-90°, 90°]$ and do not wrap around.

After transformation, all encoded angular features lie in the interval $[-1, 1]$, as shown in Figure 2.2. This representation preserves the circular geometry of angular variables and removes artificial jumps at the $0°/360°$ boundary.

Although the sine/cosine encoding slightly increases the input dimensionality, it effectively constrains the model by embedding prior geometric knowledge directly into the feature space. The network no longer needs to allocate capacity to learn periodicity from scratch, and it avoids modeling artificial discontinuities.

In conclusion, the feature transformation is not only a numerical stabilization step, but also a design choice aligned with the periodic structure of the underlying physical variables.

2.1.5 Target Feature

The target features correspond to the photon count rates recorded by the six BGO detector units of the COSI ACS.

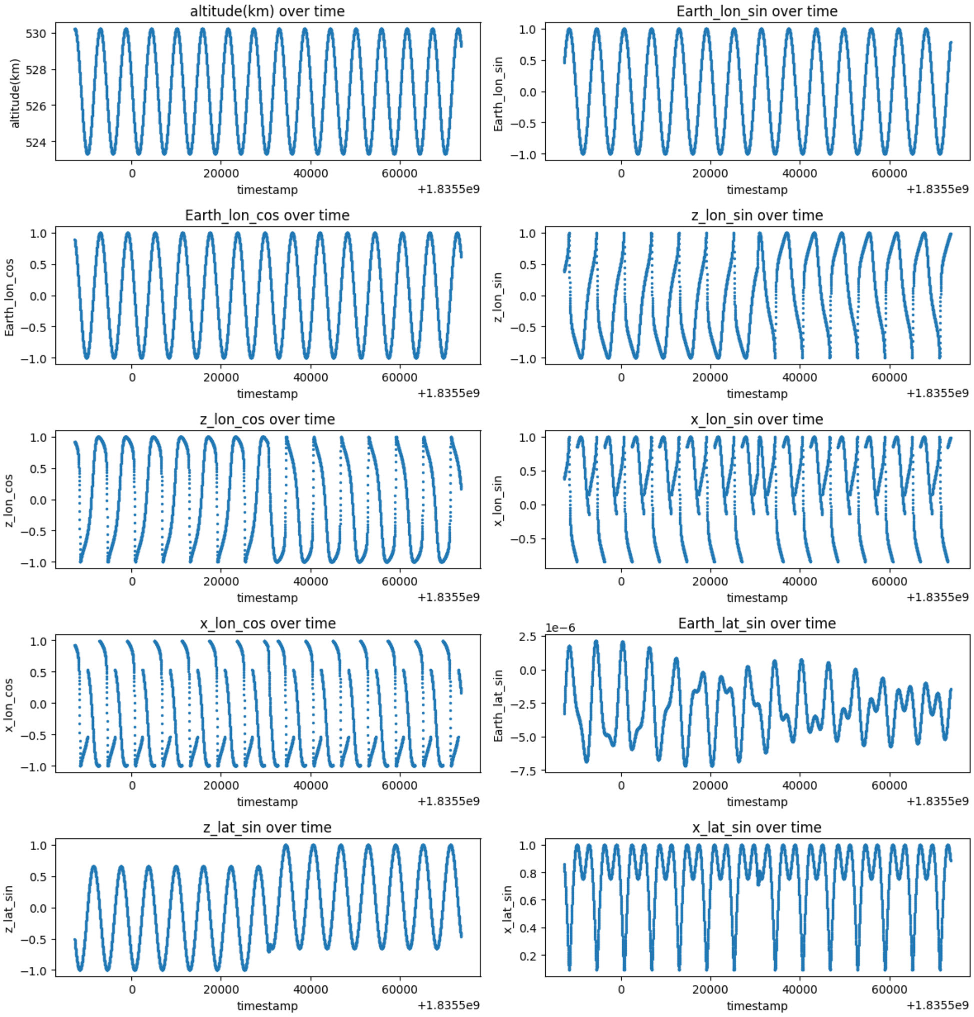

Figure 2.2: COSI ACS satellite dataframe input features after sine/cosine encoding, over 24 hours of data.

Figure 2.3 shows the 15 s binned light curves over the full 24-hour interval. All detectors exhibit a common large-scale modulation across the orbit, while their absolute count levels differ.

The Z-oriented panels (bgo_{z1} and bgo_{z0}) tend to display higher count rates, the X-oriented units intermediate levels, and the Y-oriented panels lower average counts.

A zoom over a representative one-hour segment (Figure 2.4) highlights the non-stationary nature of the background. Smooth orbital-scale modulations are superimposed with stochastic fluctuations at shorter time scales.

Figure 2.3: Count rates at 15 s time resolution over the full 24-hour interval for each detector.

Overall, the coexistence of shared orbital trends and detector-specific count levels motivates a conditional, per-detector probabilistic modeling approach rather than a single global background estimate.

2.1.6 Stabilization Phase and Rocking Motion

Only data acquired during the second half of the simulated day (starting from hour 13) are considered for this work. During the first 12 hours, the spacecraft undergoes an initial stabilization phase in which the detector response is not yet fully settled. As a result, the recorded background signal does not reflect steady-state observational conditions. This effect is clearly evident in Figure 2.3, where the count rate shows a gradual increase during the first hours of the simulation before reaching a stable regime after approximately 12 hours.

In addition, the first few orbits must be discarded because the background level is still evolving due to the activation of spacecraft materials by cosmic

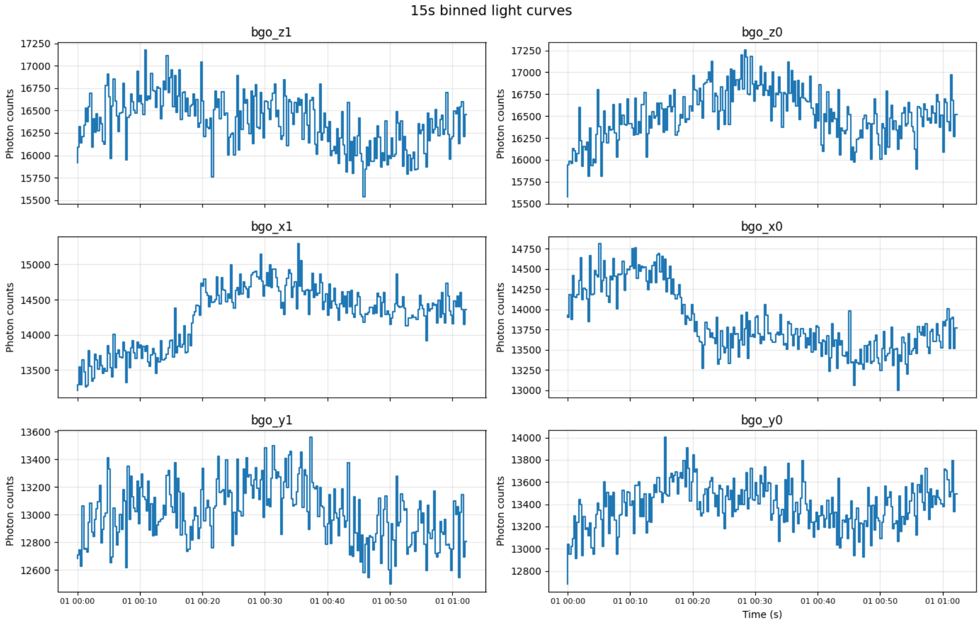

Figure 2.4: COSI ACS background light curves at 15 s resolution over an interval of approximately 62 min for each detector.

rays. This process produces short-lived radioactive isotopes within the detector and surrounding structures, whose decay temporarily increases the observed count rate until the activation-induced component reaches equilibrium. Discarding these initial orbits ensures that the analyzed data correspond to a stable background regime representative of nominal observational conditions.

Furthermore, the spacecraft performs a periodic rocking motion with a 12-hour period to systematically vary its pointing direction and achieve full-sky coverage. This scanning strategy introduces structured variations in the attitude parameters, which are clearly visible in the time evolution of angular features such as z_{lat} in Figure 2.1.

However, given the limited duration of the currently available simulation, it was not possible to perform a comprehensive analysis that simultaneously accounts for both the periodic rocking motion and the two distinct phases of the observation (the initial stabilization period and the subsequent steady-state regime).

Future simulations produced by the COSI team are expected to provide a significantly longer dataset, spanning approximately seven days of continuous observations. The availability of this extended dataset will allow a more complete characterization of the spacecraft rocking cycle and a more robust analysis of the background variability.

2.1.7 South-Atlantic Anomaly

Further inspection of the dataset reveals a non-negligible fraction of zero-count bins. This behavior is associated with the passage of the spacecraft through the South Atlantic Anomaly (SAA), a region over the South Atlantic Ocean where the Earth's magnetic field is significantly weaker [23].

Since COSI operates in Low Earth Orbit (LEO), the spacecraft periodically traverses this region during its orbit. In this area, the inner Van Allen radiation belt approaches closest to the Earth's surface, resulting in an enhanced flux of charged particles at satellite altitude. To prevent damage and avoid contamination from high radiation levels, scientific instruments are typically switched off or placed in a safe configuration during SAA crossings.

As a consequence, the ACS records zero count rates within the corresponding temporal windows. These zero-valued intervals therefore reflect operational constraints rather than physical background conditions.

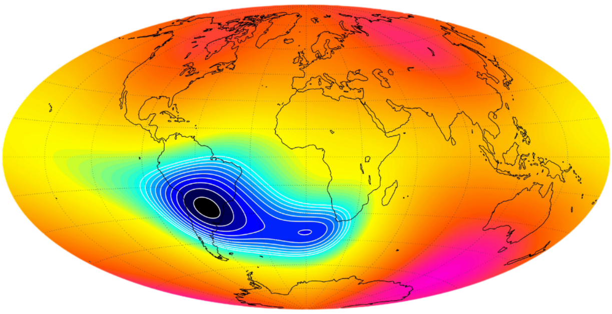

Figure 2.5: Geographical extent of the South Atlantic Anomaly. The shaded region indicates reduced geomagnetic field strength, leading to enhanced radiation exposure for satellites in Low Earth Orbit.

Because of that, we can safely delete those rows which contain zero values in their count rates columns.

2.1.8 Outlier Removal

To mitigate the impact of extreme values, a quantile-based filtering strategy was applied to all numerical features. For each variable, lower and upper bounds were computed using empirical quantiles at the 0.4th and 99.9th percentiles, respectively.

Observations falling outside this interval in any numerical column were removed from the dataset. Formally, a sample was retained only if all its numerical features satisfied

$$q_{0.004} \leq x \leq q_{0.999},$$

where q_p denotes the empirical p-quantile of the corresponding variable. This procedure removes extreme outliers while preserving the vast majority of the data distribution. The filtering is applied jointly across features to ensure consistency across correlated variables.

2.2 Distributional Analysis

In a photon counting experiment, the observed number of events in a fixed time bin is naturally described by a Poisson process [11]. Formally, for a given time bin t,

$$C_t \sim \mathrm{Poisson}(\lambda_t),$$

where C_t denotes the observed count and λ_t is the expected number of photons in that bin.

It is important to note that the Poisson mean parameter λ_t cannot be reliably estimated from a single observation C_t. This is a fundamental statistical limitation and it would represent a critical issue if the goal were to infer the Poisson parameter from isolated bins.

However, this is not the objective of our analysis. Here, we study the population-level distribution of per-bin counts, implicitly conditioning on orbital configurations.

The dataset contains a large number of bins acquired under comparable physical conditions, corresponding to similar background regimes. By aggregating bins associated with approximately stationary system states, we can empirically inspect the distribution of photon counts characterizing a given background regime. In this sense, the statistical information is not extracted from repeated measurements of the same bin, but from multiple observations collected under similar conditions over time.

This perspective allows us to meaningfully analyze the per-bin count distribution without violating the single-sample limitation of Poisson inference. The Poisson distribution is thus interpreted as a generative model for individual bins, while the empirical distribution arises from the collection of many such realizations under similar conditions.

Conditioning on the system state x_t, which encodes the relevant physical and orbital features, this relationship can be formalized as

$$P(C_t \mid x_t) = \mathrm{Poisson}(\lambda = f_\theta(x_t)),$$

where the function f_θ models the expected count rate associated with a given system configuration.

Chapter 3

Probabilistic Background Modeling

3.1 Regression Framework

We now move from distributional analysis to model estimation. The goal is no longer to inspect the empirical count distribution, but to learn a mapping from satellite state features to the parameters of a chosen parametric family defining the predictive distribution.

More generally, we assume that the conditional distribution of the counts belongs to a parametric family $\mathcal{F}$:

$$P(C_t \mid x_t) = P_{\mathcal{F}}(C_t \mid \theta_t(x_t)), \tag{3.1}$$

where $\theta_t(x_t)$ denotes the distribution parameters predicted from the input features.

In Section 2.2, the Poisson model was introduced as a first generative assumption:

$$P(C_t \mid x_t) = \mathrm{Poisson}(\lambda = f_\theta(x_t)). \tag{3.2}$$

In the following sections, we extend this framework to more flexible distributional families.

3.1.1 Train–Validation–Test Split

To evaluate the predictive models while preventing information leakage, the dataset was partitioned into training, validation, and test subsets.

Since the data form a continuous time series along the spacecraft orbit, the final 15% of the time-ordered dataset was reserved as a *contiguous* test set. This preserves the natural temporal correlations of the observations and allows model evaluation on complete orbital segments, enabling the inspection of the temporal behavior of the predicted background along full orbital light curves.

The dataset was randomly split into training and test sets, with 10% of the data reserved for testing. From the training set, 10% was further held out as a validation set for model selection and hyperparameter tuning.

Orbital temporal structure. In the temporal prediction analysis, we extract the last contiguous orbital segment of the dataset, corresponding to approximately 80 min of observations. This segment is used in Subsections 3.3.3, 3.4.6 and 5.3.2 to analyze the temporal behavior of the models on continuous data. Using a contiguous time interval allows us to evaluate how the models track the evolution of the background signal over time, rather than across randomly sampled bins. This choice is supported by the autocorrelation analysis reported in Appendix B, Section B.1, where the detector count rates exhibit a dominant periodicity consistent with the spacecraft orbital timescale.

3.1.2 Training Procedure

Once the conditional distributional family is specified, the parameters of the predictive model are estimated by maximizing the data likelihood.

In practice, the distribution parameters $\theta_t(x_t)$ are produced by a neural network (or probabilistic regression model) with learnable weights $\mathbf{w}$.

Given observed counts $C_{t,d}$ for detector d at time bin t, the model predicts the corresponding distribution parameters $\theta_{t,d}(x_t)$. Training is performed by minimizing the negative log-likelihood (NLL) over all time bins and detectors:

$$\mathcal{L} = -\sum_{t=1}^{T} \sum_{d=1}^{D} \log P_{\mathcal{F}}\Big(C_{t,d} \mid \theta_{t,d}(x_t)\Big). \tag{3.3}$$

This objective encourages the model to assign high probability to the observed counts while properly calibrating predictive uncertainty. Unlike deterministic regression, which optimizes only point estimates, likelihood-based training optimizes the full predictive distribution. This is crucial in the presence of heteroscedasticity and regime-dependent variability, as it allows the model to adapt both expected counts and uncertainty to the orbital and geometric configuration of the spacecraft.

3.1.3 Model Performance Metrics

Model performance is evaluated using the following metrics:

- Mean Absolute Error (*MAE*),
- Root Mean Squared Error (*RMSE*),
- Mean Absolute Percentage Error (*MAPE*),
- Pearson dispersion coefficient ϕ (for the Poisson model only),
- Negative Log-Likelihood (*NLL*) on the test set.

Custom MAE. The predicted mean count rate for each detector is extracted from the model output and compared with the observed counts, computing the average absolute error across all detectors and time bins.

Percentage error. To provide a scale-independent measure of the prediction accuracy, we compute the mean absolute percentage error:

$$\text{MAPE} = \frac{1}{N} \sum_{i=1}^{N} \left| \frac{y_i - \hat{y}_i}{y_i} \right| . \tag{3.4}$$

where y_i denotes the observed count and $\hat{y}_i$ the predicted mean for a given detector and time bin.[1]

---

[1] As discussed in Section 2.1.7, zero-count bins occur only during passages through the South Atlantic Anomaly, which are excluded from the analysis. Therefore, no additional numerical constant is required in the denominator.

3.2 Baseline Poisson Model

3.2.1 Model Specification

As discussed in Section 2.2, the Poisson distribution provides a natural generative model for per-bin photon counts. We now adopt this assumption as a regression baseline, modeling the conditional distribution of the counts as

$$C_t \mid x_t \sim \text{Poisson}(\lambda_t), \tag{3.5}$$

where the rate parameter is predicted as a function of the satellite state:

$$\lambda_t = f_\theta(x_t). \tag{3.6}$$

Under this model, the variance is constrained to equal the mean:

$$\mathrm{Var}(C_t \mid x_t) = E[C_t \mid x_t] = \lambda_t. \tag{3.7}$$

This assumption implies a single stationary regime per time bin and a fixed dispersion structure, which will be empirically evaluated in the following subsection.

3.2.2 Empirical Evidence of Overdispersion

Our simulated background data, however, do not satisfy the assumption of a single stationary Poisson regime. To assess the validity of the Poisson hypothesis, we compute the local mean and variance of the detector counts over sliding 10 s windows and analyze their relationship.

The use of a short temporal window allows us to estimate the dispersion properties of the signal under approximately stationary conditions, avoiding the mixing of different orbital configurations while reducing the impact of short-term statistical fluctuations.

As shown in Figure 3.1, the empirical variance is consistently larger than

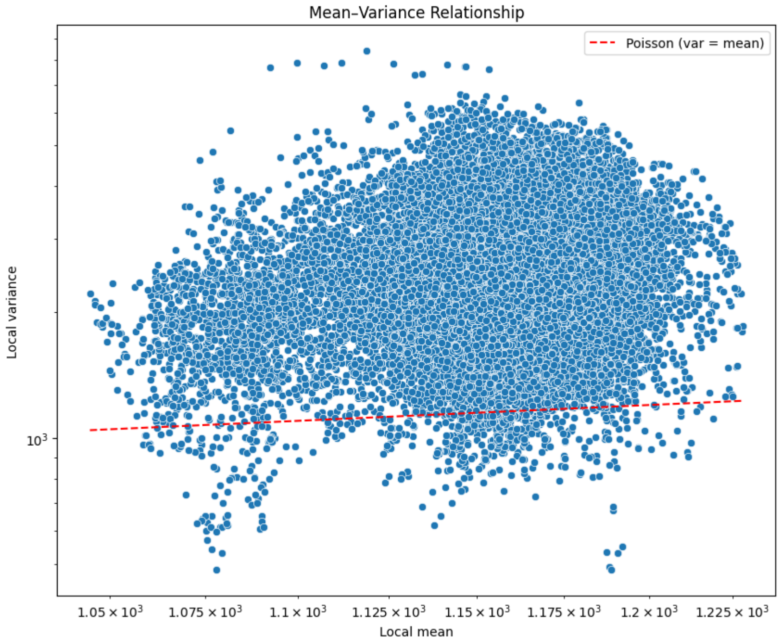

Figure 3.1: Mean–variance relationship for detector bgo_{z1} at 1 s time resolution. The dashed diagonal line represents the Poisson expectation $\mathrm{Var} = \mathrm{Mean}$. The observed variance systematically exceeds the theoretical prediction, indicating overdispersion.

the mean, particularly around the typical background count level. This systematic deviation from the Poisson identity $\mathrm{Var} = \mathrm{Mean}$ provides clear evidence of overdispersion in the detectors [10].

Such behavior reflects the presence of multiple background regimes driven by variations in spacecraft attitude, orbital position, and environmental conditions [16]. As a result, the background process cannot be described by a single stationary Poisson model and requires more flexible probabilistic formulations.

3.2.3 Limitations of the Global Poisson Assumption

The observed overdispersion has direct modeling implications. A global Poisson model assumes a single homogeneous regime and therefore underestimates the true uncertainty of the counts. When variance exceeds the mean, the model becomes overly confident, leading to biased estimates of λ_t and miscalibrated uncertainty bands.

In other words, the discrepancy is not merely statistical noise but structural: the data exhibit variability that cannot be explained by a single-parameter Poisson distribution.

This motivates the exploration of more flexible probabilistic models capable of describing heterogeneous variability in the detectors counts.

3.2.4 Baseline Poisson Regression Results

To quantify the practical impact of this mismatch, we train a conditional Poisson regression model that predicts λ_t as a function of the satellite orientation features using the 1 s resolution dataset.

The model is optimized by minimizing the negative log-likelihood (3.1.2), where the table 3.1 reports the predictive performance on the test set.

Detector	MAE	RMSE	Pearson ϕ
bgo_z1	49.55	62.08	3.366
bgo_z0	49.89	62.78	3.410
bgo_y1	40.30	50.64	2.602
bgo_y0	40.36	50.77	2.604
bgo_x1	39.73	49.67	2.623
bgo_x0	40.26	50.15	2.765

Table 3.1: Performance metrics for the baseline Poisson regression model at 1 s time resolution. Pearson dispersion $\phi \gg 1$ confirms strong overdispersion in all detectors.

Although the Poisson regressor captures the mean trend reasonably well, the estimated Pearson dispersion parameter ϕ remains significantly larger than one for all detectors. Indeed, the model achieves a MAPE of 4.22% on the test set.

This result provides a model-based confirmation of the overdispersion observed in the raw data: even after conditioning on the satellite orientation features, the variability of the counts exceeds the level predicted by a Poisson process.

Therefore, while the Poisson distribution provides a physically motivated baseline for photon counting, it remains too restrictive to describe the true variability of the background signal.

Per-detector differences in metrics. The MAE and RMSE values differ across detectors due to the varying average count rates recorded by each sensor. As clearly visible in Figure 2.3, some detectors systematically observe higher photon counts because of their geometric orientation relative to the spacecraft structure and the surrounding radiation environment. Since both MAE and RMSE are scale-dependent metrics, detectors operating at higher count levels naturally exhibit larger absolute errors even when the relative predictive accuracy remains comparable.

This behavior is expected and does not indicate worse model performance for those detectors. Rather, it reflects the intrinsic scaling of absolute error metrics with the magnitude of the observed signal. Consequently, a similar pattern will be observed for all regression models analyzed throughout this thesis.

3.3 Skew-Normal Regression

The Skew-Normal distribution, introduced by Azzalini [6] and further developed in later work [7], extends the Gaussian distribution by introducing a shape parameter that controls skewness. This extension allows the modeling of asymmetric data while preserving much of the analytical tractability of the normal distribution.

Such flexibility is particularly useful when modeling detector background counts, which often exhibit asymmetric fluctuations and local variability that cannot be adequately captured by a symmetric Gaussian model or by a simple Poisson assumption.

For these reasons, the Skew-Normal distribution provides a natural probabilistic framework for describing detector background variability. In particular, by allowing asymmetric uncertainty bands around the expected background level, it enables a more flexible representation of detector-specific behavior.

3.3.1 Characterizing the Distribution

A random variable X follows a Skew-Normal distribution if its probability density function is

$$f(x \mid \xi, \omega, \alpha) = \frac{2}{\omega}\phi\left(\frac{x-\xi}{\omega}\right)\Phi\left(\alpha\frac{x-\xi}{\omega}\right), \tag{3.8}$$

where ξ is the location parameter, $\omega > 0$ the scale parameter, and α the shape (skewness) parameter. Here, $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal PDF and CDF, respectively.

When $\alpha = 0$, the distribution reduces to a standard Gaussian. For $\alpha \neq 0$, the density becomes asymmetric, with α controlling both the direction and magnitude of skewness.

In our regression setting, the parameters $(\xi_t, \omega_t, \alpha_t)$ are predicted as functions of the satellite state features x_t. The location parameter models the expected background level, the scale parameter captures dispersion, and the skewness parameter allows asymmetric uncertainty, accommodating directional deviations observed in the background signal.

Despite this added flexibility, the model still assumes a single parametric regime. As shown in the following subsection, this assumption may remain too restrictive in the presence of multiple latent background states.

3.3.2 Implementation

The Skew-Normal regressor was implemented as a feed-forward neural network that maps the spacecraft state variables to the parameters of a detector-wise conditional Skew-Normal distribution. Given the input feature vector at time t, the network predicts, for each detector, the three parameters of the output distribution: location μ_t, scale σ_t, and shape α_t.

The model takes as input the full set of spacecraft orientation and orbital variables, and processes them through a stack of fully connected hidden layers with Rectified Linear Unit (ReLU) activation functions, defined as $\mathrm{ReLU}(x) = \max(0, x)$ [34]. The final hidden representation is then projected into three parallel output heads: one linear layer for the location parameter μ, one softplus layer for the positive scale parameter σ, and one linear layer for the skewness parameter α. The outputs of these three heads are concatenated to obtain the full parameter vector of the predictive distribution.

The hidden architecture was defined by two main hyperparameters: the initial number of units and the number of hidden layers. For a network with L hidden layers and initial width U, the default architecture progressively halves the number of units at each layer, yielding a structure of the form

$$U, \frac{U}{2}, \frac{U}{4}, \ldots$$

A variant controlled by the `flat` option was also considered. When `flat=True` and the network has at least three hidden layers, the last two hidden layers are constrained to have the same width, so that the final reduction in dimensionality is softened. This allows the model to retain a richer latent representation in the deepest part of the network.

Training framework. The model was implemented using the TensorFlow deep learning framework [1] through the `Keras` high-level API [12]. The neural network architecture follows the standard feed-forward paradigm commonly used in deep learning models [20]. Probabilistic components were handled through `TensorFlow Probability`.

The network parameters were optimized using the Adam optimizer [25], a fixed learning rate of 10^{-3}, and trained using mini-batch stochastic gradient descent.

Training procedure. Training was performed by minimizing the NLL (as defined before in subsection 3.3) of the observed counts under the predicted Skew-Normal distribution [9]. Given the model output parameters μ, σ, and α, the loss function evaluates the log-probability of the observed counts under the corresponding Skew-Normal distribution and averages it across detectors and time bins.

In practice, the standard likelihood objective was adapted to the Skew-Normal case by computing

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{N} \sum_{t,d} \log p_{\text{SN}}(y_{t,d} \mid \mu_{t,d}, \sigma_{t,d}, \alpha_{t,d}),$$

where p_{SN} denotes the Skew-Normal probability density function and the sum runs over all time bins t and detectors d. Minimizing this objective allows the model to learn the full predictive distribution of the detector counts rather than only the conditional mean.

Training employed mini-batches and early stopping to prevent overfitting. In particular, the optimization process was monitored on a validation set and automatically stopped if the validation loss did not improve for 10 consecutive epochs.

Model selection. Model selection was performed through a grid search over the following hyperparameter space:

- number of initial hidden units: {512, 256, 128},

- number of hidden layers: $\{2, 3, 4\}$,
- batch size: $\{64, 128, 256\}$,
- flat-tail architecture: $\{$`True`, `False`$\}$.

For the 50 ms dataset, the best-performing configuration was obtained with 128 initial units, 3 hidden layers, batch size 128, and `flat=True`. This corresponds to a hidden structure with widths $128 \to 64 \to 64$.

3.3.3 Results

Figure 3.2 shows the Skew-Normal regression at 50 ms resolution. At this time scale, the model captures the rapid variability of the background and produces asymmetric uncertainty bands that adapt locally to the observed fluctuations. The flexibility introduced by the skewness parameter allows the model to account for directional deviations and moderate overdispersion, yielding calibrated short-term predictions. The uncertainty bands are calibrated to a nominal 3σ coverage, but here that rule does not directly apply to skewed distributions, the interval is calibrated using Gaussian-equivalent quantiles through PPF-based quantile matching, preserving the asymmetric structure of the predictive distribution.

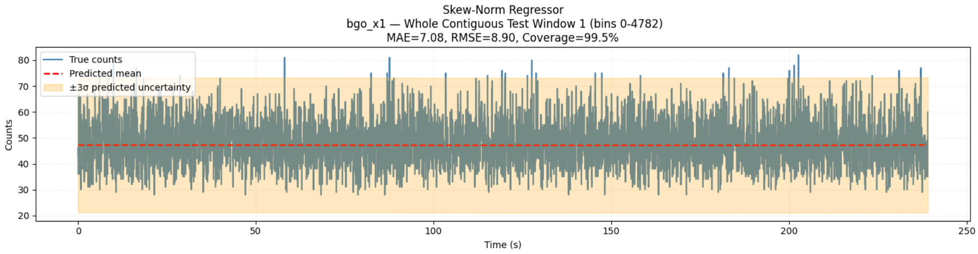

Figure 3.2: Skew-Normal predictions at 50 ms time resolution for detector $\mathrm{bgo}_{\mathrm{x1}}$, over a 4 min time window.

However, when the time resolution is increased to 1 s (Figure 3.3), a different behavior emerges. Although the model can be naturally applied at lower temporal resolutions—thus enabling the analysis of longer transients (e.g., events lasting more than two seconds)—the predictions exhibit high local variability.

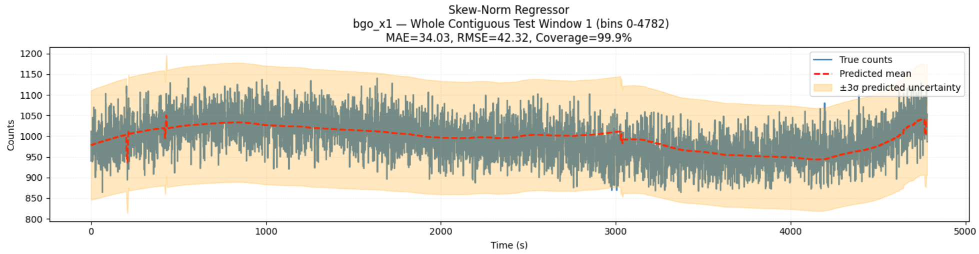

Figure 3.3: Skew-Normal background predictions at 1 s time resolution for detector $\mathrm{bgo}_{\mathrm{x1}}$ over a full orbital time window (approximately 80 min).The model adapts independently at each time bin, resulting in high local sensitivity to short-term fluctuations.

Table 3.2 shows the quantitative results for data at 1 s time resolution. The average percentage error across detectors (MAPE), is 3.46%.

Detector	MAE	RMSE
bgo_z1	42.27	52.93
bgo_z0	42.28	53.14
bgo_y1	33.07	41.43
bgo_y0	32.66	41.16
bgo_x1	32.50	40.70
bgo_x0	31.82	39.78

Table 3.2: Predictive performance of the Skew-Normal regression model on the test set at 1 s time resolution.

Because the Skew-Normal regression estimates distribution parameters independently at each time bin, effectively modeling a single regime per bin, the expected background level is entirely determined by the parameters predicted at that specific instant. As a consequence, the model reacts strongly to short-term fluctuations and may become overly sensitive to transient local variations, even when these do not correspond to structural background changes.

This behavior reflects an intrinsic limitation of single-regime parametric models: while they can adapt to asymmetric deviations, they lack an explicit mechanism to smooth across latent background states.

3.4 Gaussian-Mixture Regression

3.4.1 Handling multiple regimes

Gaussian Mixture Models (GMM; [38]) handle the single regime assumption limitation by explicitly modeling the presence of multiple background regimes. Instead of assuming a single global distribution, a mixture model represents the data as a weighted combination of several Gaussian components, each corresponding to a different latent background state. This allows the model to capture *multimodality*, regime switching, and increased variability that arise from orbital motion and detector geometry. By conditioning on the input features, Gaussian Mixture Regression [9] provides smoother and more flexible conditional predictions, effectively approximating complex, non-stationary background behavior that a single parametric model cannot represent. Related mixture-based conditional density models were previously introduced in the context of Mixture Density Networks [8].

3.4.2 Modeling in Log-Space

While detector counts are discrete and strictly positive, Gaussian mixture components are continuous and symmetric. To reconcile this mismatch, we model the logarithm of the counts rather than the raw counts directly:

$$\tilde{C}_t = \log(C_t + \epsilon),$$

where ϵ is a small constant introduced to avoid numerical instability in low-count bins.

The logarithmic transformation serves two purposes. First, it stabilizes variance and reduces heteroscedastic effects that increase with the count level. Second, it makes the empirical distribution closer to Gaussian within each latent regime, which aligns with the assumptions of mixture modeling.

This behavior becomes evident when inspecting the empirical distribution of the transformed counts. Figure 3.4 shows the histogram in log-space together with the fitted Gaussian mixture density and its individual components.

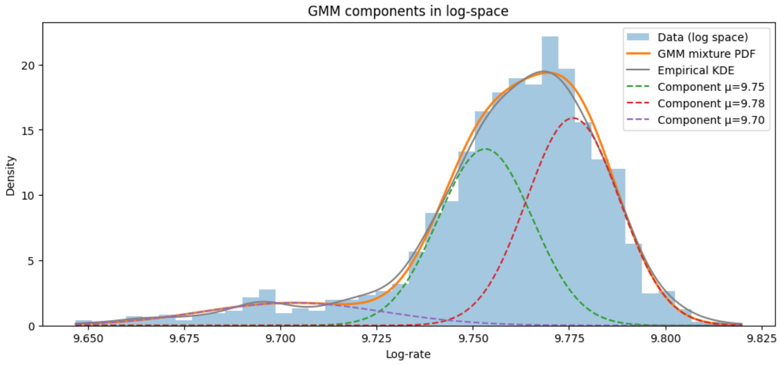

Figure 3.4: Histogram of log-transformed counts for detector bgo_{z1}, with empirical KDE (gray), full GMM density (orange), and individual Gaussian components (dashed). The mixture captures mild multimodality and asymmetric structure that cannot be represented by a single Gaussian component.

All Gaussian components are therefore fitted in log-space. Predictions are subsequently mapped back to the original scale when computing expected counts and uncertainty bands.

3.4.3 Theoretical Formulation

A Gaussian Mixture Model represents a probability density as a weighted sum of K Gaussian components:

$$p(x) = \sum_{k=1}^{K} \pi_k \mathcal{N}(x \mid \mu_k, \Sigma_k), \tag{3.9}$$

where π_k are non-negative mixture weights such that $\sum_k \pi_k = 1$, and (μ_k, Σ_k) are the mean and covariance of each component.

In our setting, we jointly model the satellite features x_t and detector counts C_t using a mixture distribution:

$$p(x_t, C_t) = \sum_{k=1}^{K} \pi_k \mathcal{N}\left(\begin{bmatrix} x_t \\ C_t \end{bmatrix} \middle| \mu_k, \Sigma_k \right). \tag{3.10}$$

Gaussian Mixture Regression is then obtained by conditioning on the observed features x_t. The conditional distribution of the counts becomes:

$$p(C_t \mid x_t) = \sum_{k=1}^{K} \tilde{\pi}_k(x_t)\, \mathcal{N}\Big(C_t \mid \mu_{k|x}(x_t), \Sigma_{k|x}\Big), \tag{3.11}$$

where the conditional means and covariances are derived analytically from the joint Gaussian components, and the mixture weights $\tilde{\pi}_k(x_t)$ depend on the posterior responsibility of each component given the input features.

3.4.4 Parameter Estimation

The mixture parameters are estimated via maximum likelihood using the Expectation-Maximization (EM; [17]) algorithm.

At each iteration:

- *E-step*: compute the posterior responsibility of component k for each observation.
- *M-step*: update mixture weights, means, and covariances using responsibility-weighted estimates.

EM iteratively increases the data log-likelihood until convergence, yielding a fully probabilistic generative model of the joint distribution.

3.4.5 Implementation

The Gaussian Mixture model was implemented using the `GaussianMixture` class from the `scikit-learn` library [39]. The model was trained on the joint dataset of spacecraft state variables and detector counts using the EM algorithm seen in the section above (3.4.4). After fitting the joint distribution, Gaussian Mixture Regression was obtained by conditioning the learned mixture model on the observed feature vectors in order to derive the predictive distribution of the detector counts.

Selection of the number of mixture components. The number of mixture components K was selected using a combination of information criteria and qualitative inspection of the predicted background curves. Figure 3.5 shows the evolution of the AIC [3] and BIC [43] scores as a function of K. Although both criteria decrease monotonically with increasing model complexity, the rate of improvement progressively diminishes beyond a small number of components. To complement this quantitative analysis, we also inspected the predicted background curves over representative time window corresponding to a full orbital period (the same of Figure 3.6). This qualitative comparison allowed us to assess the structural smoothness and physical consistency of the reconstructed background signal.

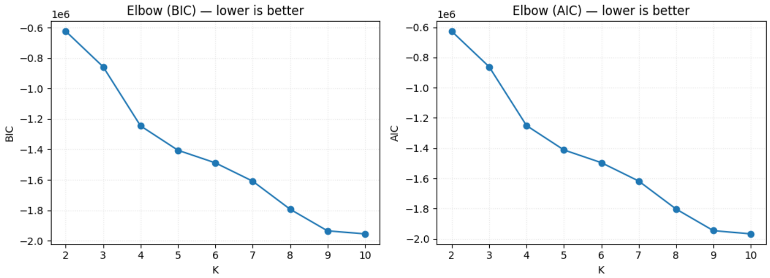

Figure 3.5: AIC and BIC scores as a function of the number of mixture components K used in the Gaussian Mixture Regression model. Lower values indicate better trade-offs between model fit and complexity.

Based on the combined evidence from the elbow analysis and the visual inspection of the regression behavior, $K = 4$ mixture components was selected as a suitable compromise between model flexibility and stability.

Model fitting. The Gaussian Mixture Regression model was fitted on the joint dataset composed of the spacecraft state variables and the corresponding detector counts. Each observation therefore consists of the feature vector x_t, describing the spacecraft orientation and orbital configuration, and the associated log-transformed count value $\tilde{C}_t$.

The joint vectors $(x_t, \tilde{C}_t)$ were modeled using a Gaussian mixture with $K = 4$ components. Each component represents a latent background regime characterized by a multivariate Gaussian distribution over both the input features and the log-count variable.

Once the joint distribution $p(x_t, \tilde{C}_t)$ was learned, the conditional predictive distribution of the counts was obtained by conditioning the mixture model on the observed features x_t. For each time bin, the model therefore produces a mixture of Gaussian conditional predictions whose weights depend on the posterior responsibility of each component given the current spacecraft state.

3.4.6 Results

Temporal behavior of the predictions. The background learned by Gaussian Mixture Regression is noticeably smoother compared to the Skew-Normal model. This behavior is structural rather than incidental.

In GMR, the predictive mean emerges from averaging across multiple background states, rather than being determined by a single set of parameters at each time bin.

This mechanism naturally introduces smoothing and structural stability. Short-term fluctuations affect individual component responsibilities but are mitigated by the weighted aggregation across regimes. As a result, GMR is less sensitive to local noise while remaining responsive to genuine regime transitions.

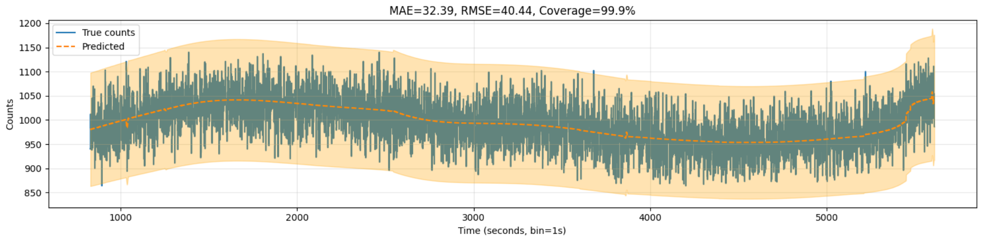

Figure 3.6: Gaussian Mixture Regression (GMR) predictions at 1 s time resolution for detector $\mathrm{bgo}_{\mathrm{x1}}$ over a full orbital time-window of approximately 80 min. The shaded region represents the nominal 3σ uncertainty band.

Figure 3.6 shows the 1 s resolution predictions. Since the predictive distribution is a mixture of Gaussians, the nominal 3σ coverage ($\approx 99.7\%$) does not correspond to a strict Gaussian mapping. Instead, symmetric bounds are computed around the predictive mean as an approximation of the conditional mixture variance.

Compared to the Skew-Normal model at the same resolution, GMR produces a more stable background estimate and more coherent uncertainty bands. The model is less reactive to isolated bin-level perturbations and better captures the underlying large-scale background evolution driven by orbital dynamics.

Overall, GMR trades local sensitivity for global robustness, providing a structurally smoother and more physically consistent background estimate.

Predictive performances. The predictive performance of the Gaussian Mixture Regression model evaluated on the test set is reported in Table 3.3. All metrics are computed using the 1 s dataset. The GMR model achieves a MAPE of 3.39%, improving by 0.07 percentage points over the Skew-Normal regressor.

Detector	MAE	RMSE
bgo_z1	41.58	51.72
bgo_z0	40.82	51.44
bgo_y1	31.85	39.77
bgo_y0	31.94	40.15
bgo_x1	31.56	39.64
bgo_x0	30.97	38.53

Table 3.3: Predictive performance of the Gaussian Mixture Regression model ($K = 4$ components) on the test set at 1 s time resolution.

The Gaussian Mixture Regression model achieves slightly better predictive performance than the Skew-Normal regressor, yielding a marginally lower percentage error (3.39% vs 3.46%).

3.5 Probabilistic Regression Summary

The three regression models considered in this section represent progressively more flexible probabilistic descriptions of the detector background.

The *Poisson model* provides the natural theoretical baseline for photon counting processes. However, its single-parameter formulation imposes the constraint $\mathrm{Var}(C_t) = E[C_t]$, which is inconsistent with the empirical variability observed in the data. Both the exploratory mean–variance analysis and the fitted regression results confirm strong overdispersion, indicating that a simple Poisson process cannot capture the full stochastic behavior of the background signal.

The *Skew-Normal regression* relaxes this constraint by introducing explicit scale and skewness parameters. This allows the predictive distribution to represent asymmetric fluctuations and moderate overdispersion. As a consequence, the model produces flexible uncertainty bands that adapt locally to short-term variations in the counts. However, since the parameters are estimated independently at each time bin, the model effectively assumes a single regime per observation and can become overly sensitive to local noise.

The *Gaussian Mixture Regression* further extends this framework by explicitly modeling multiple latent background regimes. By representing the conditional distribution as a weighted mixture of Gaussian components, the model can capture regime switching associated with orbital motion and detector geometry. The resulting predictions are structurally smoother, as the conditional expectation emerges from the weighted contribution of several latent states rather than from a single parameter set.

From a predictive perspective, both the Skew-Normal and GMR models achieve very similar accuracy in terms of MAE and RMSE on the 1 s dataset, with the GMR model showing a marginal improvement corresponding to a reduction of about 0.07% in the percentage error.

This indicates that the dominant background trends are largely driven by the spacecraft orientation and orbital configuration, which can be captured by both models.

More importantly, the improvement with respect to the baseline Poisson regression is substantial: the MAPE decreases from 4.22% to 3.39%, corresponding to a reduction of 0.83 percentage points.

Nevertheless, the predictive behavior of the Skew-Normal and GMR models differs due to their underlying statistical formulation. The Skew-Normal model prioritizes local adaptability, while the Gaussian mixture model introduces a form of implicit temporal smoothing through regime averaging. For this reason, GMR provides more stable long-timescale background reconstructions, which are particularly advantageous for anomaly detection tasks where robustness against local noise is essential.

Table 3.4 summarizes the predictive performance of the three probabilistic regressors.

Model	MAE (avg)	RMSE (avg)	Key limitation
PR	43.35	54.35	Cannot model overdispersion
SNR	35.77	44.86	High local sensitivity
GMR	34.79	43.54	Higher model complexity

Table 3.4: Summary of predictive performance for the probabilistic regression models at 1 s time resolution. Metrics are averaged across detectors. Abbreviations: PR = Poisson Regressor, SNR = Skew-Normal Regressor, GMR = Gaussian Mixture Regressor.

Overall, while the Poisson model provides a useful theoretical baseline, the empirical variability of the detector background requires richer probabilistic formulations. Among the models considered, Gaussian Mixture Regression offers the best compromise between predictive accuracy, stability, and physical interpretability of the background regimes.

Chapter 4

Explainable Background Modeling

Interpretability plays an important role in this work, particularly from a physical perspective. While predictive performance remains a primary objective, interpretable models provide an additional advantage by helping to understand how the input variables influence the predicted background signal. This is particularly useful in scientific applications, where model behavior should remain consistent with the underlying physical processes.

In the context of gamma-ray background modeling, interpretability makes it possible to verify that the model relies on physically meaningful variables related to the spacecraft state, rather than on spurious correlations arising from the training data.

In the machine learning literature, a distinction is often made between *explainability* and *interpretability*. Explainability refers to post-hoc techniques that analyze an already trained model in order to estimate how different input variables influence its predictions. Interpretability, instead, typically refers to models whose internal structure is inherently understandable, allowing the relationship between inputs and outputs to be directly inspected.

To leverage both perspectives, this work adopts two complementary approaches.

SHAP (SHapley Additive exPlanations; [32]) is used as a post-hoc explainability method to analyze the trained probabilistic regressors. In particular, SHAP is applied to both the Skew-Normal and Gaussian Mixture Regression models in order to quantify the contribution of each input feature to the predicted background counts, enabling both local and global assessments of feature importance.

In parallel, *Kolmogorov–Arnold Networks* (KANs; [31]) are explored as an intrinsically interpretable neural architecture. Unlike SHAP, which explains already trained models, KANs constitute a separate regression model whose structure explicitly represents univariate functional relationships between inputs and hidden neurons that can be directly visualized and analyzed.

Together, these tools allow us to verify that the learned models remain physically coherent, transparent, and scientifically reliable.

4.1 SHAP-based post-hoc explanations

4.1.1 Theoretical Background

SHAP is based on the concept of Shapley values from cooperative game theory [44]. Given a model $f(x)$ and an input vector $x \in R^d$, the prediction is expressed as an additive decomposition:

$$f(x) = \phi_0 + \sum_{i=1}^{d} \phi_i, \tag{4.1}$$

where ϕ_0 is the baseline prediction and ϕ_i represents the contribution of feature i to the final output.

Baseline value. The baseline is defined as the expected model prediction over a background data distribution:

$$\phi_0 = E_{x \sim p(x)}[f(x)]. \tag{4.2}$$

Since the true data distribution $p(x)$ is unknown, this expectation is approximated empirically using Monte Carlo sampling over a representative background dataset:

$$\phi_0 \approx \frac{1}{N} \sum_{x \in \mathcal{B}} f(x), \tag{4.3}$$

where $\mathcal{B}$ denotes the selected background set.

Shapley values. For a specific test sample x, the Shapley value ϕ_i is defined as the average marginal contribution of feature i across all possible subsets of features. Formally,

$$\phi_i = \sum_{S \subseteq \mathcal{F} \setminus \{i\}} \frac{|S|!(d - |S| - 1)!}{d!} \Big[f_{S \cup \{i\}}(x) - f_S(x) \Big], \tag{4.4}$$

where:

- $\mathcal{F}$ is the set of all features,
- S is a subset of features not containing i,
- $f_S(x)$ denotes the model prediction when only features in S are known and the remaining ones are marginalized.

This definition ensures a fair allocation of the prediction difference

$$f(x) - \phi_0 = \sum_{i=1}^{d} \phi_i, \tag{4.5}$$

distributing the deviation from the baseline among all features.

Interpretation. Each ϕ_i quantifies how much feature i shifts the prediction away from the expected model output. Positive values indicate that the feature increases the predicted background level, while negative values indicate a decreasing effect.

4.1.2 Implementation Choices

Feature representation for explainability. Although the predictive models are trained using the sine and cosine encoding of angular variables in order to avoid artificial discontinuities, SHAP results are discussed in terms of physically meaningful angular quantities, namely the original orientation variables and, when needed, their residualized counterparts (for reasons we will see in the upcoming paragraphs). This choice improves physical interpretability, since raw or residualized angles retain a direct connection to the spacecraft geometry, whereas the sine/cosine representation would split each physical angle into two transformed coordinates.

Reference models for SHAP analysis. The SHAP analysis is performed on both the Skew-Normal regression model and the Gaussian Mixture Regression (GMR) model, with different roles. The Skew-Normal model provides a smoother and more global view of the relationship between input variables and background prediction, making it useful for dependence analysis. The GMR model, on the other hand, allows a more detector-specific interpretation, since it can capture multiple background regimes and therefore reveals stronger variability in feature attribution across detectors.

Unless otherwise specified, dependence-style analyses are discussed using the Skew-Normal model, while detector-wise global feature rankings are shown using the GMR model.

4.1.3 Correlation Structure of the Input Features

Before interpreting SHAP values, it is important to assess the correlation structure of the input variables.

Since the residualization procedure introduced later relies on a linear regression model, correlations are evaluated using the Pearson coefficient, which measures linear dependence between variables. In practice, Pearson and Spearman correlation matrices were found to be very similar, indicating that the dominant relationships between the considered features are largely linear.

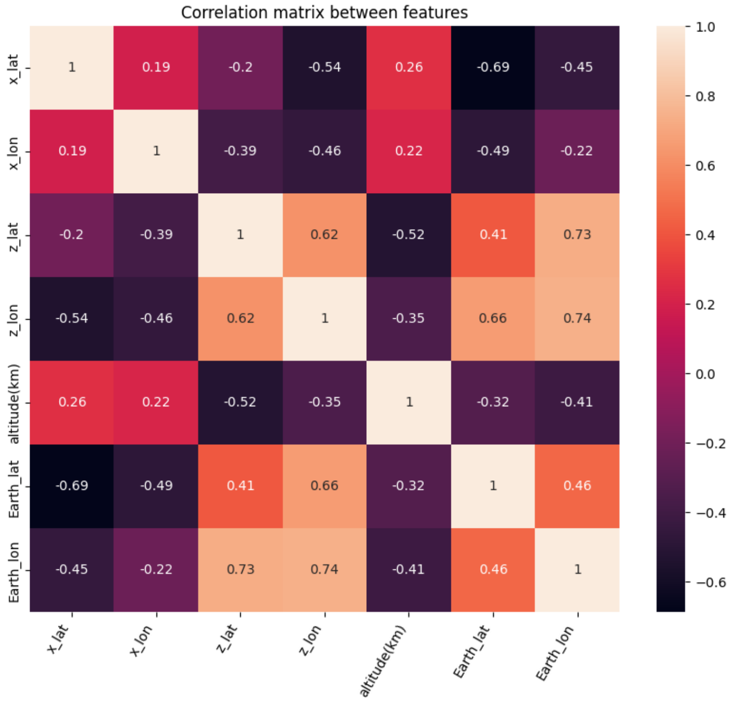

	x_lat	x_lon	z_lat	z_lon	altitude(km)	Earth_lat	Earth_lon
x_lat	1	0.19	-0.2	-0.54	0.26	-0.69	-0.45
x_lon	0.19	1	-0.39	-0.46	0.22	-0.49	-0.22
z_lat	-0.2	-0.39	1	0.62	-0.52	0.41	0.73
z_lon	-0.54	-0.46	0.62	1	-0.35	0.66	0.74
altitude(km)	0.26	0.22	-0.52	-0.35	1	-0.32	-0.41
Earth_lat	-0.69	-0.49	0.41	0.66	-0.32	1	0.46
Earth_lon	-0.45	-0.22	0.73	0.74	-0.41	0.46	1

Figure 4.1: Pearson correlation matrix of the input features. Strong correlations are observed between some orbital and pointing variables, indicating partial redundancy in the geometric information provided to the model.

Several orientation and orbital features are not independent, since they describe related aspects of the spacecraft geometry with respect to the Earth. In particular, the spacecraft pointing coordinates and the orbital longitude exhibit substantial correlation.

This correlation is physically expected. As shown in the correlation matrix in Figure 4.1, variables such as $\mathrm{z}_{\mathrm{lon}}$ and $\mathrm{Earth}_{\mathrm{lon}}$ are both linked to the spacecraft orientation relative to the Earth along the orbit, and therefore partially encode the same underlying geometric configuration. As a consequence, feature attribution methods may not assign importance uniquely, but instead split the contribution across correlated variables.

4.1.4 SHAP Dependence Analysis and Collinearity

A SHAP dependence plot shows how the contribution of a specific feature varies as a function of its value, while optionally highlighting interactions with another feature through color encoding.

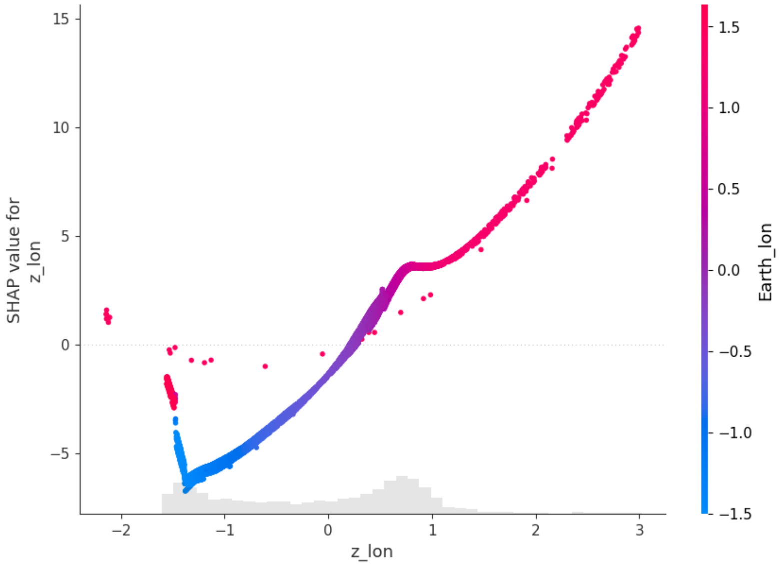

Figure 4.2: SHAP dependence plot for the feature $\mathrm{z}_{\mathrm{lon}}$, colored according to $\mathrm{Earth}_{\mathrm{lon}}$.

Figure 4.2 highlights a structured dependence between the SHAP contribution of $\mathrm{z}_{\mathrm{lon}}$and the orbital longitude $\mathrm{Earth}_{\mathrm{lon}}$. This is an important observation: the feature interaction is not visible only in the correlation matrix of the raw inputs, but also propagates to the SHAP values themselves. In other words, the attribution assigned to $\mathrm{z}_{\mathrm{lon}}$ depends in part on another strongly related geometric variable.

The dependence plot indeed reveals a clear nonlinear relationship between the value of $\mathrm{z}_{\mathrm{lon}}$ and its SHAP contribution. For negative values of $\mathrm{z}_{\mathrm{lon}}$, the contribution tends to be negative, indicating that this spacecraft orientation is associated with lower predicted background levels. As $\mathrm{z}_{\mathrm{lon}}$ increases, the SHAP value grows smoothly and eventually becomes strongly positive, meaning that certain orientations of the spacecraft boresight contribute to increasing the predicted detector background.

The color gradient further highlights the interaction with $\mathrm{Earth}_{\mathrm{lon}}$. Points with similar $\mathrm{z}_{\mathrm{lon}}$ values but different colors correspond to different orbital longitudes, showing that the contribution of the pointing direction cannot be interpreted independently of the spacecraft orbital position. In particular, the transition region around intermediate $\mathrm{z}_{\mathrm{lon}}$ values shows a continuous color variation, indicating that the model jointly uses both variables to explain the observed background variability.

This behavior suggests that the detector background depends not only on the instantaneous pointing direction but also on the relative geometry between the spacecraft orientation and the Earth position along the orbit. The SHAP dependence plot therefore provides a useful visualization of how correlated geometric variables jointly influence the predicted background counts.

This behavior reflects a well-known limitation of attribution methods in the presence of *collinearity*. When two variables carry overlapping information, SHAP may distribute the importance between them in a non-unique way.

As a result, the ranking of correlated features should not always be interpreted as a strict physical separation of mechanisms, but rather as a partition of shared explanatory power.

4.1.5 Residualization for Correlated Features

To partially address this issue, we apply a simple residualization strategy to selected pairs of strongly correlated variables. The idea is to preserve one reference feature and replace the correlated companion with its residual component, i.e. the part that cannot be linearly explained by the reference feature.

More precisely, given a correlated pair (F, G), we:

1. choose a main feature F,

2. fit a regression model such that $G \approx h(F)$,

3. replace G with the residual

$$R = G - h(F).$$

The residual feature R therefore contains only the information in G that is not already explained by F.

In our implementation, a simple linear regression is used:

$$G \approx aF + b. \tag{4.6}$$

The resulting transformed feature is then

$$G_{\text{res}} = G - (aF + b). \tag{4.7}$$

This procedure does not alter the physical signal itself, but changes its parameterization in order to reduce the ambiguity of the attribution. It should be regarded as an interpretability-oriented sensitivity analysis rather than as a guaranteed correction. Its purpose is not to improve predictive performance, but to test whether the SHAP hierarchy becomes more discriminative once overlapping geometric information is partially disentangled.

In practice, this residualization is applied to orientation variables that exhibit strong correlation with orbital longitude, such as z_{lon} and z_{lat} with respect to $\mathrm{Earth}_{\text{lon}}$.

The non-residualized SHAP analysis is used here only as a diagnostic reference to identify attribution ambiguity due to correlated variables. For clarity, the interpretability results discussed in the following sections are reported for the residualized representation.

4.1.6 Global Interpretability

We analyze the global behavior of the model using SHAP values, which aggregate feature contributions across many instances of the dataset.

Sampling strategy for SHAP analysis. For interpretability analyses, the samples used to compute SHAP values are randomly drawn from the test dataset rather than taken from a contiguous temporal segment. As a result, the instances represented in the plots originate from different orbital phases and observational conditions.

This sampling strategy ensures that the SHAP analysis captures the global behavior of the model across multiple orbital regimes, rather than reflecting the dynamics of a single orbit segment.

Beeswarm Plot. A beeswarm plot summarizes the distribution of SHAP values for each feature across the dataset. Each point corresponds to one sample, showing both the magnitude and direction of the contribution, while the color encodes the corresponding feature value.

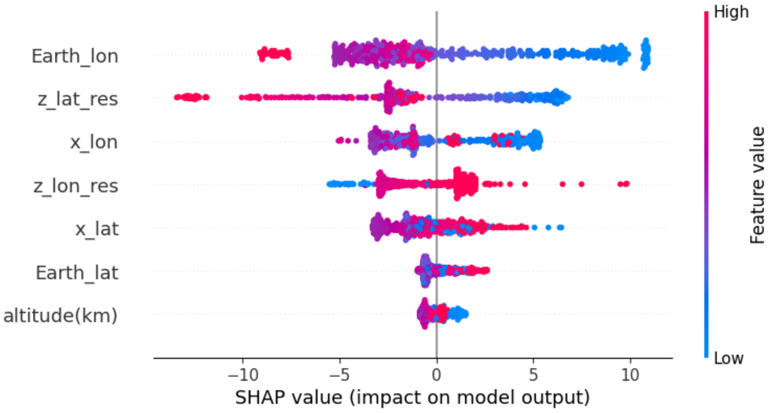

Figure 4.3: SHAP beeswarm plot computed after residualization of correlated orientation variables.

Figure 4.3 provides a global view of feature effects after mitigating the strongest collinearity. The orbital longitude $\mathrm{Earth}_{\mathrm{lon}}$ emerges as a dominant driver, consistent with an orbit-phase modulation of the background. The residualized pointing latitude $\mathrm{z}_{\mathrm{lat\_res}}$ also shows large contributions and a clear negative association with its SHAP values, indicating that higher $\mathrm{z}_{\mathrm{lat\_res}}$ tends to decrease the predicted background, while lower values increase it.

Spacecraft-frame orientation variables ($\mathrm{x}_{\mathrm{lon}}, \mathrm{x}_{\mathrm{lat}}$) provide intermediate contributions, reflecting changes in detector exposure geometry with respect to the Earth limb and spacecraft structure. The residualized longitude $\mathrm{z}_{\mathrm{lon\_res}}$ remains non-negligible, suggesting that a pointing-dependent effect persists even after removing the component linearly shared with $\mathrm{Earth}_{\mathrm{lon}}$. Finally, altitude and $\mathrm{Earth}_{\mathrm{lat}}$ contribute weakly, indicating limited influence on short-timescale modulations.

Detector-wise SHAP bar plots. While the beeswarm plot summarizes global behavior across samples, detector-wise bar plots aggregate SHAP magnitudes separately for each detector and highlight geometry-dependent sensitivities. In practice, the global importance of each feature j is computed as the mean absolute SHAP value across the analyzed samples:

$$I_j = \frac{1}{N} \sum_{i=1}^{N} |\phi_j^{(i)}|, \tag{4.8}$$

where $\phi_j^{(i)}$ denotes the SHAP value of feature j for sample i, and N is the number of samples used in the analysis. This quantity measures the average magnitude of the contribution of each feature to the model prediction.

GMR: Detector-wise SHAP bar plots. Figure 4.4 shows the aggregated SHAP feature importance obtained with the GMR model after residualization. The detector-specific patterns can be interpreted in light of the physical detector layout revealing several physically meaningful trends.

First, the orbital longitude $\mathrm{Earth}_{\mathrm{lon}}$ consistently appears among the dominant features for most detectors, and it is the primary driver for $\mathrm{bgo}_{\mathrm{x0}}$ and $\mathrm{bgo}_{\mathrm{y0}}$. This indicates that the background level is strongly modulated by the

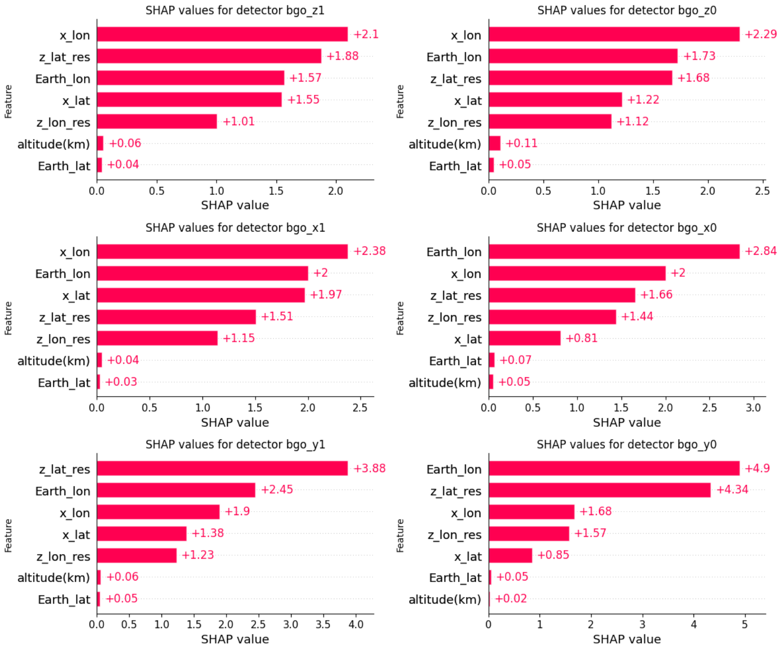

Figure 4.4: Detector-wise SHAP feature contributions for the GMR model after residualization.

orbital phase of the spacecraft, which determines the relative geometry between the instrument and the Earth environment.

Second, the residualized pointing latitude $\mathrm{z}_{\mathrm{lat\_res}}$ shows particularly strong influence for the Y detectors, where it becomes the most important feature. This suggests that variations in the inclination of the spacecraft boresight relative to the Earth direction significantly affect the exposure geometry of these detectors.

Finally, the spacecraft-frame longitude $\mathrm{x}_{\mathrm{lon}}$ consistently ranks among the most important variables across several detectors, indicating that rotations of the spacecraft around the X axis play a relevant role in shaping the detector background.

In contrast, the orbital altitude and $\mathrm{Earth}_{\mathrm{lat}}$ exhibit very small SHAP magnitudes across all detectors, suggesting that they contribute only marginally to the short-timescale background variability captured by the model.

Skew-Normal: Detector-wise SHAP bar plots. Figure 4.5 reports the corresponding detector-wise SHAP contributions for the Skew-Normal model after residualization.

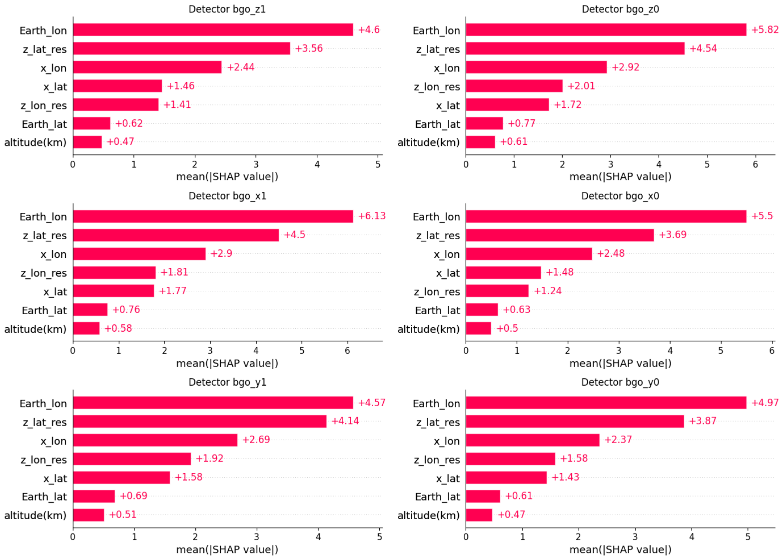

Figure 4.5: Detector-wise SHAP feature contributions for the Skew-Normal model after residualization.

Compared to GMR, the Skew-Normal model produces more homogeneous rankings across detectors. This is expected because it describes the conditional background with a single parametric regime per time bin, whereas GMR can represent multiple latent regimes. As a result, detector-specific sensitivities are more clearly exposed by the mixture model.

4.1.7 SHAP Summary and Physical Consistency

Across detectors and models, SHAP yields a physically coherent picture in which orbital phase and spacecraft orientation jointly govern the background.

1. Orbital phase: $\mathrm{Earth}_{\mathrm{lon}}$ is consistently dominant, indicating a strong orbit-dependent modulation of the background level.
2. Pointing and orientation: $z_{\mathrm{lat\_res}}, z_{\mathrm{lon\_res}}, \mathrm{x}_{\mathrm{lat}}, \mathrm{x}_{\mathrm{lon}}$ modulate the detector exposure geometry relative to the surrounding radiation environment

and spacecraft structure, with detector-specific sensitivity consistent with the known panel layout.

3. Secondary effects: $\mathrm{altitude}$ and Earth_{lat} remain marginal on short timescales, as they evolve more slowly and induce weaker geometric changes within a single orbit segment.

Overall, the agreement between attribution patterns and detector geometry acts as a physical consistency check, supporting the conclusion that the models rely on meaningful orbital and pointing drivers rather than spurious correlations.

Chapter 5

Kolmogorov-Arnold Networks: Interpretable Function Learning

Kolmogorov–Arnold Networks (KANs) are a class of neural networks designed to improve interpretability. Unlike standard neural networks, where connections are simple weights, KANs learn explicit univariate functions on each edge. This means that every connection represents a functional relationship between variables, rather than just a number. As a result, KANs allow us to directly visualize and interpret how each input feature contributes to the model's internal representation and final prediction. This makes them particularly suitable for scientific applications where understanding the structure of the learned dependencies is as important as prediction accuracy.

5.1 Theoretical Background

5.1.1 Universal Approximation and Multi-Layer Perceptrons

The Universal Approximation Theorem (UAT; [24]) states that a sufficiently wide feed-forward neural network with non-linear activation functions can approximate any continuous function on a compact domain to arbitrary precision. In standard Multi-Layer Perceptrons (MLPs; [22]), this approximation is achieved through linear combinations of fixed activation functions applied to weighted sums of the inputs. Although this guarantees expressive power, the learned representation is distributed across many parameters, making the internal structure difficult to interpret.

5.1.2 Kolmogorov–Arnold Representation Theorem

The Kolmogorov–Arnold Theorem (KAT; [27, 5]) provides a stronger structural result. It states that any continuous multivariate function $f : R^d \to R$ can be represented as a finite superposition of continuous univariate functions:

$$f(x_1, \ldots, x_d) = \sum_{q=1}^{2d+1} \Phi_q \left(\sum_{p=1}^{d} \psi_{pq}(x_p) \right), \tag{5.1}$$

where ψ_{pq} and Φ_q are continuous one-dimensional functions.

Further discussions and modern expositions of the theorem can be found in the collected works of Arnold [4].

Unlike the UAT, which guarantees approximation through layered compositions of fixed nonlinearities, the KAT expresses multivariate functions explicitly in terms of sums and compositions of univariate functions. This structural decomposition provides a theoretical foundation for constructing models whose internal components are intrinsically interpretable.

5.1.3 From KAT to Neural Architectures

Kolmogorov–Arnold Networks [31] translate this theoretical result into a practical neural architecture. Instead of using fixed activation functions applied to linear transformations, KANs learn adaptive univariate functions along each edge of the network. In this framework, the model approximates multivariate mappings by combining learned one-dimensional transformations, closely mirroring the structure suggested by the Kolmogorov–Arnold theorem.

5.1.4 General Architecture of KANs

In a KAN layer, connections between neurons are not simple scalar weights but parameterized univariate functions. Given an input feature x_i, the connection to a hidden unit is defined as

$$h_j = \sum_i \phi_{ij}(x_i), \tag{5.2}$$

where each $\phi_{ij}(\cdot)$ is a learned one-dimensional function. This replaces the classical affine transformation $Wx + b$ of MLPs with a sum of adaptive functional mappings.

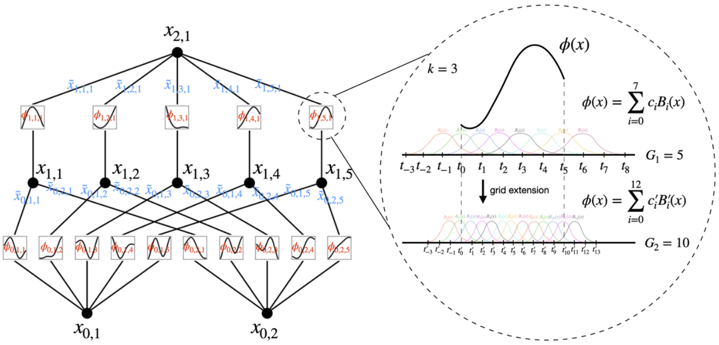

Figure 5.1: Conceptual architecture of a KAN model, adapted from [31].

5.1.5 B-spline Parameterization

In practice, the univariate functions ϕ_{ij} are parameterized using B-splines. B-splines provide a flexible and numerically stable basis for approximating smooth functions while maintaining locality: each control point affects only a limited region of the input domain.

This locality offers several advantages:

- Smooth and stable function approximation,
- Reduced interference between distant input regions,
- Direct visual interpretability of learned functional shapes.

Because each edge corresponds to a learnable curve, the model's internal structure can be directly visualized as a collection of one-dimensional response functions.

5.1.6 Intrinsic Interpretability.

Unlike post-hoc interpretability methods, KANs provide intrinsic transparency. Each learned function $\phi_{ij}(x_i)$ explicitly describes how a single input variable influences downstream activations. By inspecting these functions, it becomes possible to analyze monotonic trends, saturation effects, and nonlinear sensitivities in a physically meaningful manner.

In the context of spacecraft background modeling, this allows us to directly visualize how orbital position or pointing angles affect the predicted count rate, ensuring that learned relationships remain consistent with known physical behavior.

5.1.7 Advantages and limitations with respect to MLPs

Compared with standard MLPs, KANs offer a different trade-off between expressivity and interpretability. Their main advantage lies in the fact that nonlinear relationships are represented through explicit univariate functions on the network edges, rather than being implicitly distributed across many scalar weights and fixed activations. This makes KANs particularly attractive in scientific settings where understanding the learned functional dependencies is as important as achieving good predictive performance.

A second advantage is that the learned transformations can be directly visualized, inspected, pruned, and in some cases symbolically approximated. This enables a more structured analysis of how input variables contribute to intermediate representations and final predictions, which is much less straightforward in standard MLPs.

On the other hand, this greater transparency comes with some limitations.

KANs are typically more cumbersome to analyze and optimize than MLPs, since each connection carries a parameterized function rather than a single scalar weight. Moreover, while MLPs remain a more standard and broadly adopted baseline for generic function approximation, KANs are especially useful when the goal is not only prediction, but also the inspection of the learned functional structure.

For the present application, this trade-off is favorable: the slight increase in architectural complexity is compensated by the possibility of directly examining how orbital and pointing variables influence the predicted detector background.

5.1.8 Recent applications of Kolmogorov–Arnold Networks

Following their introduction, KANs have rapidly attracted attention across a variety of machine learning domains due to their combination of expressive power and interpretability. Recent studies have explored KAN architectures in several supervised learning tasks, including regression and classification, often reporting competitive performance with respect to standard MLPs while providing a more interpretable representation of the learned functional relationships [18].

In particular, KANs have been investigated in scientific machine learning contexts where understanding the structure of the learned model is essential [30].

Beyond standard regression tasks, KAN-based models have also been proposed for anomaly detection problems. In these settings, the ability of KANs to learn structured nonlinear transformations has been exploited to identify deviations from typical data patterns while preserving an interpretable representation of the underlying model behavior [49].

KAN architectures have also been explored in time-series forecasting and tabular learning tasks, where their flexible functional representation can capture periodic and nonlinear dependencies while maintaining a relatively compact model structure [21].

These characteristics make KANs particularly attractive for scientific applications in which predictive performance and interpretability are both important. In the context of space instrumentation and astrophysical data analysis, interpretable models can help relate learned statistical patterns to underlying physical drivers, such as orbital geometry and detector orientation. This motivates the use of KANs in the present work to analyze the dependence of detector background rates on spacecraft orbital and pointing parameters.

5.2 Implementation

The implementation of the Kolmogorov–Arnold Network models was carried out using the `PyKAN` framework, the official library released by the authors of the Kolmogorov–Arnold Networks paper [31]. The framework provides a reference implementation of KAN architectures and utilities for training, pruning, and inspecting the learned univariate functions.

5.2.1 Network Architecture

Figure 5.2 shows the structure of the Kolmogorov–Arnold Network used for background modeling after the pruning and fine-tuning stages described in the following section. The nodes at the bottom correspond to the input variables describing the spacecraft geometry and orbital configuration, while the nodes at the top correspond to the predicted background counts for each detector.

Inputs. The network receives 10 input variables describing the spacecraft orientation and orbital configuration. These include the orbital altitude, the sine and cosine encoding of the orbital longitude ($\sin(\mathrm{Earth}_{\mathrm{lon}})$, $\cos(\mathrm{Earth}_{\mathrm{lon}})$), the sine and cosine encoding of the spacecraft x-axis longitude ($\sin(\mathrm{x}_{\mathrm{lon}})$,

$\cos(\mathrm{x_{lon}})$), the sine encoding of the orbital latitude ($\sin(\mathrm{Earth_{lat}})$), the sine encoding of the spacecraft x-axis latitude ($\sin(\mathrm{x_{lat}})$), and the residualized pointing variables ($\sin(\mathrm{z_{lat\_res}})$, $\sin(\mathrm{z_{lon\_res}})$, $\cos(\mathrm{z_{lon\_res}})$).

Outputs. The network produces six output values corresponding to the predicted background counts for the ACS detectors $\mathrm{bgo_{z1}}$, $\mathrm{bgo_{z0}}$, $\mathrm{bgo_{x1}}$, $\mathrm{bgo_{x0}}$, $\mathrm{bgo_{y1}}$, and $\mathrm{bgo_{y0}}$.

Architecture. Each connection between two neurons is associated with a learnable transformation parameterized using B-spline basis functions with grid size 5 and spline order 3. This representation allows the network to approximate nonlinear relationships while maintaining smooth and interpretable functional transformations.

In Figure 5.2, each edge represents one of these learned univariate functions. The opacity of the connections qualitatively reflects the magnitude of the learned transformations, providing a visual indication of which interactions contribute most to the final prediction.

Although this global visualization provides an overview of the network connectivity, it does not reveal the explicit functional form learned along each edge. For this reason, in the following analysis we inspect individual learned functions between the input layer and the hidden layer in order to interpret how specific physical variables influence the predicted detector background.

5.2.2 Training Procedure

The Kolmogorov–Arnold Network was trained using a supervised regression setting, such that training was performed for 6000 optimization steps, using the Adam optimizer and a fixed learning rate of 10^{-3}. The model parameters were optimized by minimizing the mean squared error (MSE), corresponding to the L_2 loss between the predicted and observed detector count rates.

A small regularization penalty with strength $\lambda = 5 \times 10^{-6}$ was applied in

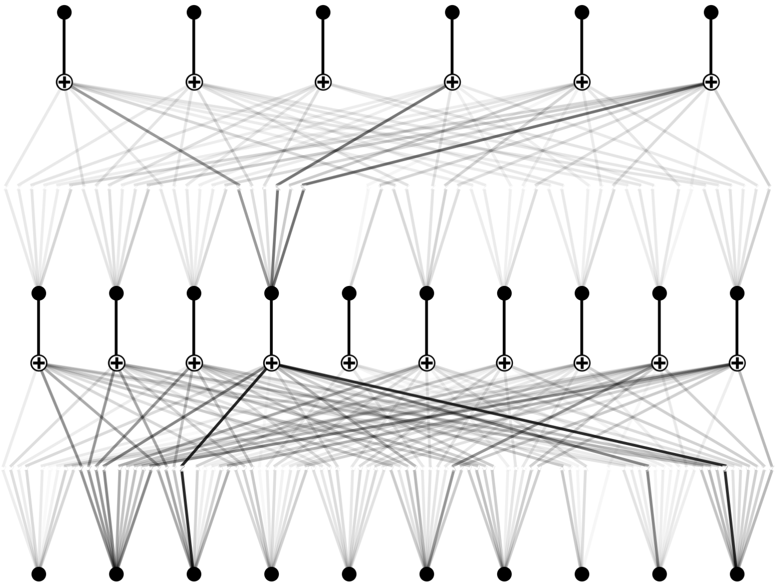

Figure 5.2: Visualization of the pruned Kolmogorov–Arnold Network used for background prediction. Bottom nodes correspond to the input features, the middle layer represents the hidden neurons, and the top nodes correspond to the predicted detector background counts. Darker edges indicate stronger learned transformations.

order to encourage smoother functional representations and reduce overfitting during training.

5.2.3 Pruning and Fine-Tuning

The initial network architecture is intentionally defined with sufficient capacity to capture potentially complex dependencies between the spacecraft configuration and the detector background rates. Since the intrinsic structure of the problem is not known a priori, this overparameterized architecture allows the model to learn a flexible functional representation during training.

After the initial training phase, a pruning procedure was applied to simplify the network structure by removing weak or redundant connections. Pruning acts as a form of structural regularization and is commonly used in Kolmogorov–Arnold Networks to identify the minimal functional structure required to approximate the target mapping, as discussed in Subsection 2.5 of [31].

The pruning process reduces the number of effective connections and highlights the most relevant functional dependencies between inputs, hidden neurons, and outputs. After pruning, a short fine-tuning phase (4000 optimization steps) was performed to allow the remaining parameters to readjust and stabilize the learned functions.

5.3 Regression results

5.3.1 Predictive performance

The pruned KAN achieves competitive predictive accuracy across all detectors (Table 5.1). Compared to the non-pruned KAN, pruning leads to a small improvement in MAPE of approximately 0.07 percentage points, resulting in a final error of about 3.39%. Since the model captures the main background variability, its learned functional structure can be meaningfully inspected for interpretability.

Detector	MAE	RMSE
bgo_z1	41.55	51.67
bgo_z0	40.89	51.51
bgo_y1	31.57	39.68
bgo_y0	30.99	38.58
bgo_x1	32.13	40.07
bgo_x0	32.10	40.38

Table 5.1: Predictive performance of the pruned KAN regressor on the test set.

5.3.2 Temporal behavior of the predictions

Figure 5.3 shows a representative light curve segment together with the background signal predicted by the pruned KAN model. The figure corresponds to a time window of 3920 bins at $1\,\mathrm{s}$ resolution, corresponding to approximately 65 min of observation.

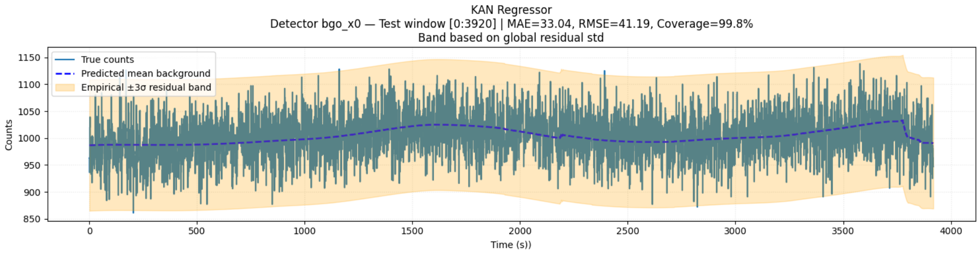

Figure 5.3: KAN background predictions at $1\,\mathrm{s}$ time resolution for detector $\mathrm{bgo}_{\mathrm{x0}}$ over a ≈ 65 min time window. The shaded region represents an empirical $\pm 3\sigma$ residual band used to visualize the scale of statistical fluctuations.

Despite being a deterministic regressor, the KAN model captures the underlying background level and produces a stable estimate of the expected count rate over time. The predicted background remains smooth compared to the rapid stochastic fluctuations visible in the observed counts, similarly to what is obtained with the probabilistic regression models introduced earlier (Skew-Normal and Gaussian Mixture Regression, Sections 3.3 and 3.4).

A clear slowly varying modulation is visible in the predicted background curve. This behavior reflects the orbital and geometrical dependencies of the detector background: as the spacecraft position and orientation evolve along the orbit, the expected background level changes accordingly. The KAN model successfully adapts to this orbital pattern, reproducing the large-scale variations of the background while leaving the high-frequency variability to the intrinsic stochastic counting noise.

For visualization purposes, we also report an empirical uncertainty band corresponding to approximately $\pm 3\sigma$, estimated from the standard deviation of the residuals on the test set. Although this band *does not represent a true predictive distribution* (since the KAN model does not explicitly model the stochastic variability of the counts), it provides a useful reference scale for the magnitude of the statistical fluctuations around the predicted background level.

Overall, the figure confirms that the KAN regressor successfully captures the main background trend driven by orbital conditions, while the residual variability is dominated by the intrinsic counting noise of the detector.

5.4 Interpretation of Learned Functions

5.4.1 Preliminary Functional Analysis

Before delving deeper into the inspection of the learned functions, it is worth highlighting that a preliminary exploratory analysis with KANs provided useful insights for feature engineering and helped validate our assumptions about the structure of the dataset.

In early experiments using the raw angular variables, the learned edge functions frequently exhibited clear sinusoidal patterns. This behavior became particularly evident after applying the symbolic regression (symbolification) procedure available in the KAN framework, which approximates the learned spline functions with analytical expressions (as discussed in Subsection 5.4.6). In several cases, the resulting symbolic forms closely resembled sine and cosine functions of the input angles.

This observation illustrates a practical advantage of interpretable models: by inspecting the learned functional relationships, it becomes possible to uncover meaningful transformations of the input variables. In our case, the emergence of sinusoidal structures confirmed that the spacecraft orientation variables exhibit an intrinsic periodic behavior.

Motivated by this result, we explicitly incorporated this prior knowledge into the preprocessing stage by transforming the angular features using sine and cosine encoding, as described in Subsection 2.1.4. This representation makes the periodic structure explicit and removes artificial discontinuities at the $\pm 180^\circ$ boundary.

After this transformation, the symbolic forms obtained from the learned functions remain trigonometric but become simpler and more stable, since the model no longer needs to reconstruct periodicity from raw angles and can instead focus on modeling the physically meaningful dependencies between spacecraft geometry and detector background.

5.4.2 Function Inspection

As discussed previously, functional representation provided by KANs allows direct inspection of the learned relationships between orbital variables and detector background.

1 s time resolution. Figures 5.4 and 5.5 show zoomed views of representative univariate functions learned by the KAN model trained on the $1\,\mathrm{s}$ time-resolution dataset. Each panel corresponds to the transformation associated with a single edge connecting an input variable to a hidden neuron.

In particular, the figures highlight the learned transformations associated with the $\cos(\mathrm{Earth}_{\mathrm{lon}})$ and $\cos(\mathrm{z}_{\mathrm{lon\_res}})$ features. These variables encode the orbital longitude of the spacecraft and the orientation of the detector boresight, respectively, both of which are expected to influence the background signal due to the changing spacecraft geometry and viewing conditions along the orbit.

Figure 5.4: Zoomed examples of learned univariate functions between the $\cos(\mathrm{Earth}_{\mathrm{lon}})$ input feature and hidden neurons for the KAN model trained at $1\,\mathrm{s}$ time resolution.

Several functions exhibit smooth nonlinear behaviors, including monotonic trends and gradual saturation patterns. Such transformations indicate that hidden neurons respond in a structured way to variations of the corresponding input variables. For instance, *monotonic* responses suggest that the background level changes systematically with the spacecraft orbital longitude or detector orientation.

Figure 5.5: Zoomed examples of learned univariate functions between the $\cos(z_{\text{lon\_res}})$ input feature and hidden neurons for the KAN model trained at $1\,\mathrm{s}$ time resolution.

More generally, the learned functions remain smooth and well-behaved, reflecting the regularization effect of the B-spline parameterization used in the KAN architecture. This functional structure allows the model to capture nonlinear dependencies while preserving interpretability, since the learned relationships between input variables and hidden representations can be directly visualized and analyzed.

50 ms time resolution. To verify the stability of the learned relationships across temporal resolutions, we repeat the same inspection on the KAN model trained at 50 ms resolution.

Figure 5.6 shows representative univariate transformations learned between the input layer and the hidden layer. Each small panel corresponds to the function associated with a specific edge of the network. Despite the higher level of stochastic noise present at this temporal resolution, the learned transformations remain smooth and structured.

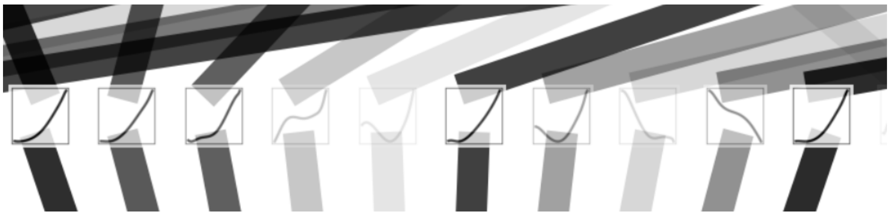

Figure 5.6: Representative univariate functions learned by the KAN model trained at 50 ms resolution between the input feature $\sin(\mathrm{Earth}_{\mathrm{lon}})$ and a hidden neuron. Each panel shows the functional transformation associated with a single edge connecting an input feature to a hidden neuron.

In particular, the first function exhibits a clear monotonic increasing pattern. This function corresponds to the $\sin(\mathrm{Earth}_{\mathrm{lon}})$ feature, indicating that variations in the spacecraft orbital longitude systematically affect the activation of the corresponding hidden neuron. This behavior is consistent with the expected dependence of the detector background on the satellite position along its orbit.

Overall, the similarity between the learned transformations at 1 s and 50 ms resolutions confirms that the KAN model consistently captures the underlying physical relationships between spacecraft geometry and detector background, independently of the temporal binning used in the training data.

5.4.3 Input $\to$ hidden contribution

To understand how the input variables contribute to the internal representation learned by the network, we visualize the attribution matrix between input features and hidden neurons. Each column corresponds to an input variable, while each row represents a hidden unit.

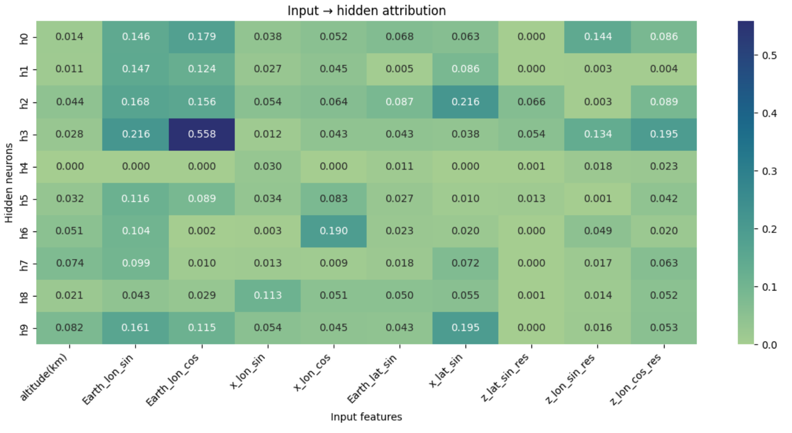

	altitude(km)	Earth_lon_sin	Earth_lon_cos	x_lon_sin	x_lon_cos	Earth_lat_sin	x_lat_sin	z_lat_sin_res	z_lon_sin_res	z_lon_cos_res
h0	0.014	0.146	0.179	0.038	0.052	0.068	0.063	0.000	0.144	0.086
h1	0.011	0.147	0.124	0.027	0.045	0.005	0.086	0.000	0.003	0.004
h2	0.044	0.168	0.156	0.054	0.064	0.087	0.216	0.066	0.003	0.089
h3	0.028	0.216	0.558	0.012	0.043	0.043	0.038	0.054	0.134	0.195
h4	0.000	0.000	0.000	0.030	0.000	0.011	0.000	0.001	0.018	0.023
h5	0.032	0.116	0.089	0.034	0.083	0.027	0.010	0.013	0.001	0.042
h6	0.051	0.104	0.002	0.003	0.190	0.023	0.020	0.000	0.049	0.020
h7	0.074	0.099	0.010	0.013	0.009	0.018	0.072	0.000	0.017	0.063
h8	0.021	0.043	0.029	0.113	0.051	0.050	0.055	0.001	0.014	0.052
h9	0.082	0.161	0.115	0.054	0.045	0.043	0.195	0.000	0.016	0.053

Figure 5.7: Input $\to$ hidden attribution heatmap for the pruned KAN model. Each entry indicates the relative contribution of an input feature to a hidden neuron.

Figure 5.7 shows that the internal representation of the network is primarily driven by a small subset of geometrical variables related to the spacecraft orbital configuration. In particular, the sine and cosine encodings of the Earth longitude ($\sin(\mathrm{Earth}_{\mathrm{lon}})$ and $\cos(\mathrm{Earth}_{\mathrm{lon}})$) exhibit consistently strong contributions across several hidden neurons. This indicates that the orbital phase of the spacecraft relative to the Earth is a dominant factor in shaping the learned latent representation of the background signal.

Some hidden neurons are strongly specialized to particular input variables. For example, unit h_3 is largely dominated by $\cos(\mathrm{Earth}_{\mathrm{lon}})$, suggesting that it captures a specific periodic modulation associated with the spacecraft orbital position. Other hidden units combine multiple inputs with moderate contributions, indicating that they encode more distributed latent factors.

Conversely, some neurons (such as h_4) receive only weak contributions from most inputs, reflecting the effect of the pruning step which suppresses unnecessary connections and leaves only the most relevant functional dependencies active.

Overall, the heatmap confirms that the latent representation learned by the pruned KAN is largely organized around orbital geometry variables, which are physically expected to drive the modulation of the detector background.

5.4.4 Hidden $\to$ output contribution

To analyze how the latent representation is mapped to the detector outputs, we visualize the attribution matrix between hidden neurons and detectors.

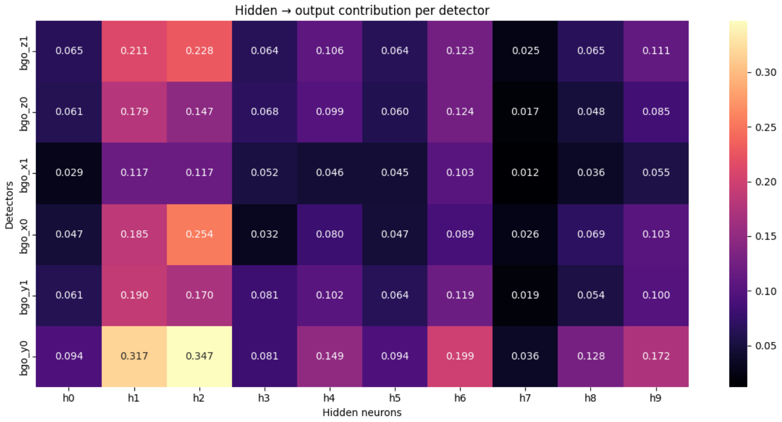

Detectors	h0	h1	h2	h3	h4	h5	h6	h7	h8	h9
bgo_z1	0.065	0.211	0.228	0.064	0.106	0.064	0.123	0.025	0.065	0.111
bgo_z0	0.061	0.179	0.147	0.068	0.099	0.060	0.124	0.017	0.048	0.085
bgo_x1	0.029	0.117	0.117	0.052	0.046	0.045	0.103	0.012	0.036	0.055
bgo_x0	0.047	0.185	0.254	0.032	0.080	0.047	0.089	0.026	0.069	0.103
bgo_y1	0.061	0.190	0.170	0.081	0.102	0.064	0.119	0.019	0.054	0.100
bgo_y0	0.094	0.317	0.347	0.081	0.149	0.094	0.199	0.036	0.128	0.172

Figure 5.8: Hidden $\to$ output attribution heatmap for the pruned multi-output KAN.

Figure 5.8 shows that the detector predictions are mainly driven by a limited subset of hidden units. In particular, neurons h_2 and h_1 consistently exhibit the strongest contributions across detectors, indicating that they encode dominant latent components of the background variability shared among channels.

Additional units contribute with smaller magnitudes, suggesting that they capture secondary corrections to the global background structure. The overall similarity of the contribution profiles across detectors indicates that the model relies on a common latent representation, while detector-specific variations arise from differences in how each output combines these hidden components.

5.4.5 Global feature importance

Path-wise aggregation of feature attributions. To obtain a single global ranking of the input variables that reflects their *effective* contribution to the predicted background, we combine the two attribution stages of the pruned KAN. First, we compute the absolute attributions from each input feature to each hidden unit (input → hidden). Second, we compute the absolute attributions from each hidden unit to each detector output (hidden → output). A global score is then derived by aggregating contributions along all input–hidden–output paths, so that an input feature is deemed important only if it both (i) strongly influences at least one hidden unit and (ii) such hidden unit contributes to at least one detector output (Equation 5.3).

This path-based aggregation is particularly *natural* in the KAN architecture, where the network is explicitly composed of univariate functions along individual edges. In contrast, in standard MLPs the nonlinear transformations entangle multiple input variables, making such a direct path-wise attribution less straightforward.

$$I_j \propto \sum_h \left| A_{h,j}^{(in\to hid)} \right| \sum_d \left| A_{d,h}^{(hid\to out)} \right| \tag{5.3}$$

The resulting scores are normalized to sum to one, yielding an interpretable global importance distribution.

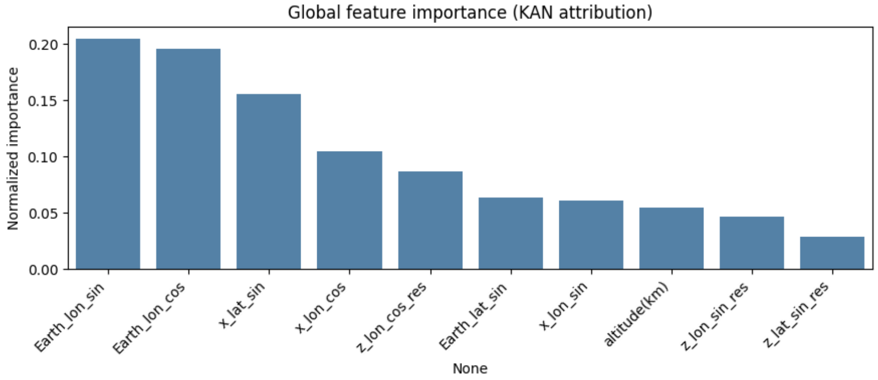

Figure 5.9: Global feature importance obtained by composing input→hidden and hidden→output attributions in the pruned KAN.

Interpretation of the global importance ranking. As shown in Figure 5.9, the dominant contributions are associated with $\sin(\mathrm{Earth}_{\mathrm{lon}})$ and $\cos(\mathrm{Earth}_{\mathrm{lon}})$, indicating that the spacecraft orbital longitude is the primary driver of the background variability captured by the model. This confirms that the background signal strongly depends on the satellite position along its orbit, which modulates the detector exposure to atmospheric albedo and other environmental radiation sources.

A second tier of influential variables includes $\sin(\mathrm{x}_{\mathrm{lat}})$ and $\cos(\mathrm{x}_{\mathrm{lon}})$, which encode the spacecraft pointing orientation. Their relatively high importance suggests that the detector background is also modulated by the instrument attitude, as different pointing configurations expose the detectors to varying levels of Earth-originating gamma radiation.

Additional variables such as $\cos(\mathrm{z}_{\mathrm{lon\_res}})$, $\sin(\mathrm{Earth}_{\mathrm{lat}})$, and $\sin(\mathrm{x}_{\mathrm{lon}})$ provide secondary contributions that refine the prediction by capturing more subtle geometric effects. Conversely, $\mathrm{altitude}$ and $\sin(\mathrm{z}_{\mathrm{lat\_res}})$ appear among the least influential inputs, indicating that within the analyzed orbital regime their contribution to background variability is comparatively weaker than the orbital-phase and pointing-related factors.

5.4.6 Symbolic approximation of learned dependencies

Although the composed input→hidden→output attribution score is not directly amenable to a compact closed-form expression, symbolic regression can still be used as a diagnostic tool to summarize the dominant *functional families* learned by the pruned KAN.

Symbolic approximation of edge functions. As a first step, the symbolification procedure was applied to the individual edge functions between the input and hidden layers. Figure 5.10 shows the symbolic approximations of the input→hidden functions previously visualized in Figure 5.4. The expressions are obtained through the KAN symbolification procedure applied to the model trained at 1 s resolution. These symbolic forms approximate the spline functions learned by the network and provide an interpretable description of how specific input variables influence the hidden representation.

Figure 5.10: Symbolic approximation of selected input→hidden functions obtained through the KAN symbolification procedure for the 1 s model. Each panel corresponds to a learned transformation between an input feature and a hidden neuron.

These symbolic forms highlight that the dominant transformations involve combinations of trigonometric, linear, and weakly nonlinear terms, reflecting the periodic and geometric structure of the spacecraft orientation variables.

Symbolic approximation of the detector response. While the figure illustrates the symbolic expressions for individual edge functions, the final detector prediction results from the composition of multiple input→hidden and hidden→output transformations.

As an illustrative example, the following expression reports a symbolic approximation of the overall mapping for the $\mathrm{bgo}_{\mathrm{z1}}$ detector output[1]:

$$\begin{aligned}\hat{y}_{\mathrm{bgo}_{\mathrm{z1}}}(x) \approx a_0 &+ 0.16\,\mathrm{z}_{\mathrm{lon}}^{\mathrm{cos,res}} + 0.15\,\mathrm{x}_{\mathrm{lon}}^{\mathrm{sin}} - 0.25\,\mathrm{z}_{\mathrm{lon}}^{\mathrm{sin,res}} \\ &+ 0.53\sin(0.41\,\mathrm{z}_{\mathrm{lon}}^{\mathrm{cos,res}} - 1.21) - 1.53\cos(0.24\,\mathrm{z}_{\mathrm{lon}}^{\mathrm{cos,res}} + 3.43) \\ &- 1.33\cos(0.83\,\mathrm{Earth}_{\mathrm{lon}}^{\mathrm{sin}} + 7.03) + 0.60\cos(0.59\,\mathrm{x}_{\mathrm{lon}}^{\mathrm{sin}} - 5.62) \\ &+ 0.16\exp(1.01\,\mathrm{x}_{\mathrm{lon}}^{\mathrm{sin}}) + 0.15\left(0.79\,\mathrm{x}_{\mathrm{lat}}^{\mathrm{sin}} + 0.99\right)^2 .\end{aligned} \tag{5.4}$$

For this detector, the symbolic approximation reveals a combination of linear, periodic, and weakly curved nonlinear terms. The most evident structure is the repeated appearance of trigonometric functions involving $z_{\mathrm{lon}}^{\mathrm{cos,res}}$ and $\mathrm{Earth}_{\mathrm{lon}}^{\mathrm{sin}}$, confirming that the predicted background strongly depends on orbital phase and pointing geometry. Additional exponential and quadratic terms involving $\mathrm{x}_{\mathrm{lon}}^{\mathrm{sin}}$ and $\mathrm{x}_{\mathrm{lat}}^{\mathrm{sin}}$ indicate that the dependence is not purely sinusoidal and includes mild curvature effects.

For clarity of presentation, only the dominant terms are reported, while very large low-frequency trigonometric terms are omitted from the displayed expression, as they mainly behave as quasi-constant offset corrections in the sampled regime.

---

[1] For readability, the notation of the input variables is slightly simplified with respect to the original feature names used in the model in order to improve clarity and layout of the symbolic expression.

5.4.7 Consistency with SHAP analysis

The global feature importance obtained from the KAN attribution analysis shows a strong consistency with the SHAP-based interpretation presented in the previous section (4.1.7). In particular, the dominant role of the orbital phase variable ($\mathrm{Earth}_{\mathrm{lon}}$) and the importance of spacecraft orientation parameters are recovered by both approaches.

While the SHAP analysis was performed on the raw orientation variables, the KAN model operates on trigonometric encodings of the same quantities (e.g. $\sin$ and $\cos$ of the angles). Despite this difference in representation, both attribution methods identify the *same physical drivers* of the background variability.

This agreement between two independent interpretability frameworks provides an additional consistency check, suggesting that the models capture meaningful physical dependencies related to orbital position and spacecraft pointing.

Chapter 6

Anomaly Detection

6.1 Transient Detection Strategy

6.1.1 Statistical formulation

The goal of the anomaly detection pipeline is to identify time intervals in which the observed photon counts significantly exceed the background expectation predicted by the probabilistic regression model.

For each time bin t, the model provides a conditional probability distribution for the background counts:

$$Y_t \sim P(Y \mid X_t), \tag{6.1}$$

where X_t represents the spacecraft orbital and attitude parameters at time t, and Y_t models the stochastic background fluctuations.

Under the background-only hypothesis,

$$H_0 : y_t \sim P(Y \mid X_t), \tag{6.2}$$

the observed counts are assumed to be generated purely by the background process. A transient astrophysical event, such as a GRB, manifests itself as a statistically significant positive deviation from this distribution.

The detection problem can therefore be formulated as a statistical compatibility test between the observed counts y_t and the predicted background distribution. Observations that fall in the upper tail of the predicted distribution are considered statistically incompatible with the background-only hypothesis at a given confidence level. In particular, deviations exceeding predefined significance thresholds (e.g., 3σ, 4σ, or 5σ) are treated as candidate anomalous fluctuations.

Because isolated fluctuations may arise from random background variations, the detection decision is not taken at the single-bin level. Instead, candidate anomalous bins are aggregated over short temporal windows in order to identify statistically consistent excesses above the predicted background level.

6.1.2 Implementation in the proposed pipeline

The probabilistic regression model provides, for each time bin t, the full predictive distribution of the expected background counts. This distribution is used to compute significance thresholds corresponding to rare background fluctuations.

Adaptive significance thresholds

In particular, we derive thresholds corresponding to Gaussian-equivalent significance levels of 3σ, 4σ, and 5σ. These thresholds are obtained directly from the predicted background distribution $P(Y \mid X_t)$.

For a given significance level $k\sigma$, the corresponding threshold $T_{k\sigma}(t)$ is defined as the upper quantile of the predictive distribution such that

$$P(Y_t > T_{k\sigma}(t) \mid X_t) = p_{k\sigma}, \tag{6.3}$$

where $p_{k\sigma}$ denotes the tail probability associated with the chosen Gaussian significance level.

An observed count y_t is therefore flagged as anomalous whenever

$$y_t > T_{k\sigma}(t). \tag{6.4}$$

To formalize this condition, we define a binary indicator variable

$$I_t = \begin{cases} 1 & \text{if } y_t > T_{k\sigma}(t) \\ 0 & \text{otherwise} \end{cases} \tag{6.5}$$

which marks whether a time bin exceeds the chosen significance threshold.

Temporal aggregation of anomalies

Since the temporal structure of gamma-ray bursts can vary significantly, the detection decision is not taken at the single-bin level. Isolated fluctuations may arise from statistical variations of the background and are therefore not considered sufficient evidence for a transient event.

To increase robustness against random background fluctuations and reduce the number of *false positives*, anomalous bins are aggregated over short temporal windows.

We therefore define an aggregated detection statistic

$$S = \sum_{t=1}^{W} I_t, \tag{6.6}$$

where W denotes the size of the analysis window.

A transient candidate is identified whenever the number of bins above threshold exceeds a minimum requirement τ:

$$S \geq \tau. \tag{6.7}$$

In other words, at least τ time bins within the analysis window must exceed the chosen significance threshold in order for the event to be classified as a candidate anomaly.

In this proof-of-concept study we set $\tau = 3$, meaning that at least three bins within the analysis window must exceed the threshold. Requiring multiple threshold crossings ensures temporal consistency of the signal and reduces the impact of isolated background fluctuations while preserving sensitivity to short-duration transients.

Multi-detector consistency

Since the COSI anticoincidence system consists of multiple detector panels with different orientations, an astrophysical transient like a GRB is expected to produce a coherent excess in more than one detector.

For this reason, we require that a candidate transient is observed simultaneously in *at least two* detectors.

This condition introduces an additional robustness criterion. Requiring a multi-detector detection significantly reduces the number of false positives generated by statistical background fluctuations affecting a single detector.

However, this constraint also introduces a *trade-off*. Depending on the burst geometry and detector orientation, not all detectors may observe the transient with the same intensity. As a consequence, requiring multiple detectors to exceed the threshold may reduce sensitivity to weak or highly directional events.

The chosen criterion therefore balances two competing objectives: reducing false alarms while maintaining sensitivity to authentic astrophysical transients.

Calibration-based threshold selection

We also implemented a calibration-based framework in which the parameter τ can be selected using a background-only calibration set by estimating the *empirical tail* distribution of the detection statistic.

This procedure enables a data-driven control of the false-alarm probability of the detection pipeline. However, since the present work focuses on demonstrating the feasibility of the method, this calibration step is not applied in the current proof-of-concept analysis and is left for future work.

6.2 Transient Data Source

To evaluate the detection pipeline, synthetic GRB events were injected into the background data using a simulation framework based on MEGAlib and the Fermi Gamma-Ray Burst Monitor (Fermi/GBM; [33]) spectral models [13].

MEGAlib The Medium Energy Gamma-ray Astronomy Library (MEGAlib; [50]) is a software framework developed for the simulation and analysis of medium-energy gamma-ray instruments. It provides tools for modeling detector geometries, simulating photon interactions through the Geant4 toolkit [2], and generating realistic detector responses for astrophysical sources.

In this work we exploit simulated GRBs produced as part of the work presented in [13]. In that study, transient gamma-ray events were simulated using a framework based on MEGAlib together with spectral models derived from Fermi/GBM observations.

These simulated detector count rates are used in the present work as transient signals that can be summed to the background light curves for the anomaly detection experiments.

6.3 Results

Anomaly detection experiment The anomaly detection experiment presented in this work evaluates the proposed background modeling framework in the context of GRB detection. Since only a *limited* number of GRB events are considered, the analysis focuses on a reduced set of representative bursts rather than on a large-scale statistical benchmark.

The objective of this section is not to perform a comprehensive study of a GRB catalog, but rather to illustrate how the probabilistic background model can be used to identify transient excesses compatible with GRBs.

For this purpose, synthetic GRB signals are summed to the background light curves and the response of the detection pipeline is evaluated.

This experimental setup therefore serves as a *preliminary evaluation* of the anomaly detection procedure based on the predicted background distributions. A complete assessment of the detection performance would require the analysis of a much larger and more diverse GRB sample, including events with different durations, intensities, and temporal structures, as well as a dedicated study of the background-only false alarm rate. Such an investigation is beyond the scope of the present work and is left for future studies.

Injection pipeline. Simulated GRB signals are summed directly to the background light curves by adding the simulated photon counts to the original background data. This approach preserves the statistical properties of the real background while introducing controlled transient events.

To make the detection task more challenging and evaluate the model under more demanding conditions, the simulated GRB counts are *scaled* before injection. In particular, a constant scaling factor of 0.2 is applied so that the injected transients remain in a relatively low signal-to-background regime.

GRB events are injected at *random* times and durations within the observation windows.

Detection results. Detection is performed by monitoring positive deviations from the model-predicted background distribution. When the observed counts exceed the significance thresholds (e.g. 3σ, 4σ, or 5σ), the corresponding time bins are flagged as anomalous.

Figure 6.1 shows an example of the detection process. In this case, the model correctly identifies a localized positive deviation from the predicted background. The observed excess spans multiple consecutive bins above the 3σ threshold.

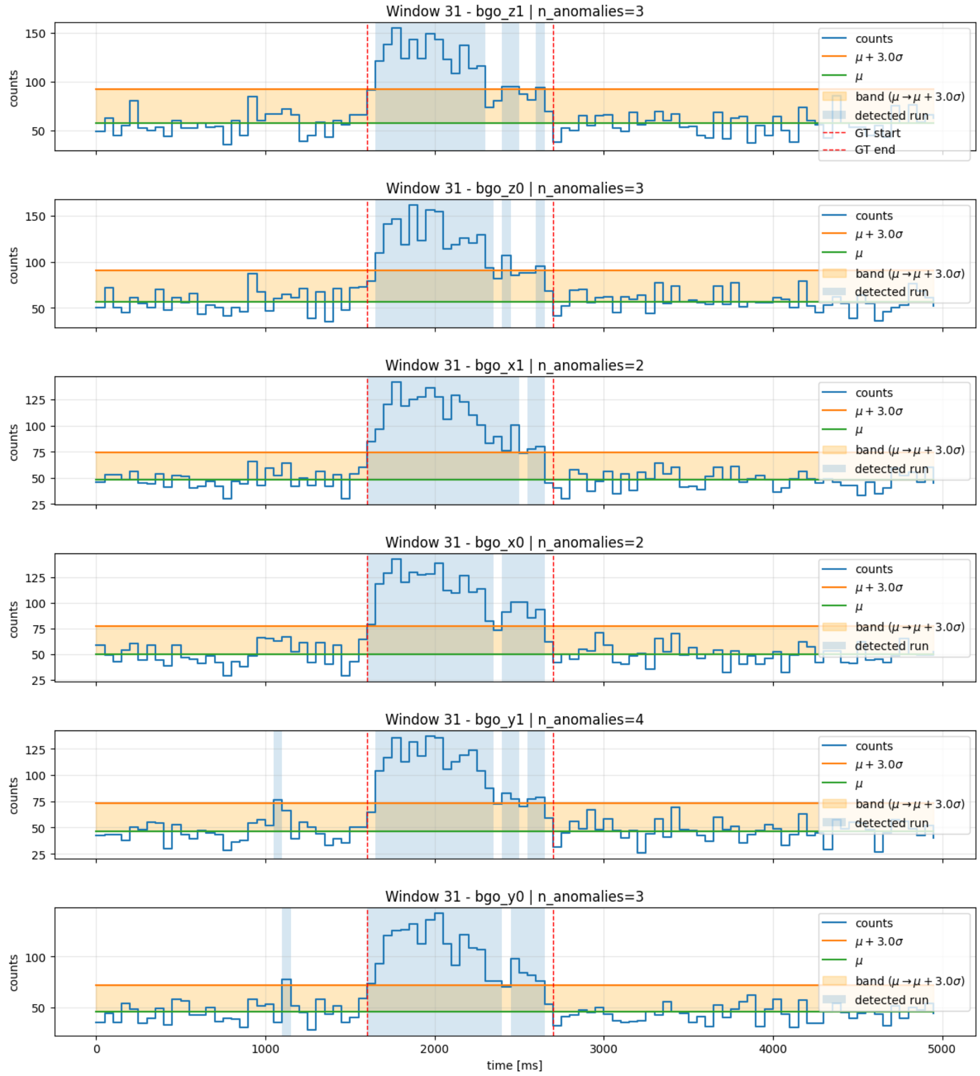
Window 31 - bgo_z1 | n_anomalies=3

Window 31 - bgo_z0 | n_anomalies=3

Window 31 - bgo_x1 | n_anomalies=2

Window 31 - bgo_x0 | n_anomalies=2

Window 31 - bgo_y1 | n_anomalies=4

Window 31 - bgo_y0 | n_anomalies=3

Figure 6.1: Example of a GRB detection task over an anomalous window.

Within our pipeline, the detection statistic is defined as the number of time bins exceeding the chosen threshold inside the analysis window. In the example shown, three bins exceed the 3σ threshold, producing a clear detection of the injected transient.

In other words, detection is performed by monitoring positive deviations from the model-predicted background distribution and applying the window-based statistic defined in the previous section.

These results demonstrate that the probabilistic background model can be effectively used to detect transient excesses without relying on explicit signal templates, providing a flexible anomaly detection framework for GRB searches.

Chapter 7

Conclusion and Future Work

7.1 Conclusions

This thesis investigated the problem of *background modeling* for the COSI Anticoincidence System (ACS), with the objective of supporting the detection of transient gamma-ray events such as *Gamma-Ray Bursts* (GRBs). Accurate background estimation represents a critical challenge for space-based gamma-ray instrumentation, as the measured count rates are strongly influenced by the spacecraft orbital position and attitude. The resulting background signal is therefore highly *non-stationary* and detector-dependent, making classical modeling approaches insufficient for robust transient detection.

In this work, we proposed a machine learning framework that combines *probabilistic regression* models with interpretable analysis techniques in order to both accurately model the ACS background and provide physical insight into the learned relationships between spacecraft parameters and detector response.

The first contribution of this work concerns the development of a probabilistic background modeling approach. Instead of predicting only a deterministic background estimate, the proposed models predict the full conditional distribution of the detector counts given the spacecraft state. This formulation allows the direct estimation of uncertainty and enables the definition of statistically meaningful detection thresholds. Several probabilistic models were investigated, including Skew-Normal regression and Gaussian Mixture Regression. These models were trained using the negative log-likelihood loss, allowing them to capture non-Gaussian features of the background count distribution.

The experimental analysis demonstrated that simple Poisson-based assumptions are not sufficient to describe the statistical properties of the ACS background. In particular, the data exhibit strong overdispersion and complex variability driven by orbital and geometrical factors. The probabilistic regressors were able to model these effects more effectively, producing predictive distributions that capture the observed variability of the detector signals.

A second key aspect of the work focused on *model explainability* and *interpretability*. Since background variations are driven by physical processes related to spacecraft orientation and orbital position, understanding the learned model behavior is essential to validate its reliability. To address this aspect, two complementary interpretability strategies were adopted.

First, post-hoc explanations were obtained using *SHAP*, which quantify the contribution of each input feature to the predicted background counts. The SHAP analysis revealed consistent physical patterns across detectors, showing that spacecraft pointing directions and orbital phase variables are the dominant drivers of the background signal, while altitude plays a secondary role. These results are consistent with known physical effects such as Earth albedo and the directional dependence of detector exposure.

Second, intrinsically interpretable models based on *Kolmogorov–Arnold Networks* (KANs) were explored. Unlike standard neural networks, KANs represent multivariate mappings through compositions of learnable univariate functions. This architecture allows direct inspection of the learned functional relationships between input variables and detector responses. The analysis of the learned functions provided additional insight into how orbital and orientation parameters influence the background behavior, demonstrating the potential of KANs as interpretable models for scientific machine learning applications.

Finally, the probabilistic background model was integrated into a transient *detection strategy*. The predicted background distributions were used to compute adaptive significance thresholds corresponding to rare background fluctuations. Observed count rates exceeding these thresholds were flagged as potential anomalies. To evaluate the detection capability of the proposed approach, simulated GRB signals generated with MEGAlib were summed to the background data. These experiments demonstrated that the probabilistic framework is able to identify transient excesses while accounting for the intrinsic variability of the background signal.

Overall, the results of this thesis show that combining probabilistic regression with interpretable machine learning techniques provides a powerful framework for modeling complex detector backgrounds in space missions. The approach enables not only accurate prediction of the background signal but also a deeper understanding of the physical drivers that shape its variability.

7.2 Future Work

Several directions can extend and improve the methodology proposed in this thesis.

A first natural extension concerns the use of more advanced probabilistic models capable of capturing even richer background dynamics. For instance, hierarchical Bayesian models or deep probabilistic architectures could provide more flexible representations of detector variability and potentially improve predictive calibration across different orbital regimes.

Another promising direction involves the integration of temporal modeling. In the present work, each time bin was treated independently, with the background predicted solely from instantaneous spacecraft parameters. However, detector background signals may exhibit temporal correlations driven by slow orbital variations, activation effects, or changes in the particle environment. Incorporating sequence-based models such as recurrent neural networks, temporal convolutional networks, state-space models, or attention-based architectures (e.g., 1D transformers) could allow the capture of these temporal dependencies and potentially improve the stability of background predictions over extended orbital segments.

Another possible extension involves incorporating physics-informed constraints or hybrid models (PINNs) combining data-driven regression with simplified physical background models.

From the interpretability perspective, future research could explore tighter integration between interpretable models and probabilistic regression. In particular, extending KANs to probabilistic outputs could combine the advantages of interpretable functional representations with uncertainty-aware predictions.

Another important future step concerns the application of the proposed framework to real mission operations. The current study was conducted primarily using simulated data and simulated GRB signals. Once observational data from the COSI mission become available, the models could be retrained and validated directly on real measurements, enabling a systematic evaluation of the detection performance under realistic conditions.

Finally, the methodology developed in this thesis is not limited to the COSI ACS and could be generalized to other high-energy astrophysics instruments. Many gamma-ray missions face similar challenges related to background variability and transient detection. The combination of probabilistic regression, interpretable machine learning, and adaptive statistical thresholds may therefore represent a general framework for background modeling and anomaly detection in space-based detectors.

In conclusion, this work highlights how modern machine learning techniques can contribute to addressing long-standing challenges in high-energy astrophysics. By combining statistical modeling, interpretability, and physical reasoning, it is possible to develop data-driven approaches that remain transparent and scientifically meaningful, ultimately supporting the discovery of transient phenomena in the high-energy universe.

Appendix A

Distribution Fitting Analysis

Purpose of this appendix

This appendix complements the distributional analysis presented in Section 2.2. The goal of the analysis reported here is not to select the final regression model, which instead estimates the conditional distribution $P(C \mid x)$, but rather to characterize the empirical properties of the observed count data.

Since the conditional distribution $P(C \mid x)$ cannot be directly observed in isolation, we first analyze the marginal distribution $P(C)$ obtained by aggregating counts across time bins. This exploratory analysis provides insight into the statistical properties of the background fluctuations (such as dispersion and skewness) and helps identify suitable distributional families for modeling the conditional background variability in the regression framework.

A.1 Selection of approximately stationary background windows

To further extend the analysis introduced in Section 2.2, it is important to emphasize that the background count rate recorded by the detectors is not strictly stationary over long time intervals. Orbital motion, Earth occultation, geomagnetic conditions, and variations in the spacecraft orientation introduce gradual changes in the expected photon rate.

In order to study the statistical properties of the background counts, we therefore identify time intervals in which the count rate can be considered approximately stationary. This allows us to analyze the intrinsic stochastic fluctuations of the background without mixing different physical regimes.

To this end, the time series is divided into consecutive windows of fixed duration. For each window we compute the dispersion index

$$\phi = \frac{\mathrm{Var}(C)}{E[C]},$$

where the mean and variance are estimated across all detectors within the window.

For a pure Poisson process the dispersion index satisfies $\phi = 1$, while values $\phi > 1$ indicate overdispersion. We therefore select the window with the smallest average dispersion index across detectors, corresponding to the interval with the most statistically stable background behavior.

This procedure mitigates the impact of slow orbital trends and allows us to isolate time intervals in which the stochastic background fluctuations can be meaningfully analyzed.

A.2 Empirical distribution of background counts

Once the most stationary time window has been identified, we analyze the empirical distribution of the photon counts recorded in that interval.

Figure A.1 shows the histogram of counts for one representative detector together with a kernel density estimate (KDE) and several parametric distributions fitted to the data.

The empirical distribution exhibits two notable features:

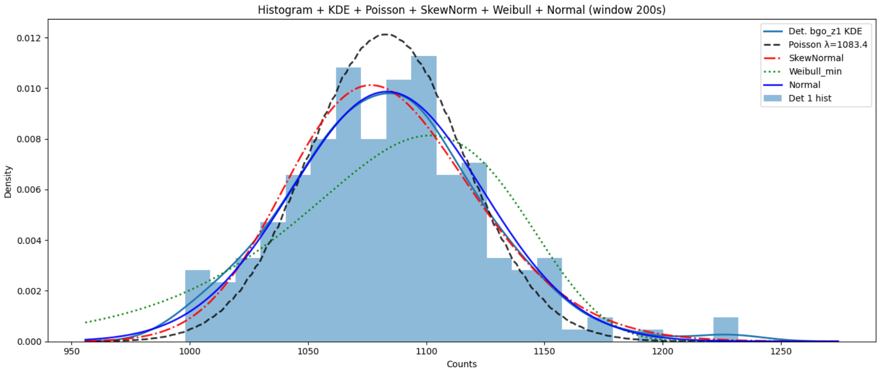

Figure A.1: Histogram of background counts for a representative detector (bgo_{z1}) in the selected stationary window of 200 seconds. Several parametric distributions are fitted to the empirical density, including Poisson, Normal, Skew-Normal and Weibull models.

- The variance of the counts is larger than the mean, indicating mild overdispersion relative to the Poisson model.
- The distribution shows a slight positive skewness and heavier tails compared to a Gaussian approximation.

As a consequence, the Poisson distribution tends to underestimate the spread of the data, while more flexible distributions provide a better approximation of the empirical density.

A.3 Stability of the count distribution across consecutive windows

To verify that the selected interval corresponds to a stable background regime, we inspect the count distributions obtained from several consecutive windows.

Figure A.2 shows kernel density estimates of the counts for multiple consecutive windows for each detector, together with the Poisson distribution corresponding to the empirical mean.

The KDE curves corresponding to different windows largely overlap, indicating that the background distribution is relatively stable within the selected

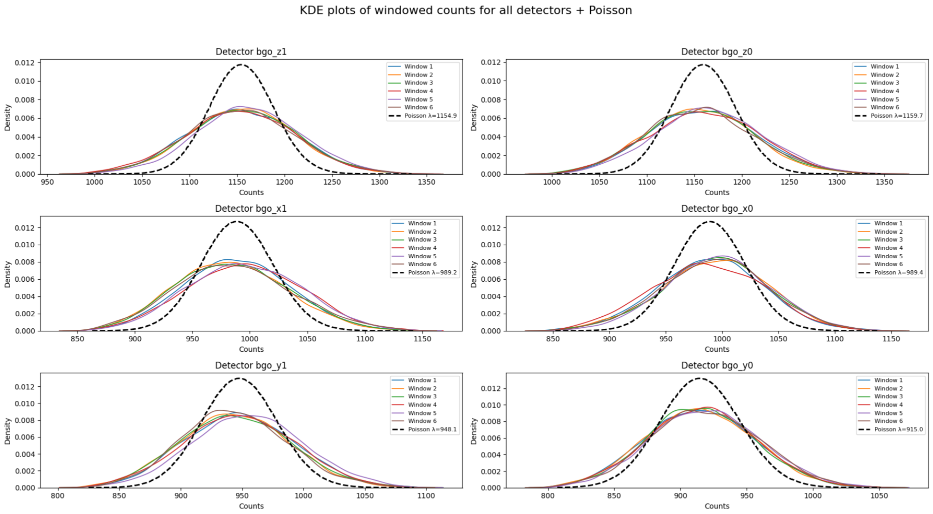

Figure A.2: Kernel density estimates of the count distributions for consecutive 200 s time windows for each detector. The dashed curve represents the Poisson distribution corresponding to the empirical mean.

interval. However, the empirical distributions consistently appear wider than the corresponding Poisson model, confirming the presence of mild overdispersion.

A.4 Parametric distribution fitting

In order to quantitatively compare different distributional models, several parametric families were fitted to the empirical count distributions. The considered distributions include Poisson, Normal, Skew-Normal, Lognormal, and Weibull models.

Model comparison was performed using the Akaike Information Criterion (AIC, [3]), defined as

$$\mathrm{AIC} = 2k - 2\log L,$$

where k is the number of parameters and L is the maximized likelihood.

Table A.1 reports the best-fitting distributions for each detector.

The Weibull distribution frequently provides the best fit among the tested

Detector	Best distribution	AIC
bgo_z1	lognormal	2048.5
bgo_z0	Weibull	2069.6
bgo_x1	Weibull	1984.4
bgo_x0	Weibull	1988.9
bgo_y1	Weibull	1965.8
bgo_y0	Weibull	2004.3

Table A.1: Best-fitting parametric distributions for each detector according to the AIC criterion.

families. This reflects its flexibility in modeling asymmetric distributions with heavier tails and its ability to approximate the aggregated distribution of counts.

A.5 Interpretation of the marginal count distribution

It is important to emphasize that the distributions analyzed in this appendix correspond to the marginal distribution of the counts,

$$P(C),$$

which aggregates observations acquired under different spacecraft configurations.

In reality, the background count rate depends strongly on the satellite orbital state, orientation, and geomagnetic environment. Each time bin is therefore generated under a different system configuration x, leading to a conditional distribution

$$P(C \mid x).$$

When counts from many time bins are aggregated, the observed marginal distribution can be interpreted as a mixture of such conditional distributions,

$$P(C) = \int P(C \mid x) P(x) \, dx,$$

where $P(x)$ represents the distribution of spacecraft configurations sampled during the observation period.

As a consequence, flexible distributions such as the Weibull may provide a good approximation of the aggregated data even if they do not correspond to the true generative process for individual time bins.

A.6 Implications for conditional background modeling

The goal of the regression models considered in this work is not to model the marginal distribution $P(C)$, but rather the conditional distribution

$$P(C \mid x),$$

where x encodes the orbital and geometrical configuration of the spacecraft.

Once conditioning on the system state is introduced, a large fraction of the variability observed in the marginal distribution is explained by changes in the expected count rate. The residual fluctuations around this conditional mean are observed to be approximately symmetric with mild skewness.

For this reason, the skew-normal distribution provides a convenient and flexible model for the conditional background fluctuations. It captures small deviations from normality while remaining numerically stable during regression training.

Consequently, although the Weibull distribution often provides the best fit to the marginal distribution of counts, the skew-normal model was preferred for the conditional probabilistic regression framework adopted in this work. This choice is also supported by quantitative evaluation: when implemented within the regression framework, the skew-normal model achieves slightly better predictive performance than the Weibull alternative, reducing the average percentage error by approximately 2.00 points (3.39% vs 5.50%). Although modest, this improvement suggests that the Skew-Normal distribution more accurately captures the conditional variability of the background counts once the dependence on the spacecraft state is taken into account.

Appendix B

Temporal Structure of the Dataset

B.1 Autocorrelation Analysis

To investigate the temporal structure of the spacecraft telemetry and detector background, we compute the autocorrelation function (ACF) of both the spacecraft state raw variables and the detector counts using the 15 s resolution dataset. The ACF measures the correlation between a signal and time-shifted versions of itself, allowing dominant periodicities in the time series to be identified.

Figure B.1 shows the autocorrelation of several spacecraft orientation and orbital variables. The periodic oscillations reflect the deterministic nature of the spacecraft motion along its orbit. In particular, strong peaks appear at regular lag intervals, indicating that similar geometric configurations recur periodically as the satellite completes successive orbits.

The same analysis applied to the detector count rates (Figure B.2) reveals a similar periodic structure. The first dominant peak occurs at a lag of approximately 320–330 bins. Considering the temporal resolution of $15\,\mathrm{s}$ per bin, this corresponds to a characteristic period of

$$T \approx 330 \times 15\,\mathrm{s} \approx 4950\,\mathrm{s} \approx 82\,\mathrm{minutes}.$$

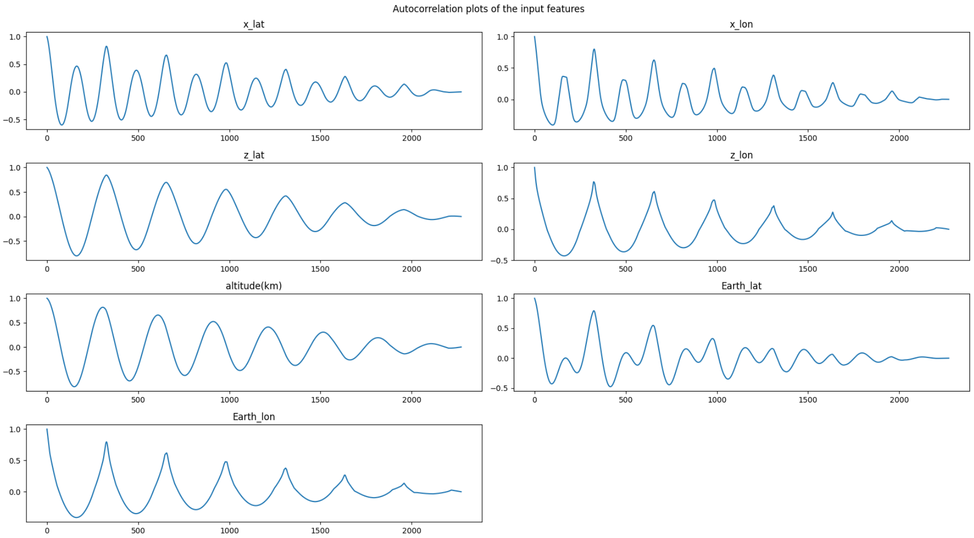

Figure B.1: Autocorrelation functions of representative spacecraft state variables describing the satellite orientation and orbital position. The curves exhibit a clear periodic structure driven by the orbital motion of the spacecraft.

This value is consistent with the orbital period of the spacecraft in low Earth orbit. Additional peaks at larger lags correspond to integer multiples of this fundamental period and arise from the recurrence of similar observational conditions across successive orbital passages.

The presence of this strong periodic component indicates that the detector background is largely modulated by the spacecraft orbital motion and the associated variations in viewing geometry and radiation environment. Consequently, the background signal exhibits a characteristic timescale of approximately 80 minutes.

This orbital periodicity motivates the evaluation strategy adopted in the main analysis, where model predictions are inspected over contiguous temporal segments corresponding to approximately one orbital period. Such windows allow the temporal evolution of the predicted background to be analyzed over a full orbital cycle while preserving the natural temporal correlations present in the data.

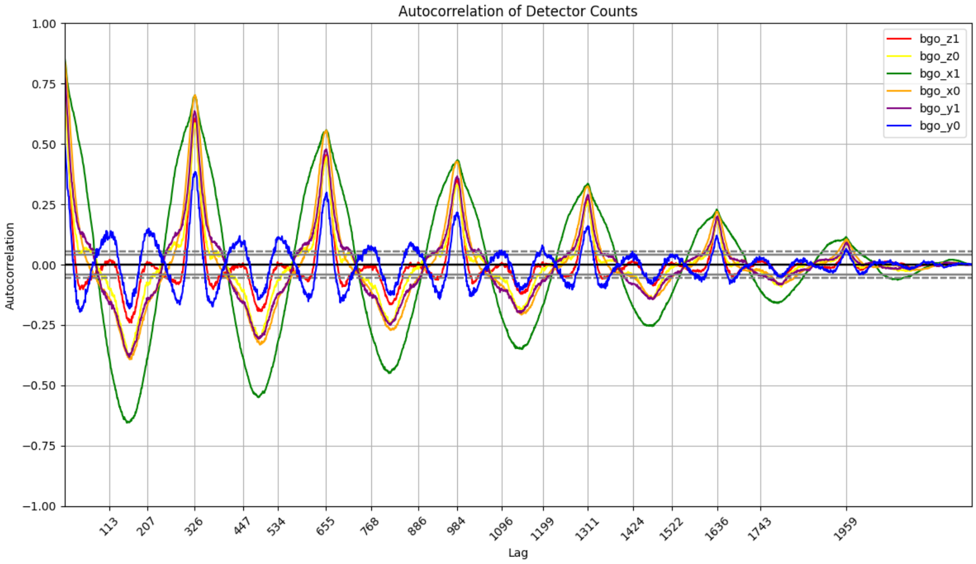

Figure B.2: Autocorrelation functions of the detector count rates for the six ACS detectors at 15 s resolution. Pronounced peaks appear at multiples of the orbital period.

Acknowledgements

I would like to sincerely thank the INAF team for their support and guidance throughout the development of this thesis and for welcoming me and making me feel like a valued member of the group.

In particular, I am deeply grateful to Dr. Nicolò Parmiggiani for his constant supervision, insightful suggestions, and invaluable guidance during all stages of this work. I would also like to thank Dr. Andrea Bulgarelli and Dr. Valentina Fioretti for their advice, encouragement, and support.

Moreover, I would like to express my appreciation to all the other researchers of the team for their help, discussions, and for creating such a stimulating and friendly research environment.

Finally, I would also like to thank Professor Lombardi for his inspiring teaching and outstanding knowledge, which greatly contributed to my academic growth during these years.

I would also like to thank my family and my friends for their continuous support and encouragement throughout this journey. Their presence and understanding have been an important source of balance during the more demanding periods of this work.

The Compton Spectrometer and Imager is a NASA Explorer project led by the University of California, Berkeley, with funding from NASA under contract 80GSFC21C0059. Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center and by National High Performance Computing (NHR) South-West at Johannes Gutenberg University Mainz. This work is supported by the Italian Space Agency, partially within contract ASI/INAF No. 2024-11-HH.0.

Bibliography

[1] Martín Abadi et al. “TensorFlow: A System for Large-Scale Machine Learning”. In: *OSDI* (2016).

[2] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/S0168-9002(03)01368-8`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900203013688`.

[3] Hirotsugu Akaike, BN Petrov, and F Csaki. “Akaike Information theory as an extension of the maximum likelihood principle”. In: *Second international symposium on information theory. Akademiai Kiado, Budapest*. 1973, pp. 267–281.

[4] Vladimir Arnold. “On the representation of functions of several variables as a superposition of functions of a smaller number of variables”. In: Jan. 2009. ISBN: 978-3-642-01741-4. DOI: `10.1007/978-3-642-01742-1_5`.

[5] Vladimir I. Arnold. “On Functions of Three Variables”. In: *Doklady Akademii Nauk SSSR* 114 (1957), pp. 679–681.

[6] Adelchi Azzalini. “A class of distributions which includes the normal ones”. In: *Scandinavian Journal of Statistics* 12.2 (1985), pp. 171–178.

[7] Adelchi Azzalini. *The Skew-Normal and Related Families*. Cambridge University Press, 2013.

[8] Christopher M. Bishop. *Mixture density networks*. English. Working-Paper 4288. Aston University, 1994.

[9] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[10] A. Colin Cameron and Pravin K. Trivedi. *Regression Analysis of Count Data*. 2nd ed. Econometric Society Monographs. Cambridge University Press, 2013.

[11] W. Cash. *Parameter estimation in astronomy through application of the likelihood ratio*. Vol. 228. 1979, pp. 939–947.

[12] François Chollet. *Keras*. `https://keras.io`. 2015.

[13] Alex Ciabattoni et al. "Benchmarking of Geant4 simulations for the COSI Anticoincidence System". In: *Exper. Astron.* 60.1 (2025), p. 9. DOI: `10.1007/s10686-025-10019-7`. arXiv: `2507.21275 [astro-ph.IM]`.

[14] *COSI Anti-Coincidence System Background Dataset*. `https://science.nasa.gov/astrophysics/programs/physics-of-the-cosmos/community/third-cosi-data-challenge-released/`. 2025.

[15] Riccardo Crupi. *Deep Learning for Gamma-Ray Bursts: A data driven event framework for X/Gamma-Ray analysis in space telescopes*. 2024. arXiv: `2401.15632 [astro-ph.HE]`. URL: `https://arxiv.org/abs/2401.15632`.

[16] A. J. Dean, A. J. Bird, N. Diallo, et al. "The Origin of Background in Gamma-Ray Telescopes". In: *Space Science Reviews* 105 (2003), pp. 285–376.

[17] A. P. Dempster, N. M. Laird, and D. B. Rubin. "Maximum Likelihood from Incomplete Data Via the EM Algorithm". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 39.1 (1977), pp. 1–22. DOI: `https://doi.org/10.1111/j.2517-6161.1977.tb01600.x`. eprint: `https://rss.onlinelibrary.wiley.com/doi/pdf/`

10.1111/j.2517-6161.1977.tb01600.x. URL: https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1977.tb01600.x.

[18] Weihao Gao et al. *TabKANet: Tabular Data Modeling with Kolmogorov-Arnold Network and Transformer*. 2024. arXiv: 2409.08806 [cs.LG]. URL: https://arxiv.org/abs/2409.08806.

[19] N. Gehrels. "The Swift Gamma-Ray Burst Mission". In: *AIP Conference Proceedings*. Vol. 727. AIP, 2004, pp. 637–641. DOI: 10.1063/1.1810924. URL: http://dx.doi.org/10.1063/1.1810924.

[20] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

[21] Md Zahidul Hasan, A. Ben Hamza, and Nizar Bouguila. *Time series forecasting with Hahn Kolmogorov-Arnold networks*. 2026. arXiv: 2601.18837 [cs.LG]. URL: https://arxiv.org/abs/2601.18837.

[22] Simon Haykin. "Neural Networks: A Comprehensive Foundation". In: 1994. URL: https://api.semanticscholar.org/CorpusID:60577818.

[23] James R. Heirtzler. "The future of the South Atlantic Anomaly and implications for radiation damage in space". In: *Journal of Atmospheric and Solar-Terrestrial Physics* 64 (2002), pp. 1701–1708.

[24] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. "Multilayer feedforward networks are universal approximators". In: *Neural Networks* 2.5 (1989), pp. 359–366. ISSN: 0893-6080. DOI: https://doi.org/10.1016/0893-6080(89)90020-8. URL: https://www.sciencedirect.com/science/article/pii/0893608089900208.

[25] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *International Conference on Learning Representations* (2015).

[26] Ray W. Klebesadel, Ian B. Strong, and Roy A. Olson. “Observations of Gamma-Ray Bursts of Cosmic Origin”. In: 182 (June 1973), p. L85. DOI: `10.1086/181225`.

[27] Andrey N. Kolmogorov. “On the representation of continuous functions of several variables by superpositions of continuous functions of one variable and addition”. In: *Doklady Akademii Nauk SSSR* 114 (1957), pp. 953–956.

[28] Chryssa Kouveliotou et al. “Identification of Two Classes of Gamma-Ray Bursts”. In: *The Astrophysical Journal Letters* 413 (Aug. 1993), p. L101. DOI: `10.1086/186969`.

[29] T.-P. Li and Y.-Q. Ma. “Analysis methods for results in gamma-ray astronomy”. In: *ApJ* 272 (1983), pp. 317–324.

[30] Ziming Liu et al. *KAN 2.0: Kolmogorov-Arnold Networks Meet Science*. 2024. arXiv: `2408.10205 [cs.LG]`. URL: `https://arxiv.org/abs/2408.10205`.

[31] Ziming Liu et al. *KAN: Kolmogorov-Arnold Networks*. 2025. arXiv: `2404.19756 [cs.LG]`. URL: `https://arxiv.org/abs/2404.19756`.

[32] Scott Lundberg and Su-In Lee. *A Unified Approach to Interpreting Model Predictions*. 2017. arXiv: `1705.07874 [cs.AI]`. URL: `https://arxiv.org/abs/1705.07874`.

[33] T. Mizuno, T. Kamae, and G. Godfrey. “Cosmic-ray interactions with the Earth’s atmosphere: gamma-ray production”. In: *ApJ* 614 (2004), p. 1113.

[34] Vinod Nair and Geoffrey Hinton. “Rectified Linear Units Improve Restricted Boltzmann Machines”. In: *ICML*. 2010.

[35] N. Parmiggiani et al. “A Deep Learning Method for AGILE-GRID Gamma-Ray Burst Detection”. In: *The Astrophysical Journal* 914.1 (June 2021), p. 67. DOI: `10.3847/1538-4357/abfa15`. URL: `https://doi.org/10.3847/1538-4357/abfa15`.

[36] N. Parmiggiani et al. “A New Deep Learning Model to Detect Gamma-Ray Bursts in the AGILE Anticoincidence System”. In: *The Astrophysical Journal* 973.1 (Sept. 2024), p. 63. DOI: `10.3847/1538-4357/ad64cd`. URL: `https://doi.org/10.3847/1538-4357/ad64cd`.

[37] N. Parmiggiani et al. “COSI Short Gamma-Ray Burst Localization Using BGO Shield Data”. In: *The Astrophysical Journal* 997.2 (Jan. 2026), p. 135. DOI: `10.3847/1538-4357/ae25fc`. URL: `https://doi.org/10.3847/1538-4357/ae25fc`.

[38] Karl Pearson. “Contributions to the Mathematical Theory of Evolution”. In: *Philosophical Transactions of the Royal Society of London. Series A* 185 (1894), pp. 71–110. DOI: `10.1098/rsta.1894.0003`.

[39] Fabian Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* (2011).

[40] S. Poolakkil et al. “The Fermi-GBM Gamma-Ray Burst Spectral Catalog: 10 yr of Data”. In: *The Astrophysical Journal* 913.1 (May 2021), p. 60. ISSN: 1538-4357. DOI: `10.3847/1538-4357/abf24d`. URL: `http://dx.doi.org/10.3847/1538-4357/abf24d`.

[41] R. D. Preece et al. “The BATSE Gamma-Ray Burst Spectral Catalog. I. High Time Resolution Spectroscopy of Bright Bursts Using High Energy Resolution Data”. In: *The Astrophysical Journal Supplement Series* 126.1 (Jan. 2000), p. 19. DOI: `10.1086/313289`. URL: `https://doi.org/10.1086/313289`.

[42] J. D. Scargle et al. “Studies in Astronomical Time Series Analysis. VI. Bayesian Block Representations”. In: *ApJ* 764 (2013), p. 167.

[43] Gideon Schwarz. “Estimating the Dimension of a Model”. In: *The Annals of Statistics* 6.2 (1978), pp. 461–464. DOI: `10.1214/aos/1176344136`. URL: `https://doi.org/10.1214/aos/1176344136`.

[44] Lloyd S. Shapley. “A Value for n-Person Games”. In: *Contributions to the Theory of Games II*. Ed. by H. W. Kuhn and A. W. Tucker. Vol. 28. Annals of Mathematics Studies. Princeton University Press, 1953, pp. 307–317.

[45] M. Tavani et al. “The AGILE Mission”. In: *Astronomy & Astrophysics* 502.3 (Jan. 2009), pp. 995–1013. ISSN: 1432-0746. DOI: `10.1051/0004-6361/200810527`. URL: `http://dx.doi.org/10.1051/0004-6361/200810527`.

[46] John A. Tomsick et al. *The Compton Spectrometer and Imager*. 2019. arXiv: `1908.04334 [astro-ph.IM]`. URL: `https://arxiv.org/abs/1908.04334`.

[47] John A. Tomsick et al. *The Compton Spectrometer and Imager*. 2023. arXiv: `2308.12362 [astro-ph.HE]`. URL: `https://arxiv.org/abs/2308.12362`.

[48] J. A. Van Allen and L. A. Frank. “Radiation around the Earth to a radial distance of 107,400 km”. In: *Nature* 183 (1958), pp. 430–434.

[49] Quan Zhou et al. *KAN-AD: Time Series Anomaly Detection with Kolmogorov-Arnold Networks*. 2025. arXiv: `2411.00278 [cs.LG]`. URL: `https://arxiv.org/abs/2411.00278`.

[50] A. Zoglauer, R. Andritschke, and F. Schopper. “MEGAlib – The Medium Energy Gamma-ray Astronomy Library”. In: *New Astronomy Reviews* 50.7 (2006). Astronomy with Radioactivities. V, pp. 629–632. ISSN: 1387-6473. DOI: `https://doi.org/10.1016/j.newar.2006.06.049`. URL: `https://www.sciencedirect.com/science/article/pii/S1387647306000972`.