




Gamma Ray Bursts: Theory and Data Analysis

Krittika Summer Projects 6.0

26 May to 29 July, 2025

Aditya Bansal

Indian Institute of Technology, Roorkee



My GitHub Repository

Contents

1	Overview of Gamma Ray Bursts	3
1.1	Discovery and Early Observations	3
1.2	Theoretical Challenges and the Origin Debate	4
1.3	Afterglow and Classification	6
1.4	Supernova Connection and Energetics	7
1.5	Recent Developments: <i>Swift</i> , Short GRBs, and Multi-Messenger Astronomy	10
2	Prompt and Afterglow Emission	13
2.1	Prompt Emission	13
2.2	Afterglow Emission	16
3	Metropolis Algorithm and Markov Chain Monte Carlo	20
3.1	Monte Carlo simulations	20
3.2	Markov Chains	21
3.3	Markov Chain Monte Carlo	22
3.4	The Metropolis-Hastings algorithm	24
4	Tasks	27
4.1	Task 1: MCMC Fitting of the Smooth Broken Power-Law Model to Afterglow Data	27
4.2	Task 2: Spectral Analysis Using the Cutoff Power Law Model	31

1. Overview of Gamma Ray Bursts

1.1 Discovery and Early Observations

Gamma-ray bursts (GRBs) are among, if not the most energetic and enigmatic events in the universe. In fact, one of the hooks that gets a person's attention is that this is the most powerful class of explosions that occur in the universe after the Big Bang. Their discovery in 1967 by the American satellite *Vela* was a byproduct of Cold War-era satellite monitoring to enforce the partial nuclear test ban treaty of 1963. The telltale sign of a nuclear test was expected to be a strong and fast (millisecond) outburst of γ -rays (as γ -rays are typically released in nuclear transitions), originating from the Earth or a known celestial body. The first GRB was initially detected as brief flashes in detectors onboard the *Vela* satellite, and was first reported to the community in 1973, after thorough analysis. Some events were simultaneously seen on the *US OSO-7*, *IMP-6*, and the Soviet *Kosmos-461* spacecrafts. Using various tests, it was confirmed that the signal was from a place that was incredibly distant from our own planet, even beyond our solar system, and it was very unlikely to be the result of any nuclear test. Such flares were completely unexpected in the community. This discovery opened up a new field in high-energy astrophysics with ongoing research into their origins and mechanisms, some of which that will be explored in this report.

The mysterious nature of these bursts led to intense debates among astronomers. Most early theories regarding the origin of GRBs suggested that they originated within our own Milky Way galaxy, but later observations, especially those with the Compton Gamma Ray Observatory, showed that GRBs are distributed uniformly across the sky. This isotropy strongly supported the idea that GRBs have a cosmological origin, arising from cataclysmic events in distant galaxies. Read more about this in the next section.




Figure 1: An artist's impression of the satellite *Vela 5B* in orbit, courtesy of Los Alamos National Laboratory.

1.2 Theoretical Challenges and the Origin Debate

Through the late 1960's and early 1970's, GRBs were detected through scintillation detectors (a device that detects radiation by using the property of certain materials called scintillators, eg. NaI(Tl), to emit light when struck by radiation) installed on satellites. They were not designed for the study of GRBs, and combined with their limited sensitivity, perhaps only a 100 of such bursts were detected. Direct comparison was incredibly difficult as the light curves (graphs that show the brightness of an object over a period of time) were often constrained to different energy bands. However, even with such limited data, scientists were able to conclude that GRBs either originated from individual dense stars, such as neutron stars or black holes, or sub-regions of larger objects, such as the cores or coronae of massive stars. This can be concluded from a light travel time argument that places strict limits on the size of the region creating the GRB, since signals crossing the region cannot do so in less than the light travel time. For $\Delta T = 1ms$, $c\Delta T = 300km$, and so the inferred region is only a few hundred km wide, hence, leading to the above argument.

Most models that explained the origin of GRBs fall into three categories. They are:

- Those which were powered by accretion power, accreting mass onto a white dwarf, neutron star or black hole, with varying accretion materials.
- Those related to stellar activity, eg. directed stellar flares.
- Those due to catastrophic destruction of stellar sized objects.

Some theories also proposed GRBs being emitted by white holes or cosmic strings, though they did not receive much support due to lack of evidence. However, it is to be noted that many proposed mechanisms have been shown to occur in the Universe, yet the vast majority do not emit GRBs.

Since the error regions obtained by triangulation to locate the sources of GRBs using satellites in the Earth's orbit were very large, it was practically impossible to pinpoint the source. The error regions consisted of thousands of stars and galaxies, any of which could have sourced the GRB. Hence, it was difficult to distinguish between the various means of GRB production. Hence, during the late 1980s, the absence of precise GRB positions intensified the controversy regarding their origin.

In 1995, a great debate on the issue was held between two camps advocating for their theories behind the origin of GRBs. Bodhan Paczynski argued for a Cosmological origin (occurring in distant galaxies across the universe), while Don Lamb advocated for a Galactic origin (occurring within or near our galaxy, namely the Milky Way). The debate was skewed towards the former due to the recent announcements from the BATSE (Burst And Transient Source Experiment) mission about the isotropy (appearing uniformly across the sky with no preference for any specific direction) of GRBs on the sky. Such isotropy naturally favors Cosmological models. However, Galactic models too remained possible due to the fact that neutron stars received natal kicks (high-speed recoil velocity given to compact stellar remnants such as neutron stars at birth due to supernova explosions) on their formation and can launch neutron stars into the galactic corona which is a sparse, extended region surrounding the Milky Way. At the end, both parties agreed that further observations were necessary.




Figure 2: Isotropic distribution of GRBs on the sky, color coded by the measured fluence (It is the time integral of the flux over the duration of the burst, acting as a measure of the burst's energy) of the burst. The isotropy is independent of the fluence. Courtesy of NASA.

1.3 Afterglow and Classification

GRB prompt emission, i.e. the intense initial burst of γ -rays and high-energy radiation that is emitted during the GRB, lasts for a very short amount of time, at most about a few minutes. This made it difficult to understand their nature. It is nearly impossible to obtain accurate source positions from the γ -rays alone, in fact, even if known in advance, it would be nearly impossible to observe the locations with enough precision during the prompt emission. Even today, scientists have only achieved this for very bright and/or long lived emissions. Hence, the idea of using GRB afterglows became popular. *BeppoSAX*, a joint Italian and Dutch satellite, offered a wide field X-ray monitor that could provide positions to as precise as arc-minutes, as well as the ability to re-point and observe the source position using an on-board X-ray telescope within a few hours. Both features were crucial to be able to observe afterglow emissions, as much precise GRB positions could now be obtained. Afterglow analysis could efficiently be done by covering the entire error region in one exposure, and limiting the candidate sources to at most a few hundred. This was first implemented on the 28th of February,

1997, for *GRB 970228*. In 1997, the process took several days. Now, with observations with *Swift*, these observations take seconds to minutes at most.

GRBs are generally classified into two classes on the basis of their duration/ T_{90} . T_{90} is the duration over which 90% of the GRB has been observed. Or more specifically, the duration between which the burst has released between 5% and 95% of its total fluence. Sometimes, tighter restrictions like T_{50} are used too.

- **Short-Hard bursts:** These are bursts that have a T_{90} of under 2s. They are called hard bursts because the peak radiation they emit lies in the high energy regions ($> 300keV$).
- **Long-Soft bursts:** These are bursts that have a T_{90} of over 2s. They are called soft bursts because most of its spectrum lies in the lower energy regions.

Though, evidence for these populations did not become widely acknowledged until the results from *BATSE*. Based on evidence from *BATSE* and *Swift*, populations of intermediate duration bursts ($2s < T_{90} < 10s$), as well as ultra-long GRBs, which last for hours. However, their relation to either the long or short GRBs remains a matter of discussion. Do keep in mind though that this classification is on the basis of the prompt emission of the GRBs.

1.4 Supernova Connection and Energetics

The connection between long GRBs and supernovae became clear with observations of *GRB 980425* in May 1998, which was accompanied by a bright supernova, *SN 1998bw* instead of a fading afterglow, in a nearby galaxy *ESO 184-G082* with a redshift (a measure of how much the wavelength of light from a distant object, such as a galaxy has been stretched as the universe expands) of $z = 0.0085$ (). This supernova was exceptionally luminous ($M_B \approx -19.3$) and showed unusual spectral features, showing no signs of either hydrogen or helium emission lines, but indicated high expansion velocities of $20000 - 30000km/s$. This provided strong evidence that long GRBs are associated with the deaths of massive stars through core-collapse events. Subsequent




Figure 3: T_{90} distribution of the first *BATSE* catalog, consisting of 222 GRBs of various durations. Courtesy of Chryssa Kouveliotou et. al..

observations of other long GRBs, such as *GRB 030329* in 2003 with $z = 0.17$, showed nearly identical supernova signatures to that of *GRB 980425*, though the former was over a 1000 times more luminous. This confirmed that long GRBs and energetic, broad-lined Type *Ic* supernovae are often connected, regardless of the total GRB energy output. Do note though that such GRBs that enable supernova analysis due to their close proximity ($z < 0.3$) are incredibly rare.

The energy released during these events is extraordinary. The parameter used for determining these energies is the isotropic equivalent energy release (E_{iso}). This is defined as the energy that the burst would have if it emitted the energy observed by us to observers in all directions. Its formula is:

$$E_{iso} = S_{\gamma} \frac{4\pi d_L^2}{1+z}$$

where S_{γ} is the measured fluence, z is the measured redshift, and d_L is the luminosity distance to the GRB. The factor of $1+z$ appears because the fluence is integrated in the observer frame, and dividing it by this extra factor brings it back to the emitter frame. For cosmological GRBs, E_{iso} can reach up to 10^{54} erg, comparable to a significant fraction of M_{\odot} emitted in only a few seconds. However, this is too large of an




Figure 4: The spectral evolution of *GRB 030329*, associated with *SN 2003dh*. The solid black line shows *SN 2003dh*, and the dashed red line shows *SN 1998bw*, associated with *GRB 980425*. The similarity in the two shows that highly energetic GRBs are associated with supernovae. Courtesy of Jens Hjorth et. al.

amount of the energy that the accretion on the black hole or neutron star might have to release the GRBs. The true energy comes out to be much lower due to relativistic beaming (the strong concentration of radiation from a moving source along the direction of its motion when it travels at speeds approaching the speed of light), which confines the emission into narrow jets. This beaming is inferred from the presence of jet breaks in afterglow light curves, and it reduces the total energy output by a factor of over 100. After correcting for beaming, the energy of long GRBs is typically around 10^{51} erg, which is still enormous but more consistent with the energy available from the collapse of a massive star. An alternative method of analysing GRB energetics arises from very late time observations in the radio band using calorimetry, where radio observations track synchrotron emission from electrons accelerated in the shock. This helps us to calculate the energy of the bursts without taking beaming into account, giving us much more accurate results directly.

1.5 Recent Developments: *Swift*, Short GRBs, and Multi-Messenger Astronomy

The launch of the *Swift* satellite on November 20, 2004 revolutionized the study of GRBs. *Swift* is equipped with the following components:

- **Burst Alert Telescope (BAT):** Responsible for γ -ray detection. Has a field of view of 1.4 sr and an energy range of $15 - 150\text{ keV}$. Has an error region of $2 - 4$ arc-minutes.
- **X-Ray Telescope (XRT):** Responsible for X-ray detection. Has a much narrower field of view, with an energy range of $0.2 - 10\text{ keV}$. Source positions are accurate upto $1 - 2$ arc-seconds.
- **Ultra-Violet and Optical Telescope (UVOT):** Responsible for UV and optical ray detection. Has a 30 cm telescope working in the range $1700 - 6000\text{ \AA}$.

These enable rapid localization and multi-wavelength follow-up of GRBs within minutes of detection. This has dramatically increased the number of well-observed GRBs (from

1 per month to 2-3 in a week), with most bursts now having X-ray, UV and/or visible afterglows. The ability to obtain accurate positions and multi-wavelength data within 100 *sec*, as compared to 10 *hrs* early on, after the burst has transformed our understanding of GRBs.

Prior to *Swift*, afterglows for short GRBs were extremely difficult to detect, due to large uncertainties in position and also due to them being dimmer than their long counterparts. But *Swift*'s capabilities have now enabled the identification of X-ray afterglows for about 70% and optical afterglows for about 25% of over 100 short GRBs detected. The first of these was *GRB 050509B* detected on the 9th of May, 2005 with a redshift of $z = 0.225$ and an X-ray afterglow. Observations have shown that short GRBs are often associated with a variety of host galaxies, including elliptical galaxies, where no long GRB had been found till then, and can be found at large offsets from their cores. This is consistent with models in which short GRBs arise from a gravitational wave driven merger of either neutron star-neutron star or neutron star-black hole binaries. The detection of a kilonova (a short-lived, optical/infrared transient powered by the radioactive decay of heavy elements like Th and U, analogous to a supernova in long GRBs) emission in a short GRB named *GRB 130603B* provides strong evidence for this merger model. Kilonovae can only be observed in relatively nearby events due to their low intensity, but their discovery has been crucial in discovering the origin of the short GRBs.

The extreme environments of GRBs make them promising sources for multi-messenger astronomy, i.e. using multiple astronomical signals to explain cosmic phenomena. We combine γ , X, and UV rays, which are electromagnetic with signals like gravitational waves and neutrinos, obtained from interactions between sub-atomic particles. The breakthrough for multi-messenger astronomy came in August 2017 with the joint detection of gravitational waves and a short GRB (*GW 170817/GRB 170817A*), observed by *LIGO*, *Virgo*, *Fermi*, and *INTEGRAL*, which was very unlikely to just be a coincidence. This event was accompanied by a kilonova in galaxy *NGC 4993*, which confirmed that neutron star mergers can produce both gravitational waves and short GRBs, and demonstrating the synthesis of heavy r-process elements, as mentioned in the above paragraph. While high-energy neutrinos have not yet been detected from

GRBs, ongoing searches by observatories like *IceCube* aim to identify such signals, which would provide further understanding of particle acceleration in GRB jets. Hence, the development of multi-messenger astronomy has enabled the study of such extreme environments that can help us understand the origin of one the most energetic phenomena discovered.




Figure 5: GRB intensity of *GRB 170817A* as measured by ESA's INTEGRAL satellite on the 17th of August 2017. Courtesy of ESA.

2. Prompt and Afterglow Emission

2.1 Prompt Emission

GRBs were first detected and named for their prompt, intense flashes of γ -ray emission, which outshine all other sources in the γ -ray sky at their peak. The prompt emission is highly variable from burst to burst, with durations ranging from a fraction of a second to several hours, and light curves that may be smooth (showing a single gradual pulse without sharp fluctuations or multiple peaks), erratic (highly irregular, with rapid, unpredictable changes), or exhibit a FRED (fast rise and exponential decay) profile. We have already seen how GRBs are classified on the basis of their t_{90} , so let us dive into further details. We have also studied that short bursts emit more high energy radiation, while long bursts somewhat less. This can be encapsulated by the hardness ratio, which denotes either the count rate, or energy contained within one band, compared to another, using their ratio. Eg. consider observing GRBs in two bands, $25 - 50 \text{ keV}$ and $50 - 100 \text{ keV}$ ranges. We can write the hardness ratio as:

$$HR = \frac{S_2}{S_1}$$

, where S_1 and S_2 are the fluences of the two bands. Typically, short GRBs have a greater HR than long GRBs. Observational constraints, such as overlap in duration distributions between long and short GRBs near the 2-sec boundary, or the dependence of measured duration on detector energy range and sensitivity complicate attempts to place GRBs with durations of $1 - 5 \text{ s}$ firmly within one class or another, and lead to the requirement of using additional diagnostics.

The spectral structure of prompt emission is typically non-thermal (or at least has a non-thermal component), often described by the Band function, named after the person who discovered it—a smoothly joined pair of power laws with an exponential cutoff.

The function can be described as:

$$F_E = \begin{cases} A \left(\frac{E}{100 \text{ keV}} \right)^\alpha e^{\frac{-E(2+\alpha)}{E_{peak}}}, & \text{if } E < \frac{(\alpha-\beta)E_{peak}}{2+\alpha} \equiv E_{break} \\ A \left[\frac{(\alpha-\beta)E_{peak}}{100 \text{ keV}(2+\alpha)} \right] e^{(\beta-\alpha)(\frac{E}{100 \text{ keV}})^\beta}, & \text{if } E \geq \frac{(\alpha-\beta)E_{peak}}{2+\alpha} \end{cases}$$

, where E_{peak} is the peak of $\nu F_\nu = EF_E = E^2 N_E$ spectrum. Simple power-laws often provide very good representations of the observations from the *Swift-BAT*, since the 15 – 350 keV range often does not contain the peak energy. In recent years, it has also been discovered that a significant fraction of the GRBs also include a thermal component, which has a temperature in the tens to hundreds of keV range. This component dominates in a few bursts, while in others it only contributes a little.

Additionally, evidence of high polarization (around $80 \pm 20\%$ in the case of *GRB 021206*) in the prompt emission, which depends on the Compton scattering cross section in the detector on the polarization angle, suggests a highly ordered magnetic field or emission geometry. In particular, the measured angular distribution of counts depends on $1 - Q\Pi\cos 2(\eta - \phi)$ where Q is the polarimetric modulation factor (a parameter that quantifies how strongly the measured angular distribution of γ -ray counts in a detector depends on the polarization angle and degree of the incident γ -rays), Π is the degree of polarization, ϕ the scattering angle and η the degree of polarization. The origin of GRB prompt emission remains a central question in




Figure 6: A modified cartoon depiction showing schematics of the fireball model and the basic components of the internal and external shocks. Courtesy of Godson Fortune Abbey et. al..

astrophysics, despite over 50 years of observation. The most widely accepted framework

is the fireball shock model (developed by Martin Rees and Peter Mészáros in the 1990s), which suggests that a catastrophic event deposits a large amount of energy (in the form of heat) in a small volume. This creates a relativistic jet with the bulk Lorentz factor Γ (a measure of the relativistic speed of the entire outflow or jet as a whole, rather than the speed of individual particles within) exceeding 100 once the thermal pressure is released. The difference in velocities of shells of material ejected at different times (due to the initially released material encountering resistance due to having to escape through heavily baryon loaded material) creates shocks. The distance of these shocks from the source can be estimated by the formula

$$R_{prompt} = c \delta t \Gamma_1 \Gamma_2$$

, where δt is the time interval between the ejection of the two shells, and Γ_1 and Γ_2 are the bulk Lorentz factors of the two ejections. And the fraction of kinetic energy dissipated in the shock can be given by

$$\varepsilon = \frac{\Gamma_1 + \Gamma_2 - 2\sqrt{\Gamma_1 \Gamma_2}}{\Gamma_1 + \Gamma_2}$$

As this outflow expands and interacts with itself or the surrounding medium, internal shocks between shells of material with different velocities dissipate kinetic energy, accelerating particles and producing the observed γ -rays via non-thermal processes such as inverse Compton radiation. Also, the acceleration of electrons due to this phenomenon, albeit due to external shocks, is the cause of the afterglow emission as per the fireball model. However, the presence of thermal components and very high-energy emission in some bursts indicates that not all prompt emission is shock-powered; some may originate from the photosphere of the fireball. The photospheric radius can be given as

$$R_{ph} = \frac{L \sigma_T}{4\pi m_p c^3 \Gamma^3} \approx 3.7 \times 10^9 \text{ m}$$

. While the fireball model provides a robust framework, many details, such as the relative roles of thermal and non-thermal emission, the origin of extremely high-energy photons, and the reasons for variability, remain unresolved. Thus highlighting the complexity of GRB prompt emission.

2.2 Afterglow Emission

The initial searches for GRB afterglows were motivated by the need to precisely locate these enigmatic events, which otherwise had sky positions with large error circles of several degrees in radius. This was because of the fact that γ -rays (obtained from the prompt emission) can't be focused, making it difficult to accurately localize them. Before the advent of modern satellites like *Swift*, accurate localization relied on triangulation using the *InterPlanetary Network*, which could sometimes narrow positions to strips of sky, but often with delays of days, which were also difficult to observe. Technological limitations—such as less sensitive X-ray satellites and small-field ground-based optical detectors—made searches challenging. The breakthrough came with *BeppoSAX* which launched in April 1996, which used coded mask cameras to localize GRBs to within a few arc-minutes, enabling the first detections of X-ray and optical afterglows and revolutionizing the field. It had a range of 40×40 square degrees and operated in the energy range of $2 - 30$ keV. The first GRB to be analysed using this telescope was *GRB 970228*. Throughout the late 1990's and early 2000's afterglows were discovered at a rate of approximately one per month, from a range of sources including *BeppoSAX*, *HETE-2*, the *IPN*, and the *Rossi X-ray Timing Explorer (RXTE)*. After the *Swift* era began, GRBs were discovered once every few days.

X-ray afterglows are very common in GRBs, especially with rapid-response satellites like *Swift*. These afterglows display a complex, multi-phase light curve:

- **Prompt Emission:** Sometimes, X-rays are detected even while the γ -ray emission is ongoing. In this case, these X-ray observations can be considered the soft end of the prompt emission. These are both temporally and spectrally variable, and can be extremely bright.
- **1st Rapid decay:** At the end of the prompt emission, visible light appears to undergo a rapid decay of t^{-5} or steeper. This is because photons generated at an angle from our direct line of sight take longer to reach the observer. This is called high latitude emission.
- **Plateau:** After the decay, X-rays appear to plateau, i.e. appear as flat or with a

slow decay of $t^{-0.5}$ which might relate to ongoing central engine activity.

- **Flares:** During the plateau phase, bursts can exhibit small or large X-ray flares. In extreme cases, these flares may have energy comparable to the GRBs themselves. **Intermediate decay:** Following the plateau, the X-ray afterglow will typically begin to fall as t^{-1} . This was the most commonly observed phase prior to the launch of *Swift*.
- **2nd Rapid decay:** After the previous phase, burst lightcurves undergo a temporal break (called the jet break) in which the decay steepens. Beyond this point, the flux decays as t^{-2} .

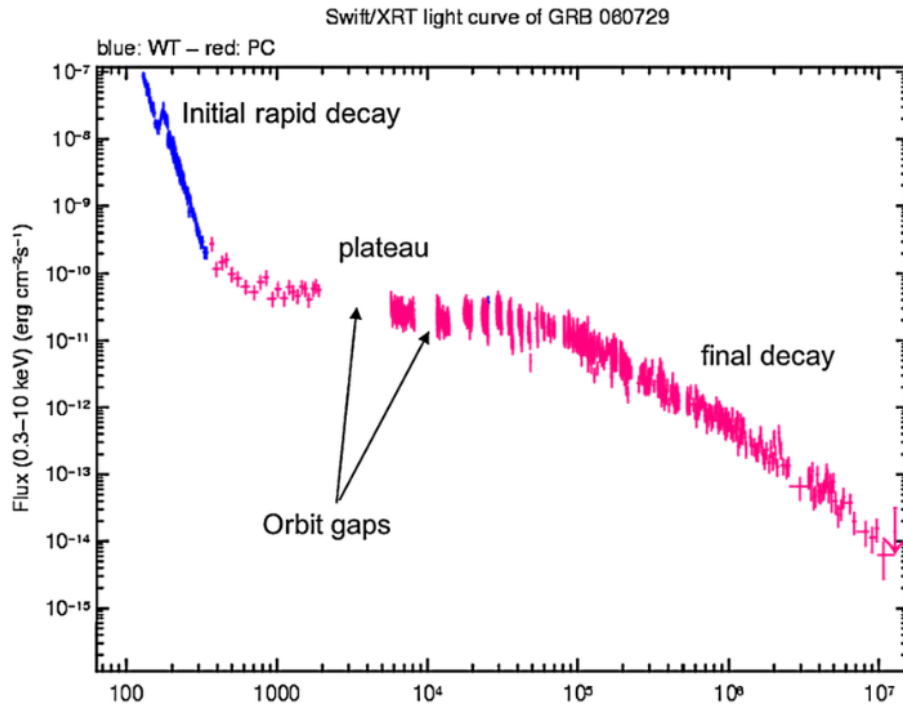


Figure 7: XRT X-ray light curve of GRB 080729, showing initial rapid decay, plateau phase with orbit gaps, and final decay. Courtesy of R. Willingale et. al..

Optical afterglows play a crucial role in GRB astronomy by providing the most accurate localizations (since the spatial resolution of optical imaging is typically far higher than for X-ray imaging), which enable the identification of host galaxies and the direct measurement of redshifts through absorption features imprinted on the afterglow spectrum. These afterglows are typically well-described by power-law behavior in both frequency and time, with flux following $F_\nu(\nu, t) = t^\alpha \nu^\beta$ reflect the underlying physics of the blast wave and its interaction with the surrounding medium. Often, $\beta = -1$ and

$\alpha = -1$ initially with eventually steepening to $\alpha = -2$. However, not all GRBs exhibit detectable optical counterparts. Such "dark GRBs" are defined by the absence of optical emission despite prompt and deep searches (we use the criterion of $\beta_{OX} < 0.5$, where β_{OX} is the optical to X-ray spectral index), often attributed to factors like dust extinction in the host galaxy, high redshift causing Lyman-alpha absorption, low circumburst density, or simply observational limitations. The majority of dark bursts are now understood to be due to dust absorption along the line of sight, as supported by high hydrogen column densities and the properties of their host galaxies. Additionally, optical afterglows provide unique insights into the structure and geometry of GRB jets through polarimetric measurements. The detection of linear polarization—sometimes highly variable in degree and direction—indicates asymmetry in the outflow and allows direct probing of jet structure, while rare observations of circular polarization present new challenges for theoretical models. Thus, optical afterglows are not only key to localizing and characterizing GRBs but also serve as powerful probes of their physical environments and jet dynamics.

Radio and sub-mm afterglows are rarer and more challenging to detect than their X-ray or optical counterparts due to the high expense and difficulty of deep radio observations. Unlike other afterglows, radio emission often peaks days to weeks after the burst and can persist for months or years. These observations are crucial for probing the total energy and geometry of GRBs, as the radio emission becomes less beamed over time and provides a more direct measure of the energy budget. The first radio afterglow detection in *GRB 970508* also yielded the first GRB redshift.

The dominant emission process in GRB afterglows is synchrotron radiation from electrons accelerated in the shock front between the relativistic outflow and the surrounding medium. The resulting spectrum is a series of connected power laws with breaks at characteristic frequencies, reflecting different physical processes such as self-absorption, peak emission, and cooling. The shape and evolution of the spectrum depend on both macroscopic (energy, geometry) and microscopic (electron acceleration, magnetic field) properties, making afterglows powerful diagnostics of GRB physics. In addition to the forward shock, a reverse shock can produce early optical and radio flares before fading away.

Strong evidence for relativistic beaming in GRBs comes from observations of jet breaks—achromatic steepenings in the afterglow light curves at X-ray, optical, and radio wavelengths. These breaks occur when the relativistic jet slows sufficiently for its edges to become visible, causing the afterglow to fade more rapidly. Jet breaks allow estimation of the jet’s opening angle and the true energy release, which is much lower than the isotropic equivalent energy. The existence of “orphan” afterglows—seen by observers outside the original jet cone—further supports the idea that GRB outflows are highly collimated and relativistic.

3. Metropolis Algorithm and Markov Chain Monte Carlo

3.1 Monte Carlo simulations

A Monte Carlo simulation is a mathematical technique that uses repeated random sampling to estimate the probability of different outcomes in a complex system or process having a significant uncertainty. In simpler terms, it is an experiment performed over and over again till we find what is the probability of a particular event. In most of the cases, it is performed on a computer due to time constraints in real life, but the principle remains the same. It works on the principle of the Law of Large Numbers, which states that if an experiment is repeated multiple times, the probability of the result will get closer and closer to the expected value. Eg. if you roll a die many times, the chance of rolling a particular number, say 2, will approach $1/6 = 0.16\dots$. Let us now test this by using a Python program.

```
1 from numpy import random
2
3 X = []
4 N = 100000
5 while len(X) < N:
6     x = random.randint(1, 7)
7     X.append(x)
8
9 n = 0
10 for i in X:
11     if i == 2:
12         n += 1
13
14 print(n/N)
```

Listing 1: Monte Carlo simulation of rolling a standard die.

As you can see after running the code, each time the result comes close to the true value, i.e. $1/6$. And also, note that if we increase the value of N , we increase our accuracy with each extra 0. Let us now discuss its pros and cons.

Advantages	Disadvantages
<ul style="list-style-type: none"> • Makes it possible to model systems without a known analytical solution. • Allows one to see how changes in multiple variables affect the probability of an outcome. • Can be used in multiple fields for a wide range of problems. 	<ul style="list-style-type: none"> • It is limited by computational power, as more accurate answers require a higher computational cost. • It is difficult to generalize the output to obtain a function for the same, as it predicts for particular values.

Table 1: Pros and Cons of the Technique

3.2 Markov Chains

A Markov chain is a stochastic process (a series of random events that happen over time) that models a sequence of events where the probability of transitioning to any future state T_{n+1} depends exclusively on the current state T_n , not the history of prior states ($T_1, T_2, T_3, \dots, T_{n-1}$)—a property known as memory-less-ness or the Markov property. Mathematically, this can be written as:

$$P(T_{n+1} \mid T_1, T_2, T_3, \dots, T_n) = P(T_{n+1} \mid T_n)$$

Consider this example: In a fictional town, the weather can be either sunny, or rainy. I.e. the weather can not be cloudy, humid, windy, etc. And the weather on the given day depends on the weather of the day before only. It is known that if a given day was

sunny, the next day would have a 50% chance of being sunny and a 50% chance of having a rainy weather. And on the contrary, if a given day was rainy, then the next day is certain to be sunny. We can clearly see that this scenario is a Markov process as it satisfies the Markov property mentioned above.

3.3 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) is a computational technique for sampling (selecting a representative subset of individuals from a large population to make inferences about the whole) from probability distributions when direct sampling is infeasible, particularly in high-dimensional spaces or complex models. It combines Markov chain properties with Monte Carlo simulation to generate representative samples from target distributions, enabling numerical integration and inference.

First, let us go over some fundamental definitions:

- **Prior ($P(\theta)$):** In Bayesian statistics, a prior pdf (Probability Density Function) represents your initial belief about a parameter before seeing new data, can be obtained from historical data or expert opinion. Eg. In a clinical study, it is estimated that a given drug works 20% of the time.
- **Likelihood ($P(D | \theta)$):** A likelihood pdf is a measure of how well the new data fits a given parameter value. For example, if 15 out of 20 patients recover with the help of the drug, the likelihood quantifies how probable this result is for different efficacy rates θ .
- **Posterior ($P(\theta | D)$):** A posterior pdf combines the prior and likelihood to yield an updated belief about the parameter after examining new data. Eg. After seeing the trial data (15/20 recoveries), the posterior updates the drug's efficacy probability. This is done with the help of Bayes' theorem. It goes:

$$P(\theta | D) = \frac{P(\theta)P(D | \theta)}{P(D)}$$

We can now begin to talk about the features of MCMC:

- MCMC is ideal for problems where you know the target pdf up to a constant factor, i.e. $P(\theta | D) \propto P(\theta)P(D | \theta)$. In other words, the ratio $P(\theta_1 | D) : P(\theta_2 | D)$ can be calculated for any two parameter states θ_1 and θ_2 .
- Note that MCMC is not recommended for optimization or exhaustive parameter-space searches. Its strength lies in sampling, not point estimation. Use an optimizer to find the optimum of the likelihood or posterior pdfs. And use a search algorithm, if you want to search all of parameter space.
- A sampling $\{\theta_k\}_{k=1}^K$ from $P(\theta)$ is a set of K draws so that the expectation values of functions, eg. $g(\theta)$ are approximated using sample averages, like

$$\mathbb{E}[g(\theta)] \approx \frac{1}{K} \sum_{k=1}^K g(\theta_k).$$

You can also choose to use $f(\theta) \propto P(\theta)$ to find the expectation values as follows:

$$\mathbb{E}[g(\theta)] \approx \frac{\sum_{k=1}^K g(\theta_k) f(\theta_k)}{\sum_{k=1}^K f(\theta_k)}$$

- The output of a Markov chain is the set of values $\theta_1, \theta_2, \theta_3, \dots, \theta_n$ with θ_{k+1} depending only on θ_k . The sequential samples are correlated, increasing the variance of estimators. The effective sample size $K_{eff} < K$ quantifies information loss. And the chains must traverse the parameter space repeatedly to ensure the sampling represents the target pdf.

Let us now talk about the advantages and disadvantages of this algorithm.

Advantages	Disadvantages
<ul style="list-style-type: none"> • It can sample from highly complex or high-dimensional probability distributions, even when direct sampling is impossible or infeasible. • It allows sampling from distributions that are known only up to a normalization constant Z, which is essential for Bayesian inference where the evidence is often unobtainable. 	<ul style="list-style-type: none"> • It can be computationally demanding, especially for large datasets or high-dimensional models. • Ensuring that the Markov chain has converged to the target distribution is non-trivial and requires careful monitoring and diagnostic checks.

Table 2: Pros and Cons of MCMC

3.4 The Metropolis-Hastings algorithm

The Metropolis-Hastings (M-H) algorithm is a foundational MCMC method for sampling from complex probability distributions. It constructs a Markov chain whose stationary distribution (essentially, a probability distribution that converges over time) converges to a target distribution $\pi(x)$. It is the simplest MCMC algorithm. It goes as follows. Here, s_1 is the k^{th} sample, which is used to find the $(k+1)^{th}$ sample. f is the probability function to be sampled. Keep in mind that we require only the ratio of f at different points to use this algorithm, and thus, we need not have a normalised probability function. And q is a proposal pdf function that can deliver samples, such that the algorithm can draw θ_{k+1} given θ_k .

Algorithm 1 Metropolis-Hastings Algorithm**Require:** Initial state x_0 , target density $f(x)$, proposal $q(x'|x)$, number of steps N

```

1:  $A \leftarrow [x_0]$ 
2: for  $n = 1$  to  $N$  do
3:   Propose  $x'$  from  $q(x'|A[-1])$ 
4:   Sample  $r \sim \text{Uniform}(0, 1)$ 
5:   if  $f(x')/f(A[-1]) \geq r$  then
6:      $A.append(x')$ 
7:   else
8:      $A.append(A[-1])$ 
9:   end if
10: end for
11: return  $A$ 

```

Here is the Python script for this.

```

1 def MH_MCMC(f, g, a, n = 100000):
2     A = [a]
3     while len(A) <= n:
4         p = A[-1]
5         q = g(p)
6         r = np.random.uniform(0, 1)
7         f_p = f(p)
8         f_q = f(q)
9         if f_p <= 0 or f_q <= 0 or np.isnan(f_p) or np.isnan(f_q):
10            :
11            A.append(p)
12            continue
13         if f_q/f_p >= r:
14             A.append(q)
15         else:
16             A.append(p)
17     return A

```

Listing 2: Metropolis-Hastings Algorithm

Let us go over some of its features.

- The algorithm generates a Markov chain by proposing new states based only on the current state (not the full history, similar to the Markov property), performing a biased random walk through the parameter space that spends time in regions proportional to the target density.
- At each step, a proposed sample is accepted with probability equal to the ratio of the target function at the proposed and current positions, ensuring detailed balance (see next point) and convergence to the desired stationary distribution.
- This algorithm uses the detailed balance condition. It ensures that the proposal distribution $q(x'|x)$ is symmetric. I.e. it is just as easy to move from x' to x as it is to move from x to x' . Mathematically, $q(x'|x) = q(x|x')$. If the proposal distribution follows this detailed balance condition, this algorithm is called the Metropolis algorithm, a special case of the more general Metropolis-Hastings algorithm, which need not use a symmetric distribution.
- The samples are generally auto-correlated, so effective sample size is less than the total number of steps, and initial samples (also termed “burn-in”) may need to be discarded to ensure stationarity.




Figure 8: Demonstration of the Metropolis algorithm used to sample from a normal distribution with mean 4 and standard deviation 2.

4. Tasks

4.1 Task 1: MCMC Fitting of the Smooth Broken Power-Law Model to Afterglow Data

The objective of this assignment is to model and fit the multi-wavelength afterglow light curve of *GRB170817* using the smooth broken power-law model. This is achieved by applying MCMC techniques to data from two observational frequencies, following the approach described in Makhatini et al. (2021). This analysis aims to estimate the parameters governing the afterglow evolution and to compare the results with published values, helping gain insights into the physical properties of the event.

The dataset analyzed in this report was sourced from this public Caltech repository. It contains multi-epoch, multi-frequency flux density measurements of the *GW170817* afterglow, spanning radio (eg. from *VLA*) to X-ray (eg. from *Chandra*) bands, along with associated uncertainties and instrument details. This comprehensive dataset, widely used in previous studies, helps with robust modeling of the afterglow evolution

and enables direct comparison with existing results in the literature.

The model used to sample the parameters for the given GRB is the Smooth Broken Power Law, which is one of the many functions used to model GRB prompt and afterglow emissions along with functions like the Band Function, Cutoff Power Law, etc. It is expressed mathematically as:

$$F(t, \nu) = \left(\frac{\nu}{3 \text{ GHz}} \right)^\beta F_p \times \left[\left(\frac{t}{t_p} \right)^{-s\alpha_1} + \left(\frac{t}{t_p} \right)^{-s\alpha_2} \right]^{-1/s}$$

where,

- ν is the observing frequency.
- s is the smoothness parameter, i.e. it determines how sharp or gradual the transition is at the break.
- F_p is the flux density at 3 GHz at the light-curve peak.
- t is the time passed post merger.
- t_p is the time where the sign of the slope of the power law changes.
- β is the spectral index, i.e. it characterizes how the observed flux changes with frequency.
- α_1 and α_2 are the power-law rise and decay slopes, respectively.

The analysis employed MCMC methods, particularly the Metropolis algorithm, to estimate the parameters of the smooth broken power-law model. It was chosen because it efficiently samples from complex, multi-dimensional posterior distributions, allowing robust parameter estimation and uncertainty quantification. Data pre-processing involved extracting relevant light curve measurements (at selected frequencies 4.5 GHz for *VLA* and 1.3 GHz for *MeerKAT*) from the multi-instrument dataset, ensuring all times were referenced to the merger event and uncertainties were appropriately included. The sampler was initialized with starting parameter values guided by literature and initial curve inspection; proposal step sizes were set to balance

exploration and acceptance rates. Each chain ran for 20,000 steps, with the first 20% of samples, or 4000 samples in all, discarded as burn-in to ensure convergence, and the final samples were used for statistical analysis of the posterior distributions.

Table 3: Parameter Estimates (Mean \pm Std)

Parameter	MeerKAT	VLA	Makhatini et al. (2021) Reference*
α_1 (rise slope)	1.156 ± 0.522	1.311 ± 0.526	~ 0.9
α_2 (decay slope)	0.553 ± 0.050	0.532 ± 0.031	~ 2.0
β (spectral index)	0.807 ± 0.096	0.609 ± 0.156	~ 0.6
s (smoothness)	14.594 ± 7.323	13.347 ± 8.748	~ 5.0
F_p (peak flux μJy)	69.696 ± 18.067	52.744 ± 5.560	20 ± 3
t_p (peak time days)	97.149 ± 49.702	205.653 ± 36.319	164 ± 12

*Values from Makhatini et al. (2021) are approximate literature references for comparison.




Figure 9: Trace plots for the MCMC sampling of model parameters in GW170817 afterglow fits: (a) MeerKAT and (b) VLA datasets. Each subplot shows the sampled values for a parameter across MCMC steps, illustrating the convergence and mixing quality achieved during parameter estimation for both instruments.

The estimated parameters from the MCMC fits show agreement with the literature for some parameters, such as the rise slope (α_1) and spectral index (β). However, notable

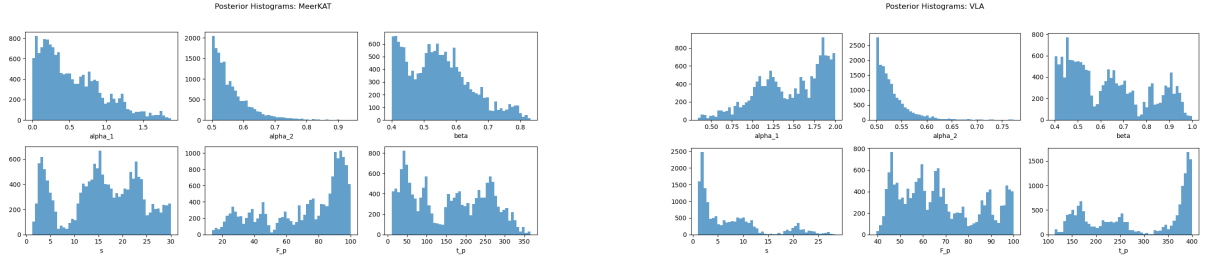


Figure 10: Posterior distributions of the model parameters for GW170817 afterglow fits: (a) MeerKAT and (b) VLA datasets. Each histogram displays the marginalized posterior for one parameter, illustrating the statistical uncertainties and differences between the two instruments' MCMC results.

discrepancies appear in the decay slope (α_2), smoothness parameter (s), and peak flux (F_p), which differ significantly from the values reported in Makhatini et al. (2021). For instance, the decay slope α_2 is found to be near 0.55 in our fits, substantially lower than the typical literature value of approximately 2.0, while the smoothness parameter s is higher, around 14–15 compared to the expected ~ 5 . The peak flux is also overestimated relative to reference values. These inconsistencies may result from issues such as choice of priors, data subsets selected for fitting, parameter space boundaries, or incomplete MCMC convergence. It is important to revisit the fitting procedure, including initial parameter guesses, chain diagnostics, and the treatment of data uncertainties, to address these differences and ensure robust, physically meaningful parameter estimates.

The MCMC fitting of the smooth broken power-law model to GW170817 afterglow data from MeerKAT and VLA recovered key parameters consistent with previous studies, confirming the model's ability to describe the afterglow's evolution. The agreement between datasets supports the reliability of the results and current physical interpretations. Future work should focus on multi-frequency fitting, more detailed jet models, and advanced Bayesian methods to better constrain parameters and address uncertainties.

4.2 Task 2: Spectral Analysis Using the Cutoff Power Law Model

Gamma-ray bursts (GRBs) are among the most luminous and energetic events observed in the universe. Traditionally, their prompt emission spectra are analyzed using the Band function, a widely adopted empirical model. In this report, we analyze the spectrum of GRB080916009 using the cutoff power law model as an alternative to the Band function. Our analysis is conducted with the 3ML framework, using the Fermi GBM data, and closely follows the methodology detailed in this link. The objective is to evaluate the effectiveness of the cutoff power law model and compare its fit quality to other functional forms, namely the smoothly-broken power law.

In this analysis of GRB080916009, the prompt emission spectrum is modeled using a cutoff power law function as an alternative to the commonly used Band function. Data from Fermi GBM detectors (NaI n3, n4, and BGO b0) are used, with source and background intervals dynamically selected by parsing the Fermi GBM catalog, offering greater flexibility compared to the hard-coded detector and interval definitions in the reference 3ML tutorial notebook. Time-Tagged Event (TTE) and CSPEC files are loaded alongside their detector response matrices, and background spectra for each detector are independently modeled using auto-determined polynomial fits. These are subtracted to yield background-corrected spectra for source intervals.

Spectral fitting is performed within the 3ML framework using likelihood-based methods and a cutoff power law model. Bayesian inference via MCMC sampling (using the “emcee” sampler) is employed to estimate parameter posteriors. The analysis documents non-critical issues such as non-OGIP compliant DRM files and duplicate TTE time tags to maintain transparency. Compared to the 3ML GRB080916C tutorial, which uses the Band function and manual selections, this approach offers dynamic configuration and model flexibility. Both methods utilize statistical tools like AIC, BIC, and DIC for model evaluation and selection.

The cutoff power law spectrum is defined as:

$$N(E) = K \left(\frac{E}{E_0} \right)^{\text{index}} \exp \left(-\frac{E}{E_c} \right)$$

where:

- $N(E)$ is the photon flux per unit energy at energy E .
- K is the **normalization constant** ($1/(\text{keV} \cdot \text{s} \cdot \text{cm}^2)$) — it sets the overall amplitude of the spectrum.
- index is the **photon index** (dimensionless) — it controls the power-law slope at low energies.
- E_c is the **cutoff energy** (keV) — the energy scale at which the exponential decay sets in.
- E_0 is a fixed reference energy, often taken as 100 keV.

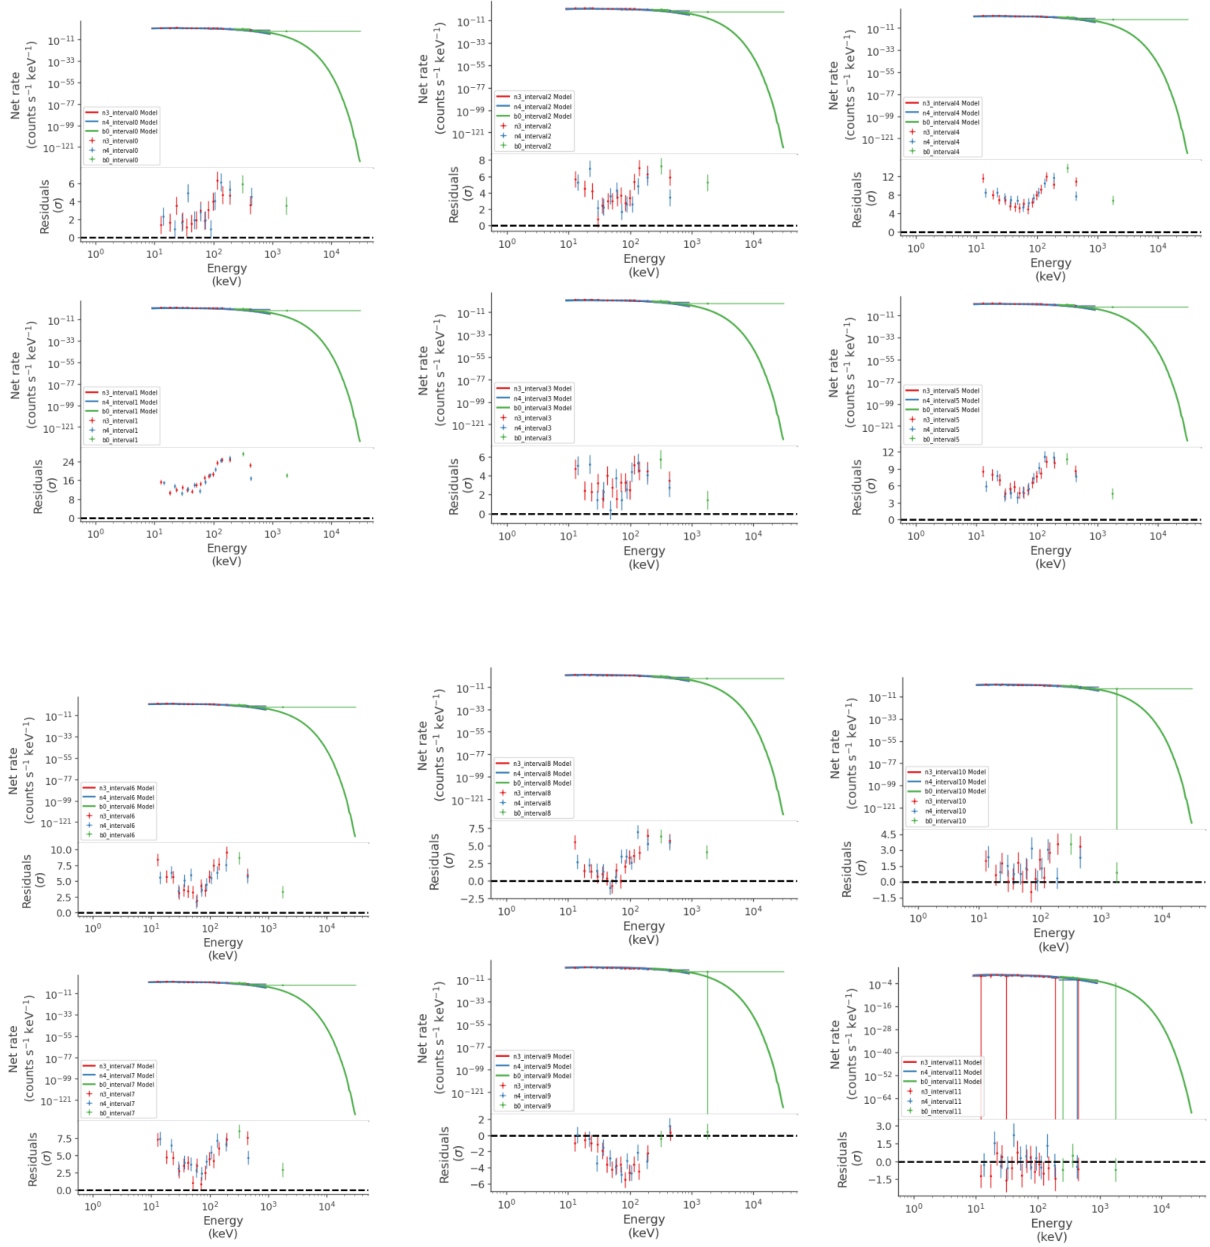


Figure 11: Time-resolved spectral fits and residuals for GRB080916009 across sequential intervals, using the cutoff power law model and Fermi GBM detectors n3, n4, and b0.




Figure 12: Spectral energy distributions for all time intervals of GRB080916009, fit with the cutoff power law model in physical flux units, showing time evolution of the prompt emission.

The cutoff power law model provides a superior fit to GRB080916009's spectra compared to the smoothly broken power law and Band function, as indicated by its significantly lower AIC, BIC, and DIC values; it achieves this with simpler parameterization, and residuals show no major systematic deviations, confirming its statistical and practical adequacy for this burst.

The cutoff power law provides a strong, statistically favored fit for GRB080916009, with significantly lower AIC, BIC, and DIC values than the Band and smoothly broken power law models. Residuals show no major systematics, confirming its suitability. While some minor data and software warnings were noted, they did not affect core results. Overall, the cutoff power law is an efficient and justified alternative for modeling this burst's prompt emission.

The analysis finds that the cutoff power law model delivers an excellent and statistically robust fit to the prompt emission spectra of GRB080916009, outperforming both the Band function and the smoothly broken power law according to information criteria (AIC, BIC, DIC). The model shows no major systematic deviations in fit residuals and achieves a close match to the observed data across all intervals.

These results demonstrate that, for GRB080916009, the cutoff power law provides both an efficient and interpretable description of the burst spectrum. Its relative simplicity, combined with strong statistical support, indicates that it is a suitable and insightful choice for modeling GRBs of this spectral shape, offering reliable parameter estimates and supporting clear physical interpretation without unnecessary model complexity.

References

- [1] David W. Hogg and Daniel Foreman-Mackey. Data analysis recipes: Using markov chain monte carlo. *The Astrophysical Journal Supplement Series*, 236:11, May 2018.
- [2] Andrew J. Levan. Gamma-ray bursts, 2018.
- [3] S. Makhathini, K. P. Mooley, M. Brightman, K. Hotokezaka, A. J. Nayana, H. T. Intema, D. Dobie, E. Lenc, D. A. Perley, C. Fremling, J. Moldòn, D. Lazzati, D. L. Kaplan, A. Balasubramanian, I. S. Brown, D. Carbone, P. Chandra, A. Corsi, F. Camilo, A. Deller, D. A. Frail, T. Murphy, E. J. Murphy, E. Nakar, O. Smirnov, R. J. Beswick, R. Fender, G. Hallinan, I. Heywood, M. Kasliwal, B. Lee, W. Lu, J. Rana, S. Perkins, S. V. White, G. I. G. Józsa, B. Hugo, and P. Kamphuis. The panchromatic afterglow of gw170817: The full uniform data set, modeling, comparison with previous results, and implications. *The Astrophysical Journal*, 922:154, 12 2021.