Krittika Summer Project 2025

Gamma Ray Bursts: Theory and Data Analysis

Pratham Srivastava

Mentors: Mehul Goyal and Yashowardhan Rai

Department of Electrical Engineering
IIT Bombay

June - July 2025

Contents

1 Introduction

Gamma-ray bursts (GRBs) are extraordinary astrophysical events characterized by brief and intense flashes of gamma-ray radiation, representing some of the most energetic phenomena known in the universe. These extreme electromagnetic emissions are second to the Big Bang as the most energetic and luminous phenomenon ever known.

Figure 1: An illustration of a gamma-ray burst. Gamma ray bursts happen when neutron stars collide or when giant stars explode into black holes, releasing jets of super-energetic photons that look like a narrow beam of a flashlight.

GRBs are now a major astrophysical tool to probe the final stages of stellar evolution and the creation of supernovae, to examine the creation and propagation of relativistic jets, and their impact on the Universe at large, and to use as lighthouses to identify galaxies across cosmic time – including some in the earliest epochs of the Universe.

Their enigmatic nature has intrigued the scientific community for decades, prompting extensive research to decipher their origins and the underlying physical processes involved. In the near half century since the announcement of their discovery, GRBs have moved from a niche question to the astrophysical mainstream and are one of very few classes of source to be identified from the very high-energy gamma-ray to the low frequency radio regimes.

Part I: Theoretical Analysis of Gamma Ray Bursts

---

2 Origins and Overview: Understanding GRBs and Their Fundamentals

2.1 The "Dark" Era (1973(67) – 1991)

The serendipitous discovery of GRBs traces back to the late 1960s when the US military Vela satellites were launched, each equipped with rudimentary γ-ray detectors, in order

to monitor Soviet compliance with the Partial Test Ban Treaty of 1963, which prohibited all test detonations of nuclear weapons, except for those conducted underground. In particular, the telltale sign of a nuclear test was expected to be a strong, fast (millisecond) outburst, originating from the Earth or a known celestial body.

At 14:19 UTC on 2 July 1967, the Vela 3 and Vela 4 satellites recorded an unprecedented gamma-radiation burst distinct from any previously identified astrophysical source. Despite the limited temporal resolution of the Vela satellites, GRB670702 exhibited fundamental characteristics typical of GRBs (such as a duration of approximately 10 s, a two-pulse light curve and a peak emission around the MeV energy band), very different from the properties for a nuclear test in space (a short-duration, hard, nonstructured X-ray burst).

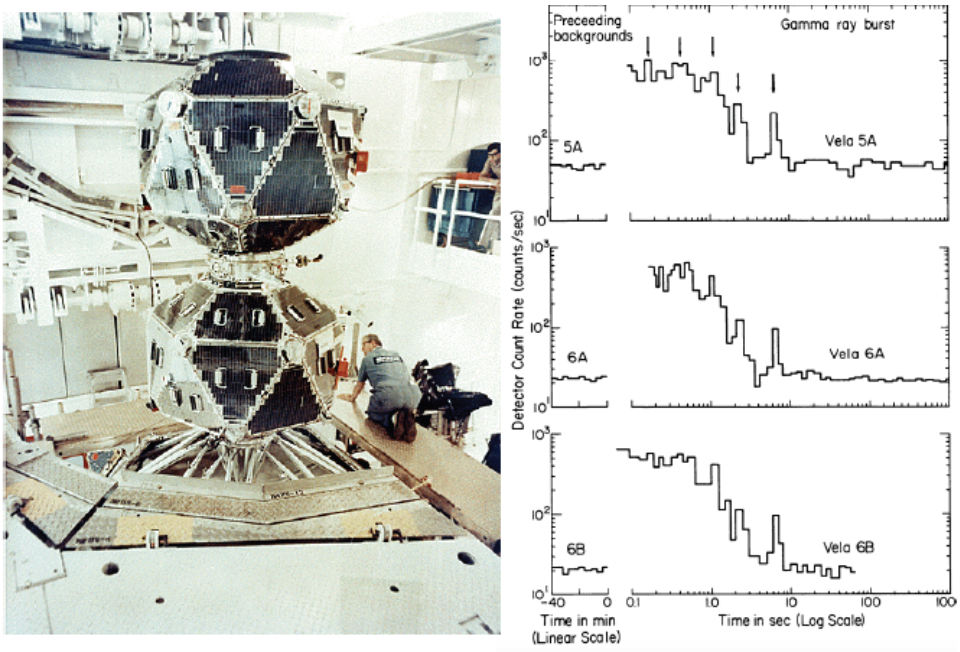

Figure 2: Left: A US Vela satellite, responsible for the first detections of GRBs in the late 1960s, Right: The lightcurve of GRB 700822 (22nd August 1970),

The poor localization capability of gamma-ray detectors made it very difficult to discover the electromagnetic counterparts of GRBs at lower frequencies. Still, it was clear that these signals, not thermal in origin, were not coming from Earth and not even from the Solar System plane. The proposed models for GRB creation could be broadly divided into two categories: Galactic and Cosmological models.

Galactic models, favored by the lower energies required in the processes, posited that the GRBs were produced either by the accretion/fall of material onto compact objects (supernova fallbacks, accretion from a binary companion onto white dwarfs or neutron stars, later observed as X-ray Novae) or through magnetic reconnection/quakes in the crusts of magnetars (neutron stars with magnetic fields reaching 1015 G).

In this debate, Bodhan Paczynski argued for a Cosmological origin. The BATSE missions had revealed about the isotropic distribution of GRBs in the sky. Such isotropy

naturally favours cosmological models, as it is inevitable in essentially all of them. Due to the large distances, extragalactic models in general had very high energy demands (with energies up to 1051 –1053 erg). This, in addition to implying a particularly low rate of events (10–6 –10–5 events per year in a galaxy with the luminosity of the Milky Way), has led to the proposal of exotic phenomena and/or objects (collapsing white dwarfs, neutron star mergers, active galaxy cores) among the possible origins of GRBs.

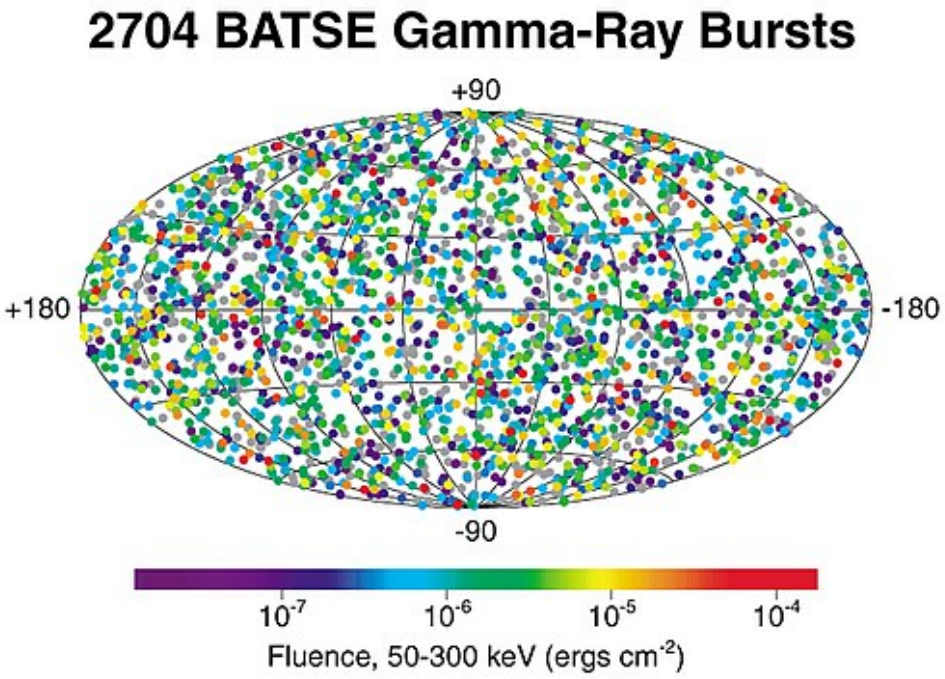

Figure 3: The distribution of BATSE GRBs on the sky, colour coded by the measured fluence of the burst. The bursts are distributed isotropically on the sky, and show no preference for any direction. This is true at any fluence level.

2.2 GRB phenomenology

A GRB is simply a burst of gamma-rays, an extremely energetic event which occurs in distant galaxies – the brightest and the most powerful class of explosion in the universe. These can last from a fraction of a second up to several hours. After the initial flash of gamma rays, a long-lived afterglow is emitted, usually in the longer wavelengths of X-ray, ultraviolet, optical, infrared, microwave or radio frequencies.

The promptness in the emissions has placed stringent limitations on the size of the region creating these GRBs. For $\Delta T = 1$ ms, c Δ T = 300 km, and so the inferred region is only a few hundred km across, greatly restricting the range of astronomical bodies that could be responsible for GRB creation. The light travel time arguments meant that theories of GRB creation needed to concentrate on compact regions, either individual dense stars (e.g. neutron stars or black holes) or subregions of larger objects, such as the cores of massive stars, or stellar coronae.

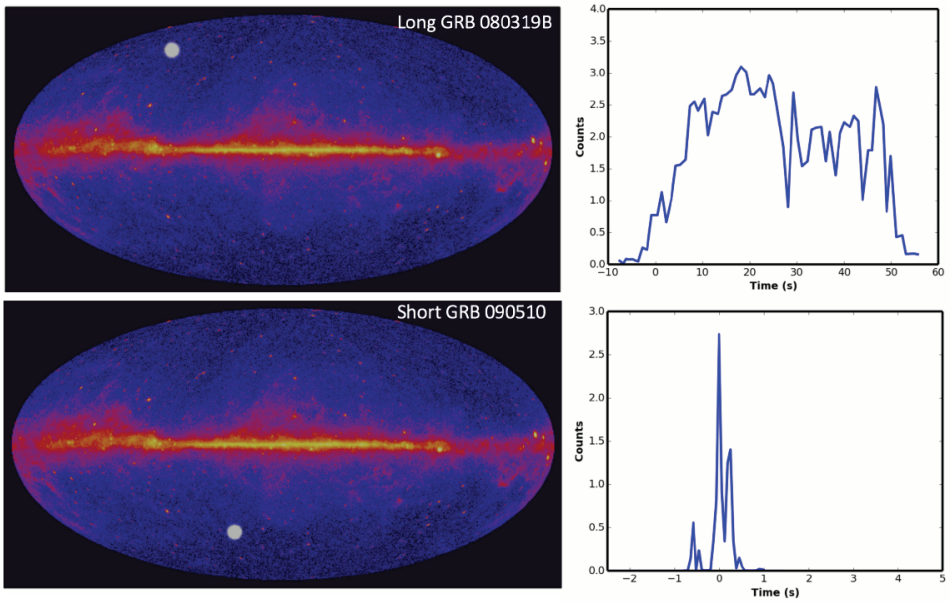

Figure 4: Simulations of the appearance of GRBs on the γ-ray sky (background from the Fermi Large Area Telescope). The bursts appear as the brightest sources in the γ-ray sky and then fade away. The right hand figures show the lightcurves of the two GRBs.

Another important feature is that GRBs do not repeat, they are one-off events, after which the system creating them is either destroyed or enters a much longer quiescent period (most likely the former is the case).

During the bursting phase the apparent brightness of the bursts exceeds the brightness of any other source in the γ-ray sky. It is particularly notable that GRBs emit the majority of their energy release in the γ-ray regime. In particular, they emit most of their power at photon energies of tens to hundreds of keV, while some have been detected at GeV energies. The actual flux output of a burst can be compared by looking at its energy spectrum, often referred to as the νF_ν or EF_E spectrum.

2.3 Improving GRB Positioning

Since the error regions determined by triangulation methods relying on satellites in Earth orbits were typically large, it was impossible to pinpoint the bursts on the sky. The typical error regions contained many thousands of Galactic stars and background galaxies, any of which could have hosted the GRB. Hence, these locations were not sufficient to provide means of distinguishing between the ever increasing number of models for GRB production.

One solution to this problem was to increase the baseline (distance between satellites) used in triangulation. By doing this, the differences in the time-of-flight for photons became increasingly longer, and so the region to which the burst can be located commensurately smaller. This was a major motivation to create the InterPlanetary Network (IPN). Since γ-ray detectors are relatively simple, and not too heavy, it was possible to place them on numerous spacecraft. With baselines $>$ 1AU, it would in principle be possible to get positions orders of magnitude better than previously possible. The hope was that these precise positions would be the decisive step towards the origin of GRBs.

2.4 The BATSE Era (1991 – 1996)

On 5 April 1991, the US spacecraft Compton Gamma Ray Observatory (CGRO) was launched in low Earth orbit. On board, there were four instruments that covered a huge energy range, ranging from 20 keV X-rays to 30 GeV gamma rays. Among these instruments, the Burst and Transient Source Experiment (BATSE) has led to substantial improvements in our understanding of GRBs. It consisted of eight scintillation gamma-ray detectors. Each module was placed on one of the corners of the satellite in order to obtain complete coverage of the sky, achieving a burst detection sensitivity of 3×10^{-8} erg cm^{-2} for a 1 s burst. BATSE was thus able to obtain approximate localizations of GRB signals based on the ratios of the count rates between the eight detectors.

Perhaps the most important result from BATSE was that GRBs appeared to be located isotropically on the sky – they come from no preferred direction. This marked the beginning of the end of galactic models. Still, they were not ruled out completely, and this led to the Great Debate in 1995.

2.5 Duration distributions of Gamma Ray Bursts

In 1996, Norris and colleagues found that the light curves of individual pulses are well fitted with a fast-rise exponential decay (FRED) profile with an average rise-to-decay ratio of 1:3 :

$$I(t) = \begin{cases} A \exp\left(-\left(\frac{|t-t_{\max}|}{\sigma_r}\right)^{\nu}\right) & t \leq t_{\max} \\ A \exp\left(-\left(\frac{|t-t_{\max}|}{\sigma_d}\right)^{\nu}\right) & t > t_{\max} \end{cases}$$

where t_{max} is the time of the pulse's maximum intensity; A is the peak value; σ_r and σ_d are the rise (t $\leq$ tmax) and decay (t $> t_{max}$) time constants, respectively; and ν is the "peakedness", a measure of the pulse's sharpness.

BATSE observed GRBs with very disparate durations, from $\sim$ 6 ms up to $\sim$ 2000 s. To classify the light curves, a duration measurement was introduced: T_{90}, the duration over which 90 % of the total fluence of a given burst is recorded. In essence, T_{90} can be obtained by integrating the observed GRB lightcurve, and then determining the times within the lightcurve when 5 % and 95 % of the total fluence has been observed. Furthermore, T_{90} is always defined in the observer frame, not in the rest frame of the burst. In principle this means that bursts observed at 2 s at some T_{90} at z = 0 would last for $T_{90}\times(1 + z)$ at some higher redshift.

With this definition, Kouveliotou and his colleagues showed that a GRB can be divided into two distinct groups: long bursts with $T_{90} > 2$ s and short bursts with $T_{90} < 2$s. When examining the spectral properties of the BATSE bursts, a distinction became clear: the shorter GRBs emit more high energy emission than the longer bursts. This leads to the identification of two populations of GRB – short hard bursts, and long soft bursts.

2.6 Fluence distributions of gamma-ray bursts

The large number of bursts detected by BATSE also allowed the distribution of the total fluence of the bursts to be studied . Of particular value is the so-called $\log N$ - $\log S$

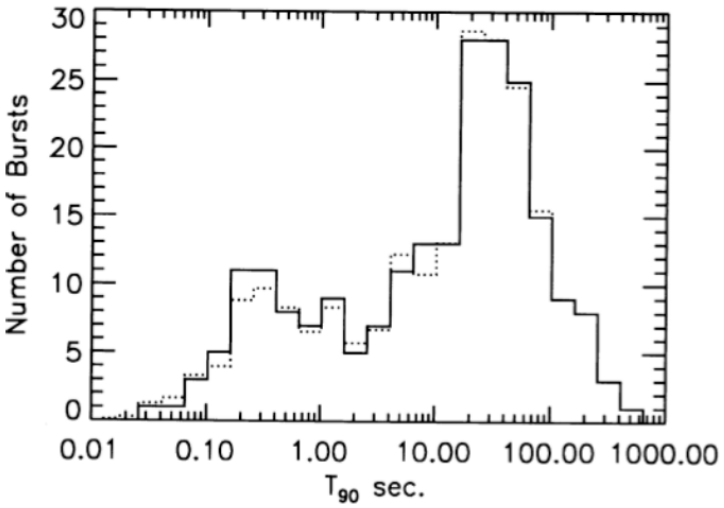

Figure 5: The duration distribution of BATSE GRBs. The short and long GRB populations are clearly visible.

distribution which describes the number of bursts (N) of a given fluence (S).

The principle of the $\log N$ - $\log S$ test is simple; imagine a set of standard candles, distributed uniformly in space. As the distance to the source increases the number of sources (N) would increase with the volume (as the cube of the distance). However, the relative fluence (S) observed from each source would decrease following the expected inverse square-law for light (1/distance2). Hence, the slope of the $\log N$ - $\log S$ would be expected to be -3/2.

The $\log N$ - $\log S$ for GRBs shows the expected -3/2 slope for brighter bursts, but deviates away at the fainter end, providing fewer faint bursts than expected. This deviation suggests that we are seeing the edge of the observable GRB population — a cosmological boundary.

This was one of the early key pieces of evidence that GRBs are not local (e.g., in the Milky Way), but instead come from very distant galaxies — they are cosmological in origin, and the effects of the expanding universe (e.g. redshift, time dilation, cosmic evolution) are at play.

2.7 The BeppoSAX Era (1996 – 2004)

As we have seen, much progress was made during the BATSE era in the investigations of GRBs. Still, regardless of the large amount of data acquired, the lack of multiwavelength observations and precise localization capabilities, together with the impossibility of estimating the distance of the sources, meant that it was not yet possible to understand the energy scales of these phenomena and completely rule out any of the two model paradigms that were proposed. This was aggravated by "the no-host problem": the impossibility of identifying the host galaxies of such signals. In order to move forward and solve these problems, it was therefore necessary to carry out a new type of observation. The key to making multiwavelength observations and estimating distances was to obtain the very precise localizations of the GRBs.

BeppoSAX was an Italian–Dutch satellite for X-ray astronomy. Launched in 1996, it offered two crucial enhancements for the detection of GRBs. The first was a wide field X-ray monitor that could deliver positions that were accurate to arcminutes, and the the second was the ability to repoint and observe the source position within an on-board X-ray telescope on a timescale of a few hours.

The BeppoSAX mission achieved a breakthrough in February 1997, within a year of its launch, by detecting the gamma-ray burst GRB 970228. Once the event's localization was pinpointed, subsequent deep imaging revealed a faint, very distant host galaxy in the GRB's location. This swiftly resolved the longstanding debate about the distance scale of GRBs, confirming them as extragalactic events originating within faint galaxies situated at staggering distances.

Since the X-ray telescopes and WFCs (Wide Field Cameras) were not aligned, BeppoSAX had to be slewed after locating each GRB. The procedure had to be carried out from the ground, leading to slow replacements. For this reason, in this era, predominantly (almost exclusively) long GRBs were observed.

2.8 The GRB-Supernova Association

The detection of the long-lasting, fading, multiwavelength afterglow emissions allowed the identification of the host galaxies, solving the "no host problem" and firmly establishing the origin of GRBs as extragalactic. The ability to obtain accurate localizations showed that the GRBs are located toward the edges of the host galaxies, i.e., in the star formation zones, and not toward the center. This fitted well with one of the most investigated scenarios, i.e., the GRB-supernova association.

The first hint of such an association was the discovery of SN 1998bw, a type Ic supernova in a galaxy near z = 0.0085. Furthermore, a supernova red bump was discovered in the optical light curves of several other GRBs. These findings led to the development of various models for the progenitors.

One of the earlier models, first proposed by Woosley in 1993, was the collapsar model. This proposed that the long GRBs are produced by the core collapse of massive, rapidly rotating stars. In 1998, Paczynski proposed a variant of Woosley's collapsar model, the so-called hypernova model: a rapidly rotating star collapses into a black hole surrounded by a thick accretion disk (or "torus"), which magnetically launches a relativistic jet, which, in turn, powers the observed GRB. Another variation of this scenario was proposed in the Black-Hole Accretion Disk (BHAD) Models, where the binary merger of two compact objects or the collapse of a rotating star produces a rapidly accreting disk around a "hyperaccreating" black hole.

2.9 The Swift Era (2004 – Ongoing)

In the early 2000s, the GRB phenomenon seemed almost completely decoded. BATSE observations had led to excellent modeling of gamma emissions, identifying two subpopulations of GRBs depending on their duration. Subsequent observations by BeppoSAX

led to the identification of multiwavelength afterglows and, through accurate localization, allowed long GRBs to be correlated with supernova emissions. These discoveries were accompanied by impressive and rapid theoretical advances that led to the development of, among others, the fireball shock model and the possible central engines of GRBs.

Still, there were some "details" of these phenomena that were not yet well understood. First of all, even though BeppoSAX had been able to observe a large number of GRB afterglows, these were just from long-duration bursts due to the requirement to slew the telescopes to observe the signals with the main cameras once the wide-field cameras were triggered. This operation, carried out from the ground, was very slow and therefore limited the observation possibilities only to long-duration GRBs and meant that there was a gap in the data acquired between the prompt-emission signals that triggered the wide-field cameras and afterglow observations after slewing. To fill these observational gaps, the Neil Gehrels Swift Observatory mission was launched on 20 November 2004.

Swift contains three main instruments, the Burst Alert Telescope (BAT), the X-ray Telescope (XRT) and the UV and Optical Telescope (UVOT). It is also able to autonomously slew rapidly to the locations of any detected GRB. This ability to quickly repoint the telescope and observe the initial stages of the afterglows has led to a true revolution in our understanding of these emissions.

The initial observations turned out to be rather different than expected: while long GRBs are generally found in blue, regular, and highly star-forming host galaxies and are located precisely in the star-forming regions of these galaxies, short GRBs, on the other hand, are hosted mainly by elliptical or irregular galaxies, far from star-forming regions. This led to the suggestion that short GRBs may have different progenitors than long GRBs, not being associated with the death of massive stars but with the coalescence of compact objects (such as neutron stars or black holes).

2.10 The Fermi Era (2008 – Ongoing)

In 2008, another NASA gamma-ray mission, the Gamma-ray Large Area Space Telescope (GLAST), later renamed the Fermi Gamma-Ray Space Telescope (FGST), was launched. Fermi is a purely γ-ray observatory, and does not have the ability to rapidly localise GRB afterglows. However, it is equipped with extremely powerful -ray detectors that have excellent sensitivity, and crucially a far wider energy range than the Swift-BAT.

Fermi has also provided a much better view of the broadband properties of the prompt γ-ray emission, where its wide energy coverage enables hitherto undetected features to be found. Perhaps most importantly, data has shown that the GRB spectrum extends to very high energies up to tens of GeV in energy. Interestingly, these photons often appear delayed with respect to the rest of the prompt emission, sometimes by thousands of seconds.

In addition to this, observations from Fermi show a clear thermal component with a temperature of tens to hundreds of keV. This goes in contrast with the non- thermal synchrotron processed GRBs as suggested by BATSE. This, at least in principle, provides a new route to understand the creation and evolution of the prompt emission.

2.11 The Birth of Multi-Messenger Era (2017 – Ongoing)

Multi-messenger emitters are astrophysical objects or events that produce multiple types of "messenger signals" – such as EMR, gravitational waves, neutrinos, cosmic rays; which we can detect on Earth to study the universe in more depth. GRBs have been believed to be multi-messenger emitters since 1989.

The crucial breakthrough for multimessenger astronomy occurred in August 2017, when Fermi detected a short-duration GRB approximately two seconds after LIGO and VIRGO observed the first merger of a binary neutron star system. Since short GRBs are rare, the connection between the two based on timing alone was highly suggestive. This became increasingly clear as location information became available. The position triangulated by the two LIGO detectors and VIRGO was in agreement with that from Fermi and with a refined electromagnetic position enabled by an independent detection of the GRB by the INTEGRAL satellite.

These pivotal observations finally opened the opportunity of multi-messenger astronomy using gravitational waves and visible light as a route to probe the final moments, and aftermath of the merger of a compact binary.

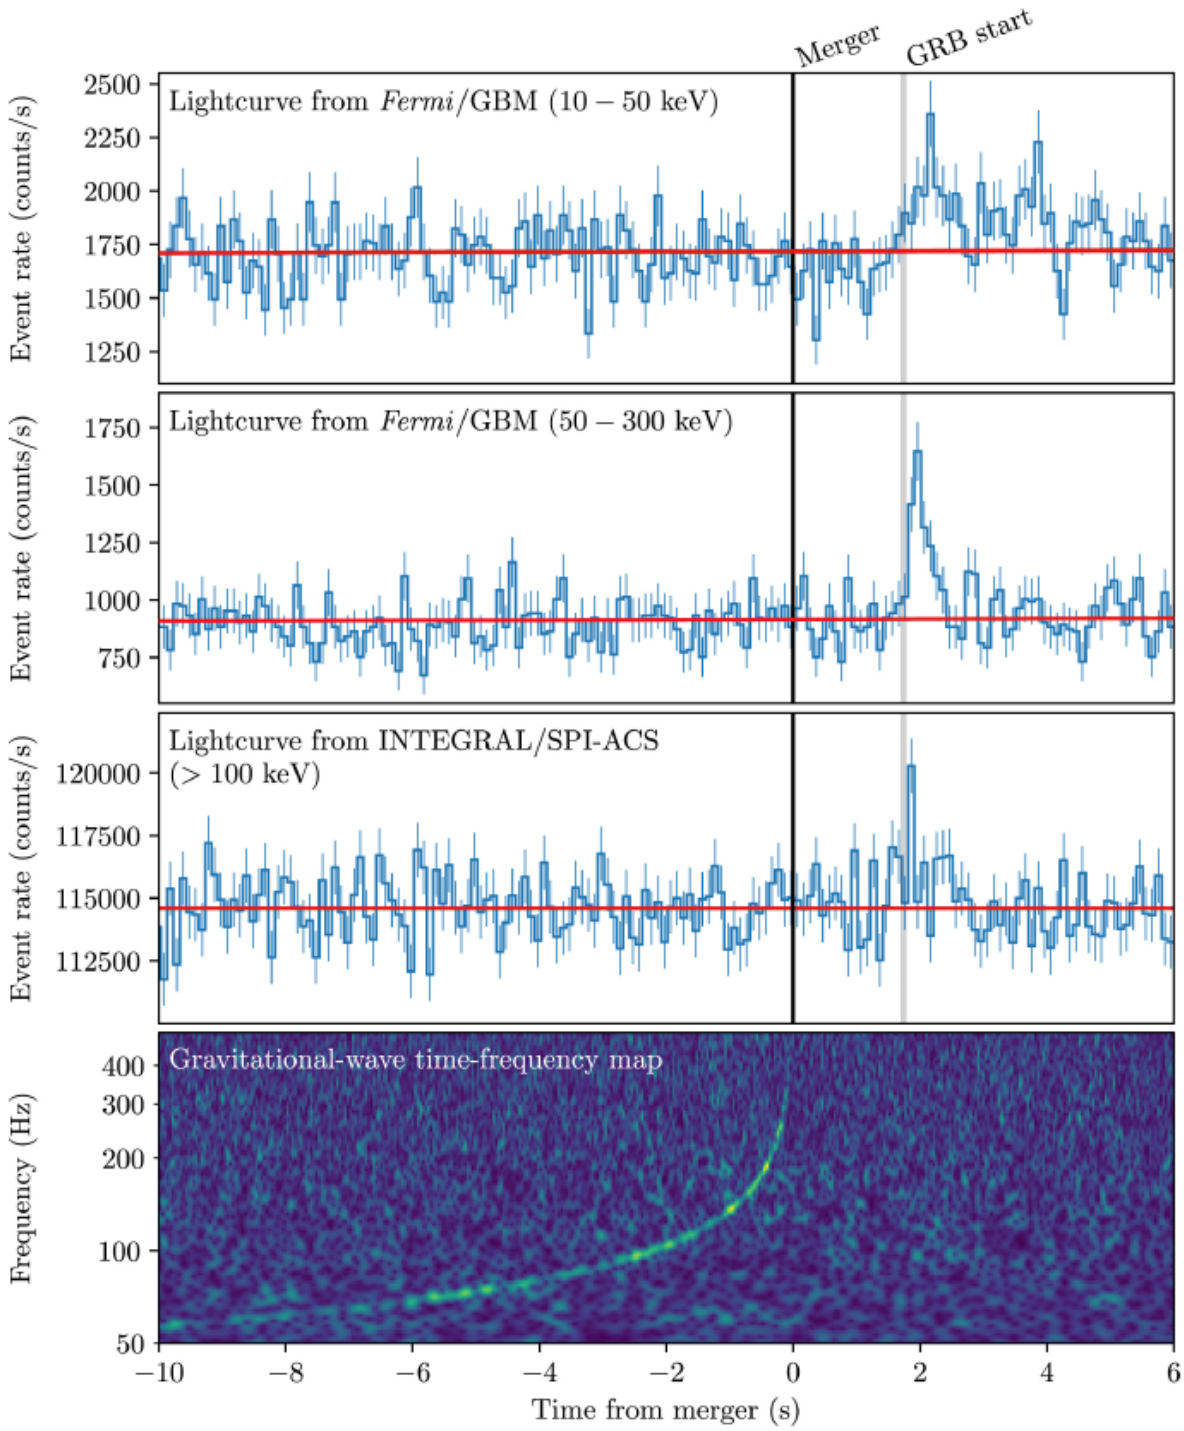

Figure 6: The detection of GW170817 and GRB170817A. The lower panel shows the frequency of the gravitational wave signal plotted against the time of the merger; the sweeping line is often referred to as the gravitational wave "chirp". The upper panels show the γ-ray light curves as recorded by Fermi (in two different energy bands) and Integral. The GRB occurred within ~ 2 seconds of the merger.

3 Prompt Emission

The prompt γ-ray emission is the defining characteristic of a gamma-ray burst. It is readily detected, even by rudimentary space-based gamma-ray detectors, due to the extreme high energy photon budget that GRBs exhibit. Indeed, at peak GRBs outshine all other sources within the γ-ray sky, including the Sun.

3.1 Temporal Structure

There is an often spoken adage in the GRB community: "When you've seen one GRB, you've seen one GRB". This reflects that the GRB prompt emissions vary hugely from burst to burst. No two bursts have the same duration and lightcurves.

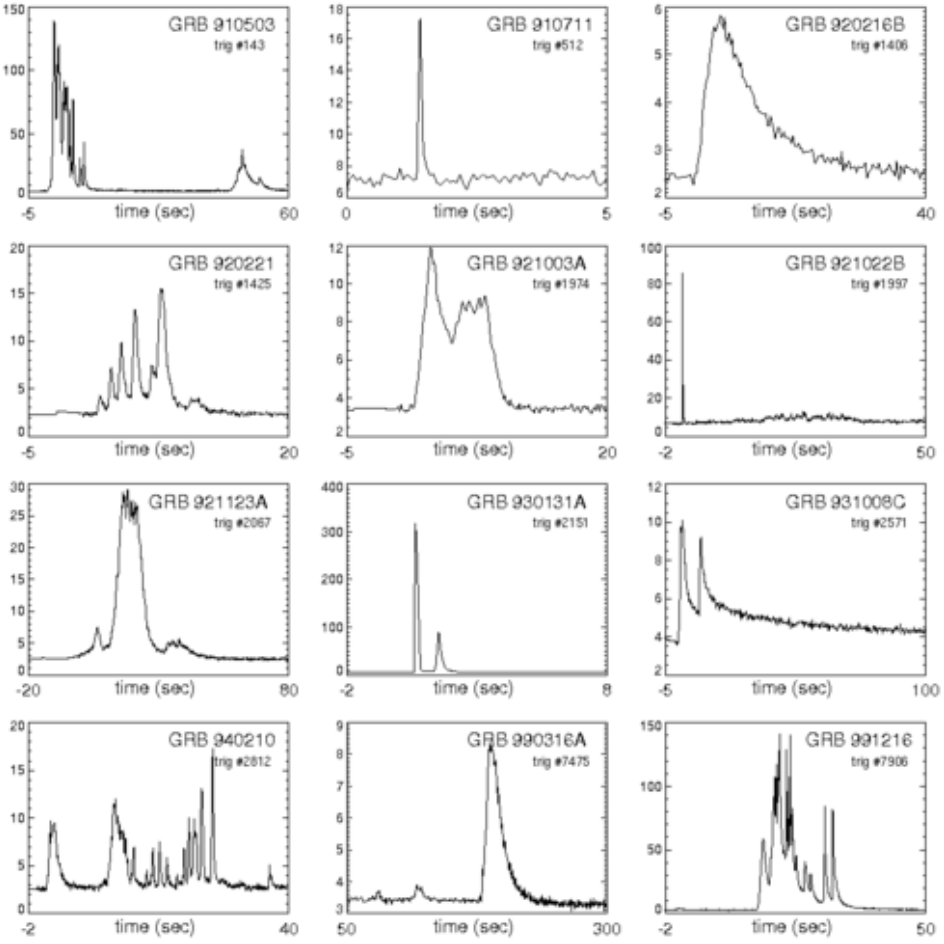

Figure 7: Prompt gamma-ray lightcurves for a set of Swift bursts. The wide range of durations and burst morphologies can be clearly seen, highlighting the diversity of GRB prompt emissions.

The prompt gamma-ray lightcurves of several well-known GRBs are shown in Figure 6. Some bursts are smooth, others highly erratic, some last for a fraction of a second, others for several hours. This includes highly variable spiky bursts, and those which exhibit a fast rise and exponential decay, a so-called FRED (Fast Rising Exponential Decay) profile.

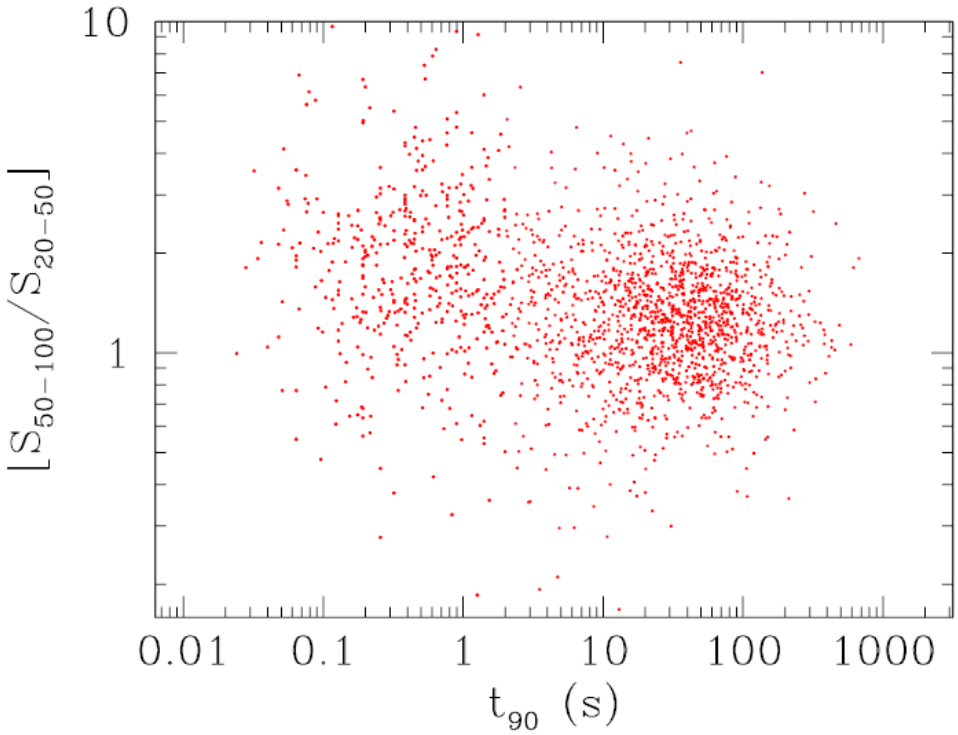

Figure 8: The hardness duration diagram for BATSE GRBs demonstrating the divide between short-hard bursts (t_{90} ¡ 2s and long-short bursts ($t_{90} >$ 2s). It is apparent that the median hardness of the short bursts is significantly higher than for the long burst population.

Perhaps the most natural property to consider is the duration of the burst, measuring how long the gamma-rays are visible. The most commonly used measure of duration is the so-called T_{90}. The advantage of a duration measurement is that it is independent of the morphology of the lightcurve. From the measured durations, it became apparent that there are at least two broad populations of GRBs: long bursts with T90 $>$ 2 s and short bursts with T90 $<$ 2s.

Examination of the spectral properties of the different populations suggests that short bursts emit more high energy radiation, while long bursts somewhat less. This is often encapsulated in a so-called hardness ratio, a ratio of photons observed in two BATSE channels: channel 3 (100–300 keV) counts divided by channel 2 (50–100 keV) counts.

$$HR_{32} = \frac{S_3(100-300\,\mathrm{keV})}{S_2(50-100\,\mathrm{keV})}$$

The two groups also have different spectral characteristics, with short GRBs tending to be spectrally “harder” (i.e. exhibiting more higher energy photons) than long GRBs.

There is evidence that in the long bursts, the higher energy emission arrives earlier than the lower energy emission. This so-called “spectral lag” is generally positive (high energy leads, low energy lags) for long bursts, and is normally zero for short bursts. Hence, a GRB might look shorter in high-energy gamma rays, and longer in lower-energy X-rays. This means that the duration of a burst as recorded is a function of the energy range and sensitivity of the detector used to make the observations.

3.2 Spectral Structure

Spectral studies of BATSE GRBs have shown that the emission is non-thermal, with spectra generally characterized by a double power law and a smooth break. Band and colleagues proposed a phenomenological function, now known as the Band function, to fit GRB spectra:

$$F_{\text{Band}}(E) = \begin{cases} A\left(\frac{E}{100\,\text{keV}}\right)^{\alpha} \exp\left(-\frac{E(2+\alpha)}{E_{\text{peak}}}\right) & E < E_c \\ A\left[\left(\frac{(\alpha-\beta)E_{\text{peak}}}{100\,\text{keV}(2+\alpha)}\right)^{\alpha-\beta} \exp(\beta-\alpha)\left(\frac{E}{100\,\text{keV}}\right)^{\beta}\right] & E \geq E_c \end{cases}$$

$$\text{where} \quad E_c = (\alpha-\beta)\frac{E_{\text{peak}}}{2+\alpha} \equiv (\alpha-\beta)E_0$$

α is the low-energy index, β is the high-energy index, E_c is the break energy, and E_{peak} is the peak energy of the νF_ν spectrum. No specific theoretical model predicts this spectral shape; however, the Band function provides an excellent fit to most observed spectra.

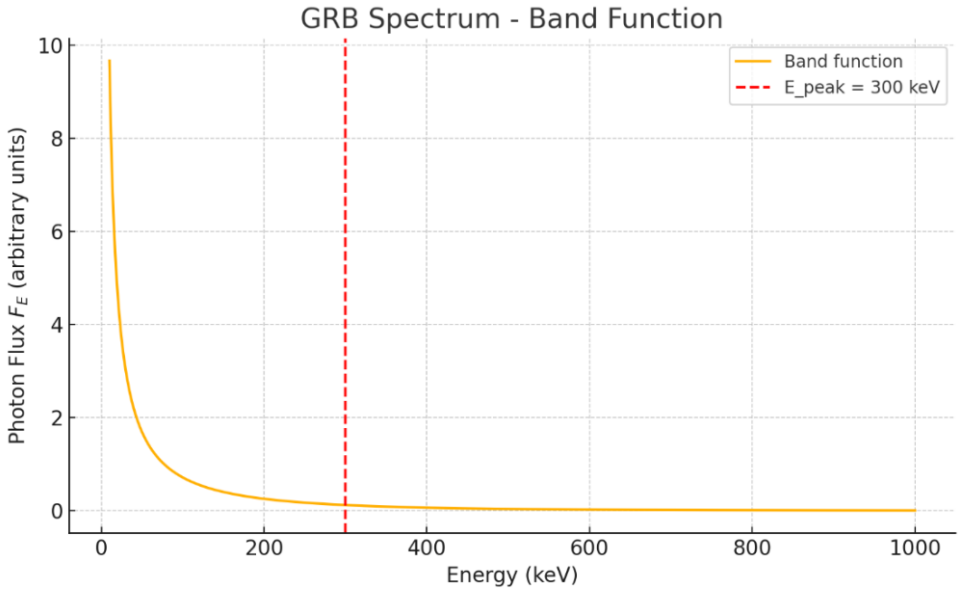

Figure 9: The band function smoothly transitions between a low-energy power law with index $\alpha = 1$ and a high-energy power law with index $\beta = 2.5$.

An example of such a spectrum from the famous burst GRB 990123 which reached a peak visual magnitude of ~ 9 is shown in Figure 9. This demonstrates the broken power-law appearance and the clear peak energy. It also shows that over wide ranges of energies the source will be observed as a single power-law.

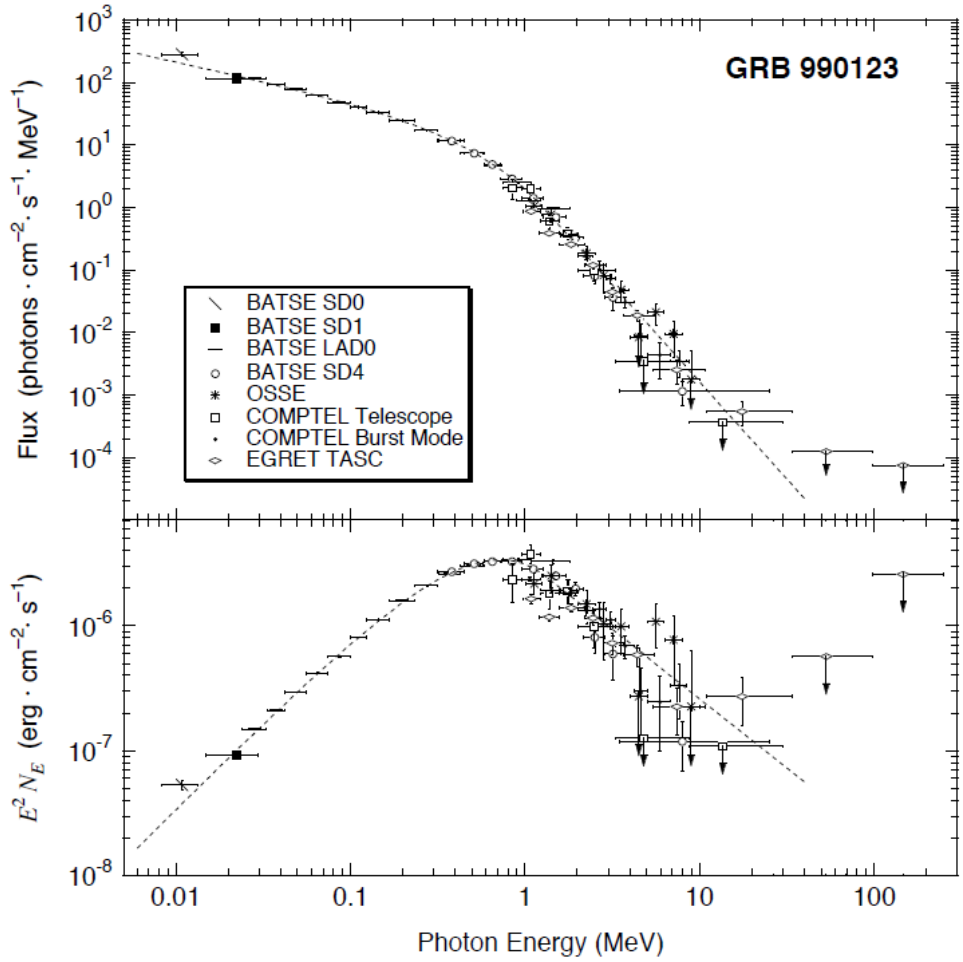

Figure 10: The BATSE spectrum of the famous GRB 990123. The burst shows an excellent example of the Band spectrum characterised by a low energy slope of α = -0.6 and β = -3.1, with a peak energy of E_{peak} = 720 keV. Instruments such as the Swift-BAT (energy range 15-350 keV) often only sample one of the power-law regions.

One crucial discovery made by Fermi was that GRBs do exhibit extremely high energy emission. This is not naturally expected from the Band model alone, since it predicts that the energy above the peak should cut-off rapidly. Instead, the very high energy emission would appear to form a separate component. (See Figure 10).

3.3 Origin of the prompt emission: the "fireball" shock model

The evidence (in particular the spatial distribution) gathered in the early 1990s lent credence to cosmological models of GRBs. However, in this scenario, the burst energies are particularly challenging, with isotropic energy releases of up to 1054 erg placed within a region only a few km across; the energy density would be extreme.

This was solved with the development of the fireball shock model. According to this scheme, the observed gamma-rays are emitted when an ultra-relativistic energy flow is converted to radiation. The "central engine" (possibly a black hole or neutron star) that produces the relativistic energy flow is hidden from direct observations. The central engine produces a relativistic "fireball" – a highly energetic, expanding shell of plasma

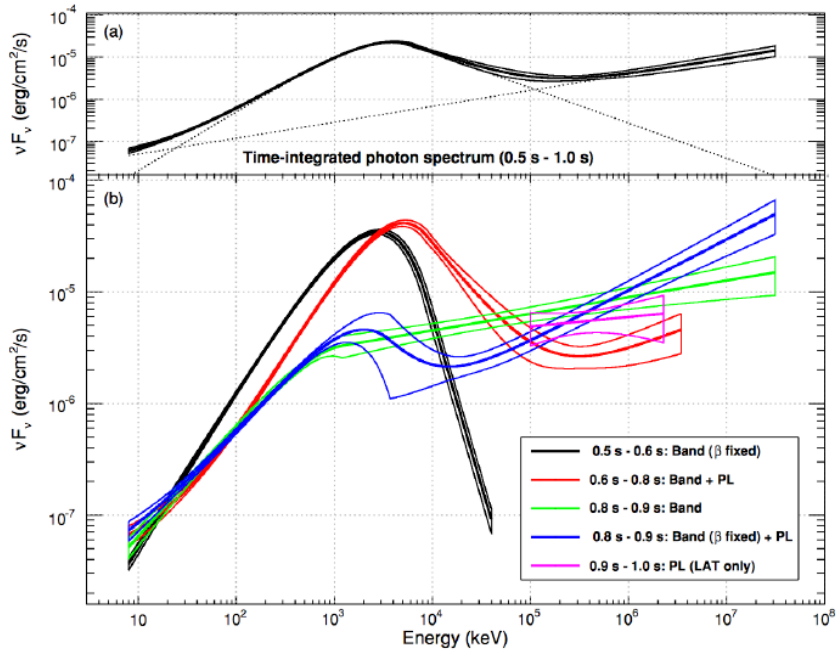

Figure 11: The spectrum of the short GRB 090510, as observed by the Fermi, with a notable further upturn at very high GeV energies, inconsistent with the extrapolation of Band function.

travelling at nearly the speed of light. Although there is little matter entrained within the jet, the high velocities mean there is significant kinetic energy, and so much of the thermal energy initially deposited at the base of the jet has been converted.

The fireball is not perfectly smooth and variations in the outflow velocity lead to collisions and formation of internal shocks. Since all of these velocities are very close to the speed of light, these shocks do not occur close to the energy source, but instead at much larger radii. These shocks accelerate the particles to relativistic speeds, and the accelerated particles interact with magnetic fields to produce gamma-ray radiations. . If these are emitted with a time gap of δt and with Lorentz factors $\Gamma 1$ and $\Gamma 2$ then time in which they will interact is given by $\delta t \Gamma 1 \Gamma 2$ and the distance from the source (given that the velocity is approximately the speed of light) is given by $c\delta t \Gamma 1 \Gamma 2$. This, for a time difference of 100 ms, and $\Gamma 1 \approx \Gamma 2 \approx 100$, corresponds to 3×10^{11} m ~ 2 AU. When two shocks combine, some fraction of the initial kinetic energy is dissipated in the shock, and this is given by

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

Not all the energy of the relativistic shell is converted to radiation by internal shocks. As the jet continues through the circumstellar medium, it sweeps up the material forming a further external shock between the head of the jet and the surrounding matter. This interaction creates shocks in which electron acceleration is possible. The remaining kinetic energy is most likely dissipated via these external shocks that produce an "afterglow" in different wavelengths.

Indeed, while GRBs are extremely compact their afterglows are generated at far larger

radii, but appear at the same time as the burst (or more accurately shortly thereafter) because they are generated by material travelling towards the observer at near the speed of light.

3.4 GRB Energetics

At the distances now measured for GRBs, they are extremely energetic. A common metric for determining these energies is the *isotropic equivalent energy release* ($E_{\rm iso}$); this is the energy that the burst would have if it emitted the energy observed by us to observers in all directions. It can be calculated from the measured fluence (S_γ) following:

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

where d_L is the luminosity distance to the burst and z is the redshift. At the typical redshifts of GRBs, the inferred isotropic energies of long bursts are in the region $10^{52} < E_{\rm iso} < 10^{54}$ erg. In other words, they suggest orders of a solar mass of material is rapidly transferred into radiation, on timescales of only a few seconds. The solution lies in the highly relativistic and beamed nature of GRBs.

4 Afterglow Emission

Gamma-ray burst afterglows are long lived (seconds-days-weeks) emission viewed in the aftermath of the burst. They appear to emit synchrotron emission (produced when relativistic (near-light-speed) charged particles, such as electrons, spiral around magnetic field lines) across the electromagnetic spectrum – often leading to high polarisation of the prompt emissions. The afterglows are generally accepted as the signature of the shock formed between the outflowing relativistic jet and the circumstellar medium around the progenitor system. It is the afterglows that precisely locate GRBs on the sky, enable redshift measurements and identify host galaxies.

4.1 X-ray Afterglows

X-ray afterglows are now almost common to GRBs. A typical X-ray afterglow can be broken down into several different phases: (See Figure 11)

1. *Prompt emission*: This is typically extremely bright, and can be highly variable (both temporally and spectrally).

2. *Rapid Decay*: At the end of the prompt emission the visible light appears to undergo a rapid decay, often to t^{-5} or steeper.

3. *Plateau*: Following the rapid decay, the X-rays often appear to plateau, appearing as flat, or with a slow decay ($< t^{-0.5}$).

4. *Flares*: During this plateau phase (or potentially earlier or later), it is possible for bursts to exhibit small or large X-ray flares.

5. *Intermediate Decay*: At the end of the plateau, the X-ray afterglow will typically begin to fall as t^{-1}.

6. *Rapid Decay*: After this slower decay, burst lightcurves will eventually undergo a temporal break in which the decay steepens. Beyond this point, the flux of the counterpart typically decays as t^{-2}. This temporal break is often referred to as the "jet break".

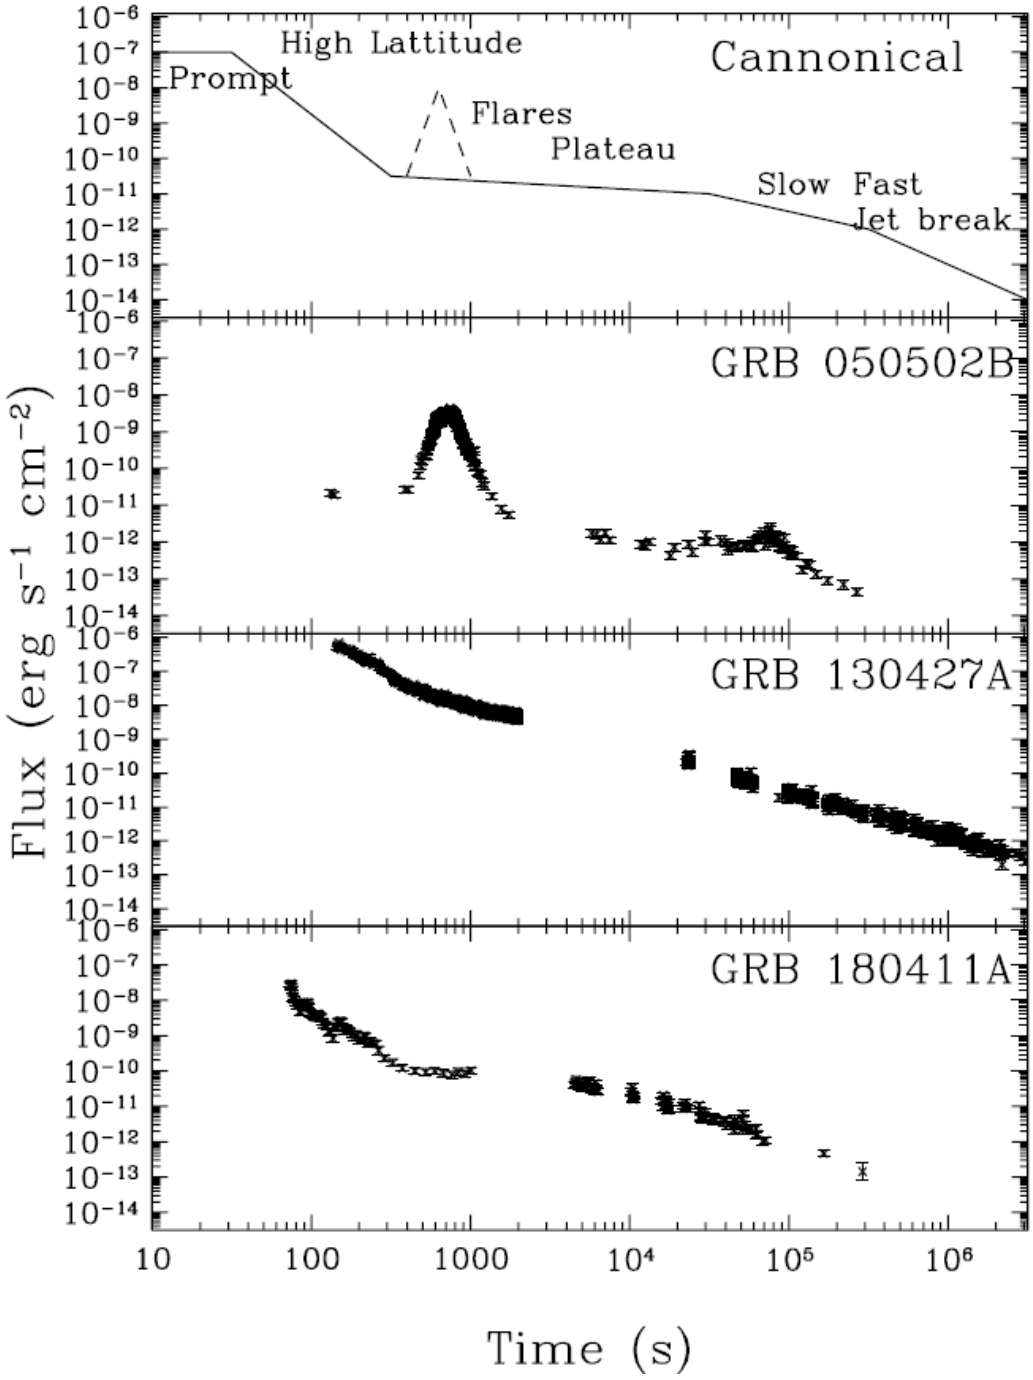

Figure 12: The canonical lightcurve of an X-ray afterglow (top) with typical features marked. Also shown are several GRBs which exhibit some of these features.

4.2 Radio Afterglows

Longer wavelength counterparts to GRBs are rather rarer than those seen in the X-ray or the optical/IR. Despite this, a good number of radio counterparts have been found, beginning with GRB 970508, the burst which also yielded the first GRB redshift.

Unlike the X-ray and optical afterglows, which are typically at their brightest at the time of the first observations; radio afterglows often rise with time, reaching their peak brightness days to weeks after the burst itself.

4.3 Synchrotron emission

The dominant emission process that shapes GRB afterglows is the emission of synchrotron radiation from the shock-front where the ultra-relativistic outflow impacts the surrounding circumstellar medium about the progenitor. At this shock front relativistic electrons are accelerated and gyrate about the magnetic field. As these electrons gyrate around magnetic field lines they emit synchrotron radiation, and this produces emission with a range of photon energies that spans the electromagnetic spectrum.

When the shockwave propagates through the medium, electrons are swept up and shocked, this results in a distribution of Lorentz factors amongst the electrons, $N(\gamma_e) \propto \gamma_e^{-p}$, where $p > 2$ to avoid infinite energy when integrating the electron energy, and is typically taken to be in the region $2 < p < 2.5$.

At low frequencies, the source behaves essentially as a black body because emitted photons are re-absorbed. This occurs at frequencies lower than the so-called *self absorption frequency*, ν_a. Below this frequency, the dependence on the specific flux per unit frequency (F_ν) with frequency is given as $F_\nu \propto \nu^2$. At higher frequencies, the emission scales as $F_\nu \propto \nu^{1/3}$ up to the *peak frequency*, ν_m. The final break frequency is related to the cooling of the electrons as they gyrate around the field, the *cooling frequency* ν_c. The electrons are shock accelerated, but cool as they radiate. At frequencies higher than the cooling frequency the emission scales as $F_\nu \propto \nu^{-p/2}$.

4.4 Reverse shock

In addition to the outgoing forward shock that ploughs into the interstellar medium, there is also a so-called reverse shock, which propagates backward through the ejecta. In many ways "reverse" shock can be misleading, this shock is still moving outward from the restframe of the star at relativistic velocity, but is ploughing backward through the ejecta in their frame, and is slowing the expansion.

The identification of signatures of the reverse shock is challenging. They occur early in the afterglow when the forward shock is strong, and for long bursts while the prompt emission may continue. A clear identification hence requires the accurate decoupling of these differing emission processes, and hence comprehensive observations across the electromagnetic spectrum from early to late times.

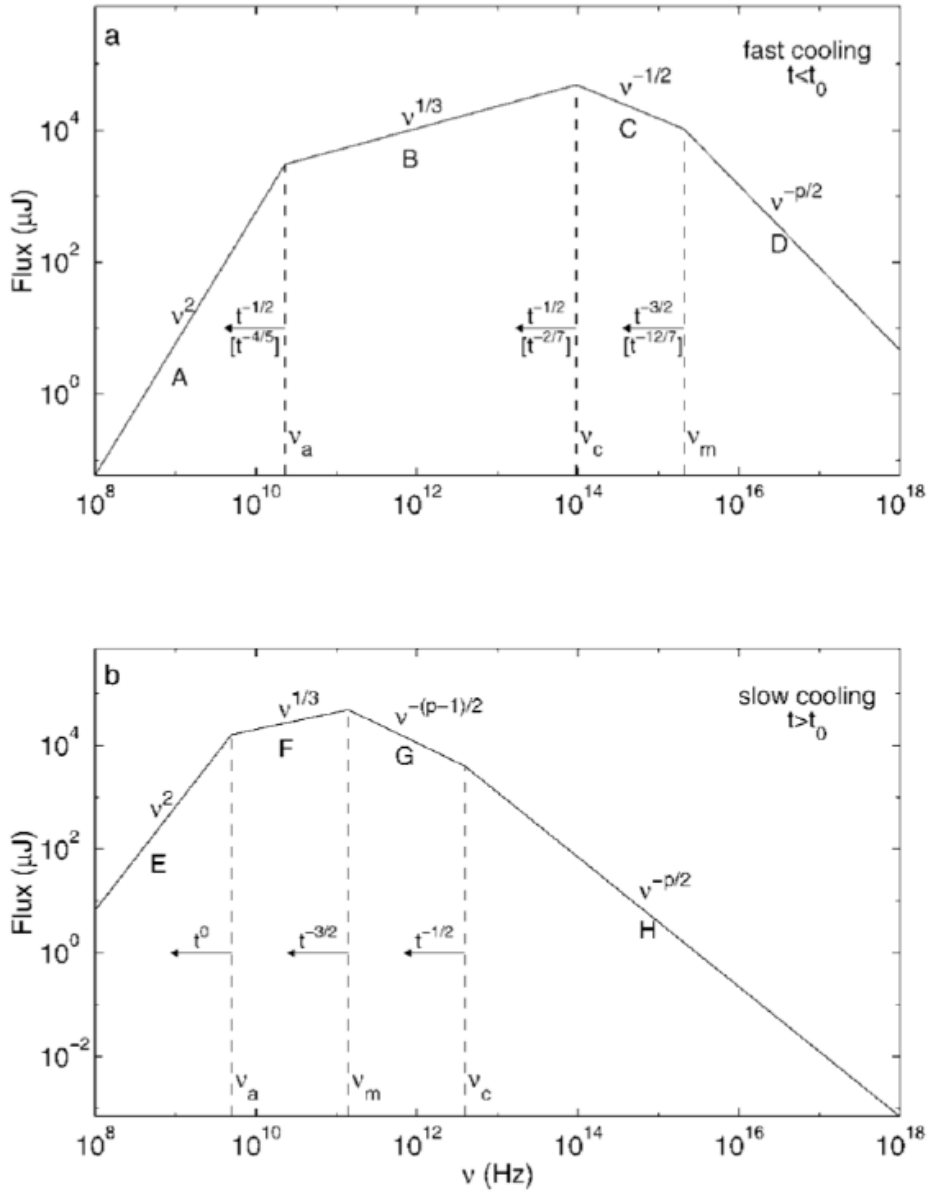

Figure 13: The predicted synchrotron spectral energy distributions for GRB afterglows, flux plotted against frequencies in Hz.

4.5 Relativistic Beaming

The above discussion is valid for isotropic emission, although for the early phases of the GRB afterglow still provides a good description even if the GRB is highly collimated. Because of relativistic beaming, they only illuminate a small fraction of the sky, which can be expressed through the beaming factor, $f_b \sim \theta_j^2/2$, where θ_j is the half opening angle of the jet (the angle from the polar axis to the edge of the jet) in radians. The correction from the measured energy of the burst $E_{\rm iso}$ to its true energy E_γ is given by

$$E_\gamma = E_{\rm iso}(1 - \cos\theta_j) \approx E_{\rm iso}\frac{\theta_j^2}{2},$$

where the beaming fraction f_b is given as $(1-\cos\theta_j)$. As time passes and the jet slows down there eventually reaches the point where $\Gamma = 1/\theta_j$. At this point, the two sides of the jet are in causal contact – i.e., light signals (or pressure waves) can travel from one side to the other within the jet. Beyond this crucial point, the afterglow is expected to fade more rapidly, yielding a steepening of the slope, and a so-called "jet break" – a sudden steepening in brightness decay.

The most simple model of the relativistic jet is of a so-called top hat jet, where the energy is the same across the jet, and drops to zero at sharp edges. In particular, in the post jet-break regime the jet becomes visible to viewers at progressively larger angles, including those who were not originally within the beam.

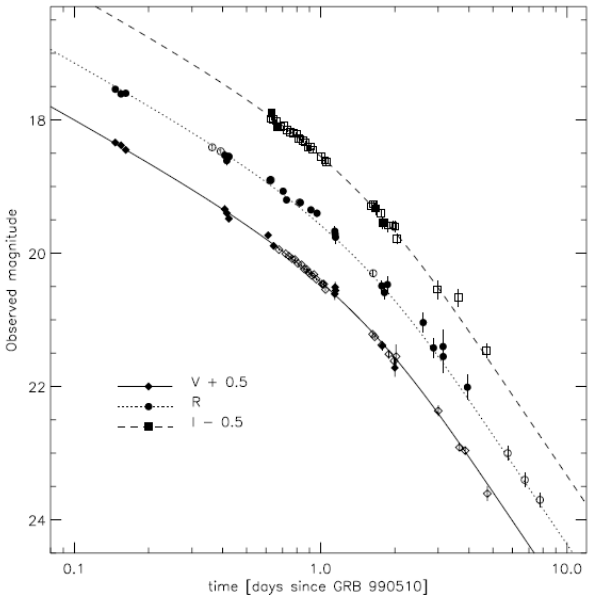

Figure 14: An early example of a jet-break in a GRB afterglow, GRB 990510. The clear, achromatic steepening of the afterglow is interpreted as the jet break at the point that $\Gamma = 1/\theta_j$, and provides a direct route of measuring GRB collimation.

Part II: Data Analysis of Gamma Ray Bursts

The study of Gamma-Ray Bursts (GRBs) relies heavily on the careful analysis of observational data collected across multiple wavelengths and timescales. From the prompt gamma-ray emission to the long-lasting afterglows observed in X-ray, optical, and radio bands, each phase of a GRB encodes critical information about the burst mechanism, energetics, and surrounding environment. Data analysis plays a pivotal role in extracting physical parameters such as the burst duration, spectral features, redshift, fluence, and variability, which in turn help classify GRBs and test theoretical models. With the advent of high-resolution instruments aboard satellites like *Swift*, *Fermi*, and *INTEGRAL*, modern GRB analysis has become increasingly sophisticated, involving time-resolved spectroscopy, lightcurve modeling, and population studies. This section outlines the key methodologies and tools used in GRB data analysis, laying the groundwork for interpreting the rich dataset these cosmic explosions provide.

5 MCMC Techniques for GRB Parameter Estimation

5.1 Introduction to Markov Chains

Markov chains are mathematical models that use concepts from probability to describe how a system changes from one state to another. These were developed by a Russian mathematician A. A. Markov about 100 years ago.

A Markov chain is a stochastic process describing a sequence of possible events in which the probability of each event depends only on the state attained in the previous event. Informally, this may be thought of as, *"What happens next depends only on the state of affairs now."*

Mathematically speaking, a **Markov Chain (MC)** is a sequence of random variables $X = \{X_0, X_1, X_2, \ldots\}$ with the following properties. For each $k \in \{0, 1, 2, \ldots\}$, X_k is defined on the sample space Ω and Ω takes values in a finite set S. Thus,

$$X_k : \Omega \to S$$

Also, for $k \in \{0, 1, 2, \ldots\}$, and $i, j, i_1, i_2, \ldots, i_k \in S$,

$$\mathbb{P}(X_{k+1} = j \mid X_k = i, X_{k-1} = i_{k-1}, \ldots, X_0 = i_0) = \mathbb{P}(X_{k+1} = j \mid X_k = i) \tag{1}$$

and the transition probabilities

$$\mathbb{P}(X_{k+1} = j \mid X_k = i) = p_{ij} \text{ are independent of } k. \tag{2}$$

Essentially, we have a system that changes from one state to another over time. The various states of the system are contained in the set S. The set S is called the **set of states or state space** of the Markov chain X. If the set S has m elements, then X is called an m**-state Markov chain**.

- **Condition (1)** means that, if the system is in state i at time k, then the probability it will change to state j at time $k + 1$ does not depend on what happened in earlier times $k-1, k-2, \ldots, 0$. This is usually called the **Markov property**. One might say, loosely, that the next state of the system depends only on the present state and not on earlier states — the process is **memoryless**.

- **Condition (2)** means that, if the system is in state i at time k, then the probability that it will change to state j at time $k+1$ does not depend on k. We say the transition probabilities p_{ij} do not vary with time (k). This is called the **stationary condition**.

An example of a Markov chain is the game "*Snakes and Ladder*".

5.2 What are Monte Carlo methods?

Monte Carlo methods are a broad class of computational algorithms that rely on repeated random sampling to obtain numerical results. The underlying concept is to use randomness to solve problems that might be deterministic in principle.

For example, consider a circle inscribed in a square. Given that the ratio of their areas is $\pi/4$, the value of π can be approximated using the Monte Carlo method:

1. Draw a square, then inscribe a circle within it.
2. Uniformly scatter a given number of points over the square.
3. Count the number of points inside the circle.
4. The ratio of the inside-count and the total-sample-count is an estimate of the ratio of the two areas, $\pi/4$.
5. Multiply the result by 4 to estimate π.

There are two important considerations:

1. If the points are not uniformly distributed, the approximation will be poor.
2. The approximation improves as more points are randomly placed in the whole square.

Hence, the use of Monte Carlo method requires large amounts of random numbers, benefitted greatly from pseudorandom number generators.

5.3 Markov Chain Monte Carlo (MCMC) Sampler

In principle, Monte Carlo methods can be used to solve any problem having a probabilistic interpretation. By the law of large numbers, integrals described by the expected value of some random variable can be approximated by taking the empirical mean (a.k.a. the 'sample mean') of independent samples of the variable. When the probability distribution of the variable is parameterized, mathematicians often use a Markov chain Monte Carlo (MCMC) sampler. The central idea is to design a judicious Markov chain model with a prescribed stationary probability distribution. That is, in the limit, the samples being generated by the MCMC method will be samples from the desired (target) distribution. By the ergodic theorem, the stationary distribution is approximated by the empirical measures of the random states of the MCMC sampler.

The Markov Chain Monte Carlo (MCMC) sampling method was employed to explore the posterior distribution of model parameters in the analysis of Gamma-Ray Burst (GRB) data. MCMC is particularly well-suited for this task due to its ability to handle

non-linear, multi-parameter models commonly used in GRB studies, such as the Band function or broken power law.

The MCMC method constructs a chain of samples from the parameter space, with each sample weighted by its likelihood given the observed data. The posterior distribution $P(\theta|D)$ is derived using Bayes' theorem:

$$P(\theta|D) \propto \mathcal{L}(D|\theta)P(\theta) \tag{3}$$

where:

- θ represents the vector of model parameters,
- D denotes the observed spectral data,
- $\mathcal{L}(D|\theta)$ is the likelihood function,
- $P(\theta)$ represents the prior probability distribution for the parameters.

The Bayesian approach through MCMC sampling offers significant advantages for GRB spectral analysis. By construction, it provides robust parameter estimation even in cases of parameter degeneracy or poorly constrained fits, as the full posterior distribution captures all plausible solutions weighted by their probability. The method naturally incorporates prior knowledge through the $P(\theta)$ term, which is particularly valuable when analyzing low-count spectra where physical constraints can guide the fit. The complete characterization of parameter uncertainties via the posterior distributions enables more rigorous error analysis compared to traditional fitting methods. Furthermore, the approach allows for direct computation of credible intervals for all derived quantities through sampling of the posterior, providing a statistically rigorous framework for propagating uncertainties through any subsequent calculations or model comparisons.

5.4 Metropolis-Hastings MCMC Implementation

The Metropolis-Hastings algorithm, a foundational Markov Chain Monte Carlo (MCMC) method, was employed for Bayesian parameter estimation in the GRB spectral analysis. This approach constructs a Markov chain whose stationary distribution converges to the target posterior distribution $P(\theta|D)$ through an iterative accept-reject scheme:

$$P(\theta|D) \propto \mathcal{L}(D|\theta)P(\theta)$$

Algorithm Specification

The implementation proceeded as follows:

1. **Initialization**: Starting parameters θ_0 were drawn from the prior distributions $P(\theta)$
2. **Proposal Mechanism**: At each iteration t, new parameters θ' were generated from a symmetric proposal distribution $q(\theta'|\theta_t)$, implemented as a multivariate Gaussian:

$$q(\theta'|\theta_t) = \mathcal{N}(\theta_t, \Sigma) \tag{4}$$

where Σ is a covariance matrix tuned for optimal acceptance rates.

3. **Acceptance Criterion**: The proposed state was accepted with probability:

$$\alpha = \min\left(1, \frac{\mathcal{L}(D|\theta')P(\theta')}{\mathcal{L}(D|\theta_t)P(\theta_t)}\right) \tag{5}$$

There are many MCMC methods more advanced than M–H. However, they all share characteristics with M–H (initialization, tuning, judging convergence), such that it is very valuable to understand M–H well before using anything more advanced. Because many problems involve a huge amount of dynamic range in the density function f(θ), and we like to avoid underflows and overflows (in, say, ratios computed for accept-reject steps), it is often advisable to work in the (natural) logarithm of the density rather than the density. When working in logarithmic density, the accept–reject step would change from a comparison of a random deviate with a ratio of probabilities to comparison of the logo of a random deviate with a difference of log probabilities.

This implementation provided efficient sampling of the typically correlated, non-Gaussian parameter spaces encountered in GRB spectral modeling, while maintaining detailed balance to ensure convergence to the true posterior distribution.

6 MCMC Sampling using emcee package

To perform efficient Bayesian inference on the GRB data, we employed `emcee`, a widely-used Python library for Markov Chain Monte Carlo (MCMC) sampling. `emcee` implements the affine-invariant ensemble sampler proposed by Goodman and Weare (2010), which is particularly effective for exploring high-dimensional parameter spaces with complex, correlated posteriors. Unlike traditional MCMC methods such as Metropolis-Hastings, `emcee` uses an ensemble of walkers that simultaneously sample the target distribution, leading to faster convergence and improved sampling efficiency. This made it especially suitable for our parameter estimation tasks, where likelihood surfaces can be non-trivial and multi-modal. The sampler's Pythonic interface and built-in parallelization further simplified its integration into our data analysis pipeline.

The affine-invariant ensemble sampler `emcee` [**?**] was employed for Bayesian parameter estimation of GRB spectral models, following the methodological framework demonstrated in the line-fitting tutorial but extended to accommodate the complexities of gamma-ray burst spectroscopy.

6.1 Ensemble Sampling Architecture

The implementation utilizes `emcee`'s parallel walker architecture:

$$\mathcal{P}(\theta|D) \propto \prod_{i=1}^{N_{\text{chan}}} \text{Poisson}(D_i|M_i(\theta)) \times P(\theta) \tag{6}$$

where:

- N_{chan} is the number of energy channels
- $M_i(\theta)$ represents the predicted counts for spectral model M with parameters θ
- Poisson likelihood replaces Gaussian likelihood for proper count statistics

6.2 Prior Handling and Parameter Transformations

The spectral analysis incorporated physically motivated priors through strategic parameter transformations, ensuring both sampling efficiency and physical plausibility.

For the Band function parameters, the low-energy spectral index α was transformed as $\log(\alpha - \alpha_{\min})$ to enforce the physical constraint $\alpha > \alpha_{\min}$, where $\alpha_{\min} = -2.0$ represents the hardest theoretically allowed slope for synchrotron emission. Similarly, the peak energy E_p was sampled in log-space ($\log(E_p/\text{keV})$) to better handle the orders-of-magnitude dynamic range typical in GRB spectra (1 keV to 10 MeV) while maintaining positive definiteness. These transformations automatically satisfy the physical boundaries of the parameters while improving the sampler's performance in high-curvature regions of parameter space.

The high-energy spectral index β was similarly constrained with $\beta < -2.0$ through a sigmoidal transformation. These choices reflect the known physical limits of non-thermal emission mechanisms while remaining uninformative within the allowed ranges, striking a balance between physical realism and minimal prior influence on the posterior distributions.

6.3 Results

We fitted the light curve data for GRB170817 with a model called the smooth broken power-law model, which is parametrized as follows:

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

where the physical parameters were constrained through Bayesian inference:

- **Frequency Dependence**:
  - ν: Observing frequency (3 GHz for VLA data)
  - β: Spectral index (constrained by broadband SED when available)
- **Temporal Evolution**:
  - t: Time since merger event
  - t_p: Peak time of light curve
  - α_1: Rising slope (> 0 prior for physically plausible models)
  - α_2: Decaying slope (< 0 prior)
  - s: Smoothness parameter (log-uniform prior between 0.1-10)
- **Flux Normalization**:
  - F_p: Peak flux density at 3 GHz (log-uniform prior)

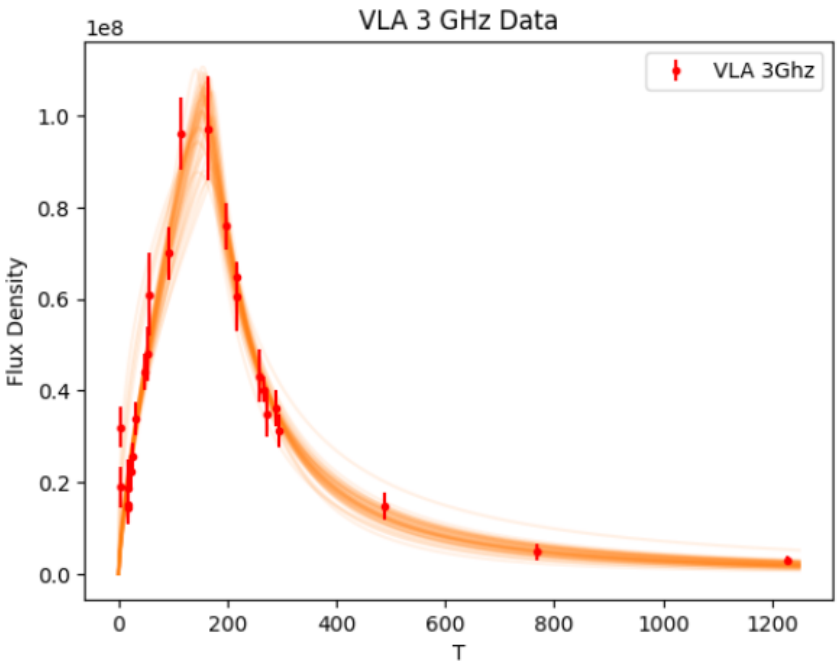

Figure 15: 3 GHz radio light curve of GRB 170817A observed with the Karl G. Jansky Very Large Array (VLA). The flux density measurements (red points) show the temporal evolution of the afterglow emission following the neutron star merger event GW170817. The solid curve represents the posterior predictive distribution from the MCMC analysis performed using the `emcee` package

The corner plot shows all the one and two dimensional projections of the posterior probability distributions of the parameters: `beta`, `alpha1`, `alpha2`, `F_peak`, `t_peak`, `log_s`, `log_f`. This is useful because it quickly demonstrates all of the covariances between parameters. Also, the way that we find the marginalized distribution for a parameter or set of parameters using the results of the MCMC chain is to project the samples into that plane and then make an N-dimensional histogram. That means that the corner plot shows the marginalized distribution for each parameter independently in the histograms along the diagonal and then the marginalized two dimensional distributions in the other panels.

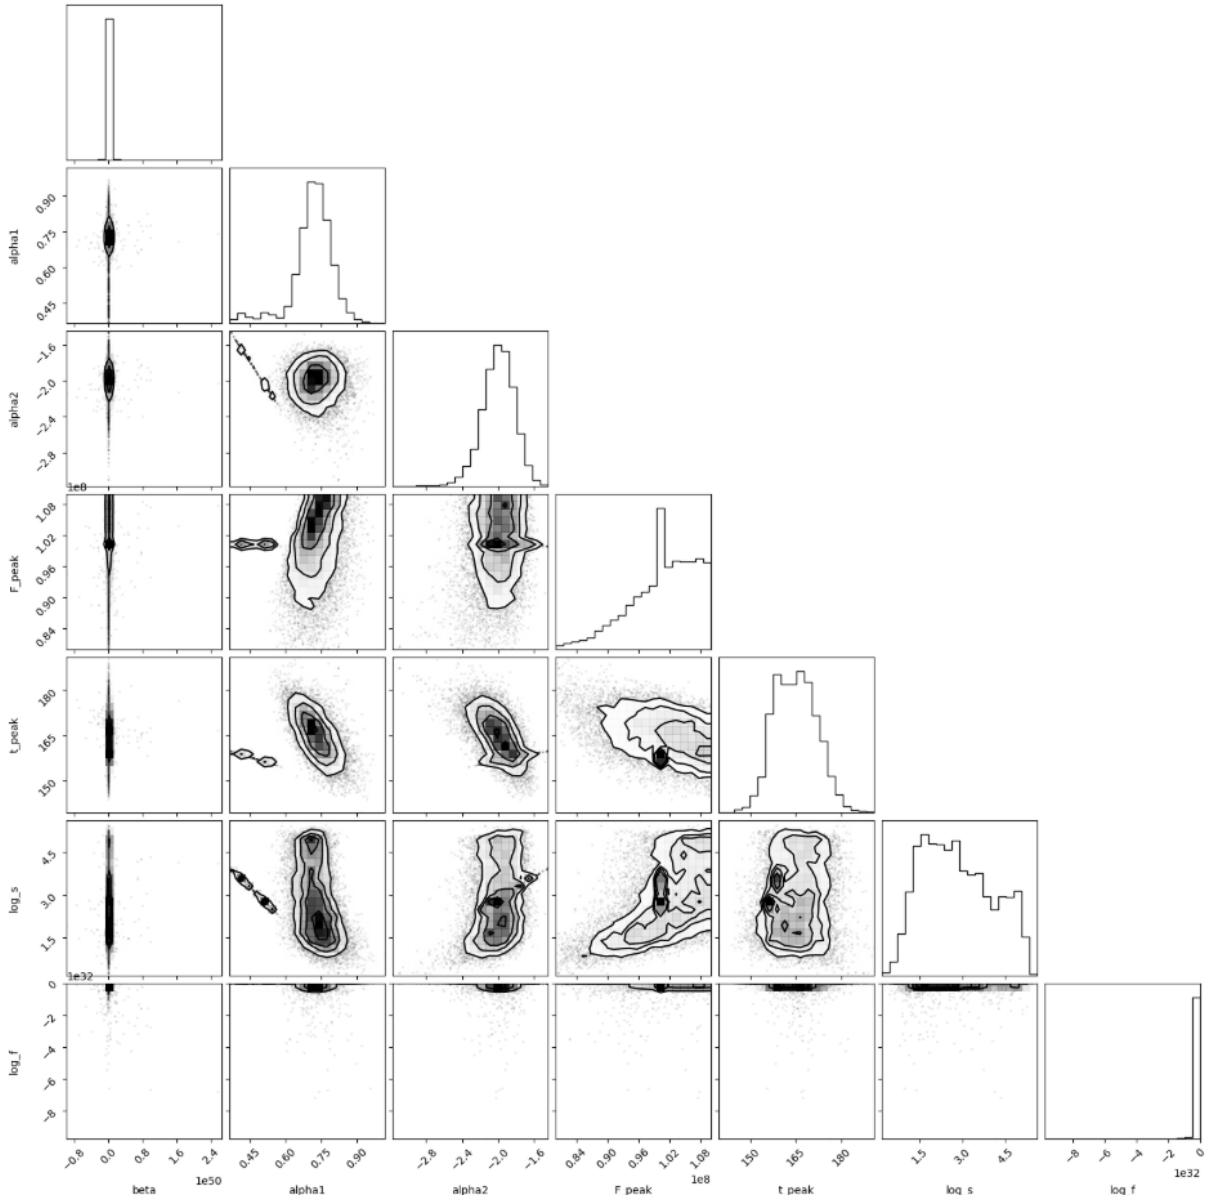

Figure 16: Corner plot showing marginalized posterior distributions for GRB170817.

7 Multi–Mission Analysis with `ThreeML`

`ThreeML`[1] (The Multi-Mission Maximum Likelihood framework) is an open-source Python package designed to perform joint, likelihood-based analyses of data coming from multiple high-energy instruments—ranging from radio to very-high-energy γ-rays—within a single, coherent interface. Internally, `ThreeML` separates the problem into three layers:

(a) **Instrument plug-ins** wrap the official analysis software (or native `FITS`/TTE readers) for each telescope, so that users can ingest *Swift* - BAT/XRT, *Fermi* - GBM/LAT, *INTEGRAL*, *HAWC*, and many other data sets without re-formatting. Each plug-in returns a standard `DataList` object that is agnostic of the underlying mission-specific quirks.

[1] `https://threeml.readthedocs.io`

(b) **Model definition** is provided by the companion library `astromodels`. Spectral and temporal models (e.g. Band, cutoff power law, smoothly-broken power law) are combined into point or extended-source objects and can be assigned informative or uninformative priors for Bayesian work.

(c) **Optimisation and sampling engines** can be swapped on the fly. For maximum-likelihood (frequentist) fits, one can use `iminuit` or `nlopt`. For Bayesian inference, `ThreeML` offers direct hooks to `emcee`, `ultranest`, and `MultiNest`, allowing seamless exploration of highly correlated, multi-modal posteriors typical of GRB spectra.

A minimal GRB spectral fit therefore requires only a few lines:

```
from threeML import *
# 1. Load data ------------------------------------------------------------
gbm = FermiGBMTTELike("TTE", filename="grb230123456_tte.fits")
data = DataList(gbm)

# 2. Define Band model ----------------------------------------------------
spec = Band()
source = PointSource("grb", ra, dec, spectral_shape=spec)
model  = Model(source)

# 3. Fit or sample --------------------------------------------------------
jl = JointLikelihood(model, data)
jl.fit()                            # frequentist
# - or -
bayes = BayesianAnalysis(model, data, sampler="emcee")
bayes.sample(n_iterations=3_000)
```

Key advantages from using this framework include:

- **True multi-mission joint fits**: a single likelihood built from simultaneous BAT, GBM, and LAT data sharpens constraints on $E_{\rm peak}$, α, and β and captures cross-instrument systematics consistently.
- **Transparent engine switching**: the same model can be optimised with `iminuit` for quick point estimates, then re-sampled with `emcee` to obtain full posteriors, without rewriting any code.
- **Reproducibility**: configuration files (`.yaml` or `.json`) store every model parameter, prior, sampler setting, and random seed, making the entire analysis pipeline shareable and version-controlled.

Because `ThreeML` is built entirely in Python, it integrates smoothly with the rest of the GRB analysis workflow—plotting with `matplotlib`, posterior diagnostics with `corner.py`, and high-level orchestration via Jupyter notebooks or `Snakemake` pipelines—providing a robust, end-to-end environment for modern high-energy transient science.

8 References

1. Andrew J. Levan. *Gamma-ray Bursts*. Radboud University.

2. Alessandro Armando Vigliano and Francesco Longo. *Gamma-ray Bursts: 50 Years and Counting!*. University of Udine and Trieste.

3. KC Chan and CT Lenard. *An Introduction to Markov Chains*. La Trobe University.

4. David W Hogg and Daniel Foreman-Mackey. *Data Analysis Recipes: Using Markov Chain Monte Carlo*.

5. Giacomo Vianello and Robert J. Lauer. *The Multi-Mission Maximum Likelihood framework (3ML)*. Stanford and University of New Mexico

6. Emcee Documentation: https://emcee.readthedocs.io/

7. 3ML Documentation: https://threeml.readthedocs.io/