

An internal shock model calibrated with real gamma-ray burst light curves using a genetic algorithm

M. Maistrello^{1,2,*}, C. Guidorzi^{1,2,3}, S. Kobayashi⁴, and R. Maccary^{1,2}

¹ Department of Physics and Earth Science, University of Ferrara, Via Saragat 1, I-44122 Ferrara, Italy

² INAF – Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Piero Gobetti 101, I-40129 Bologna, Italy

³ INFN – Sezione di Ferrara, Via Saragat 1, I-44122 Ferrara, Italy

⁴ Astrophysics Research Institute, LJMU, IC2, Liverpool Science Park, 146 Brownlow Hill, Liverpool L3 5RF, UK

Received 19 May 2026 / Accepted 9 July 2026

ABSTRACT

Context. The origin of prompt emission in gamma-ray bursts (GRBs) remains a fundamental open question. The internal shock (IS) model is a leading mechanism proposed to explain the dissipation of kinetic energy from relativistic ejecta into gamma rays. However, the model parameters have yet to be fully optimised to reproduce the diverse morphological properties observed in real GRB light curves (LCs).

Aims. Utilising a machine-learning framework, we evaluated the IS model through parameter optimisation, comparing the statistical properties of simulated LCs against those of observed data.

Methods. Our dataset consists of three GRB catalogues (*Swift*/BAT, *Fermi*/GBM, and *CGRO*/BATSE). By adopting a model for the GRB formation rate as a function of redshift, we employed a genetic algorithm to optimise the IS model parameters. The algorithm minimises a total loss function based on six independent metrics, representing both the average properties and the statistical distributions of real LCs.

Results. The calibrated IS model successfully reproduces the average post-peak GRB temporal profile, together with the corresponding root-mean-square and third-moment temporal profiles, as well as the average autocorrelation function. Furthermore, it recovers the observed distributions of duration, signal-to-noise ratio, peak count per burst, peak flux, and fluence. We find that a generalised Zipf distribution governs the number of shells emitted per GRB, while rest-frame emission times follow a negative exponential distribution.

Conclusions. The optimised formulation of the IS model reproduces a wide range of observed GRB LC properties, despite its simplified treatment of radiation physics. Moreover, it provides two key insights into the central engine's activity: (i) the emission times suggest a stochastic process where all shells within a given burst have an identical, independent, and constant probability of ejection per unit time; and (ii) the heavy-tailed distribution of the number of shells per GRB mirrors the frequency-magnitude distribution of earthquakes, known as the Gutenberg-Richter law. Finally, this optimised model serves as a predictive tool for the GRB populations expected to be detected by future missions.

Key words. radiation; dynamics – relativistic processes – methods: data analysis – methods: statistical – gamma-ray burst: general

1. Introduction

Gamma-ray bursts (GRBs) are intense flashes of gamma rays, ranging in duration from milliseconds to thousands of seconds, originating at cosmological distances. They are linked to catastrophic stellar events, including the collapse of massive, hydrogen-stripped stars – known as ‘collapsars’ (Woosley 1993; Paczyński 1998; MacFadyen & Woosley 1999) – and the merger of compact binaries (Paczynski 1986; Eichler et al. 1989; Abbott et al. 2017). In both scenarios, the progenitor leaves behind a central engine – either a hyper-accreting stellar-mass black hole (BH; Popham et al. 1999) or a rapidly rotating, highly magnetised neutron star (NS; Usov 1992) –, which launches an ultra-relativistic jet. While GRB emission is believed to arise from internal dissipation within this outflow, several fundamental questions remain. The nature of the central engine, jet-launching mechanism, jet composition, and specific processes that convert internal energy into gamma rays remain under debate. This uncertainty motivates the development of physically

grounded models capable of reproducing the diverse temporal variability observed in GRB light curves.

One of the earliest models proposed to explain GRB emission – commonly referred to as prompt emission – is the so-called ‘internal shock’ (IS) model (Rees & Meszaros 1994). Within this framework, an initially hot fireball composed of photons, electron-positron pairs, and a small amount of baryons converts most of its thermal energy into kinetic energy, which is subsequently partially dissipated through mildly relativistic shocks. These shocks are thought to develop within a highly variable relativistic outflow driven by an intermittent central engine. Electrons are accelerated at the shock fronts, and the resulting internal energy is radiated via synchrotron and inverse Compton processes, powering the observed prompt emission. The identification of a thermal component superimposed on the non-thermal spectrum in a few GRBs has spurred the development of photospheric models. In these models, a fraction of the prompt emission energy originates at the photosphere, where the jet becomes optically thin, resulting in emission at significantly smaller radii than predicted by the standard IS model (see Pe’er 2015; Beloborodov & Mészáros 2017 for reviews).

* Corresponding author: mstmm1@unife.it

In contrast to alternative scenarios, such as the ICMART (internal-collision-induced magnetic reconnection and turbulence; Zhang & Yan 2011) model, the magnetic field in the IS framework is not assumed to play a dynamically dominant role in the ejecta. It thereby provides a physically motivated framework for linking the variability of the central engine to the observed temporal structure of GRB light curves (LCs). Previous studies have indeed shown a strong correlation between the occurrence time of a pulse in a GRB LC and the ejection time of the corresponding shell from the central engine. The number of pulses in a GRB LC reflects the number of shells ejected by the engine, while the separation between pulses traces the time intervals during which the engine was inactive. Thus, the observed temporal structure directly encodes the activity of the central engine (Kobayashi et al. 1997, hereafter K97).

In our previous studies (Bazzanini et al. 2024; Maistrello et al. 2025; hereafter M25), we presented a well-calibrated workflow to optimise the parameters of a stochastic pulse-avalanche model for GRB LCs (Stern & Svensson 1996), starting from real data. The workflow employs a machine-learning technique known as the genetic algorithm (GA), enabling the generation of synthetic LCs that successfully reproduce a set of average temporal properties of real GRB LCs observed by the Burst Alert Telescope (BAT; Barthelmy et al. 2005) aboard the Neil Gehrels Swift Observatory, the Gamma-ray Burst Monitor (GBM; Meegan et al. 2009) aboard *Fermi*, and the Burst and Transient Source Experiment (BATSE; Paciesas et al. 1999) aboard the Compton Gamma Ray Observatory (CGRO; 1991–2000). However, the stochastic pulse-avalanche model lacks a direct physical connection to the underlying dissipation mechanism. A key open question is therefore whether a physically motivated model, such as the IS scenario, can reproduce the observed temporal properties when directly calibrated against real data.

To our knowledge, this is the first systematic optimisation of an IS model directly constrained by multiple statistical properties of observed GRB LCs. To this end, we implemented a version of K97’s IS model, optimised following the prescription described in M25. The paper is organised as follows. In Sect. 2, we describe the sample selection and the metrics adopted for the model optimisation. The model is presented in Sect. 3. Our results, discussion, and conclusions are reported in Sects. 4, 5, and 6, respectively.

2. Data analysis

2.1. Sample selection

We considered three complementary GRB datasets, drawn from several instruments and energy passbands. These include (i) 2024 GRBs from the 4B BATSE catalogue (Paciesas et al. 1999); (ii) 1389 GRBs detected by *Swift*/BAT between January 2005 and November 2023; and (iii) 2356 GRBs from the fourth *Fermi*/GBM catalogue (von Kienlin et al. 2020). Hereafter, these datasets are referred to as the BATSE, *Swift*, and *Fermi* samples, respectively.

For all samples, we analysed background-subtracted LCs in the total instrumental passbands, namely 25–2000 keV for BATSE, 15–150 keV for *Swift*/BAT, and 8–1000 keV for *Fermi*/GBM, adopting a uniform time resolution of 64 ms. For the *Fermi* sample, we used the full NaI scintillator passband and performed background subtraction with the publicly

available GBM Tools¹, following the procedure described in Maccary et al. (2024b).

Only long-duration GRBs were retained by imposing a threshold of $T_{90} > 2$ s. Additional cuts were applied on the signal-to-noise ratio (S/N), requiring $S/N \geq 15$, ≥ 10 , and ≥ 15 for the BATSE, *Swift*, and *Fermi* samples, respectively. Consequently, the samples (i)–(iii) were reduced to 1389, 635, and 1120 GRBs, respectively. The adopted thresholds were chosen to preserve the diversity of LC morphologies while ensuring adequate statistical quality.

Showing compelling evidence for a compact-object merger origin, events – GRB 060614 (Della Valle et al. 2006; Jin et al. 2015); GRB 211211A (Rastinejad et al. 2022; Troja et al. 2022; Yang et al. 2022); GRB 191019A (Levan et al. 2023; Stratta et al. 2025) – were excluded from each dataset. This choice preserves the homogeneity of the sample by excluding events with a different physical origin from the bulk of the population, although the number of long-duration GRBs associated with compact-object mergers may be larger than what common wisdom assumed thus far (Maccary et al. 2026).

2.2. Statistical metrics

To quantify the agreement between simulated and real LCs, we adopted the set of metrics introduced in M25. These metrics capture complementary aspects of GRB LCs. Briefly summarised below (see M25 and references therein for a detailed discussion), these consist of

- the average post-peak time profile, $\langle F/F_p \rangle$, obtained by aligning LCs at their brightest peak and averaging over the first 150 s, where F_p denotes the peak count rate;
- the average third moment of the post-peak profile, $\langle (F/F_p)^3 \rangle$, sensitive to profile asymmetries;
- the average autocorrelation function (ACF) of the LCs;
- the distribution of LC durations expressed in terms of $T_{20\%}$, defined as the time interval during which the signal exceeds 20% of its maximum value (Stern & Svensson 1996);
- the distribution of the S/N of the LCs.

Additionally, we included the distribution of the number of peaks per GRB, which show evidence for the existence of two regimes in long GRBs – namely peak-poor and peak-rich behaviours – consistently in all three datasets (Guidorzi et al. 2024; Maccary et al. 2024b). It complements the original set by constraining the temporal complexity of the simulated GRBs.

3. Internal shock model

The LCs were generated using an IS model originally introduced by K97. In the following, we outline the basic elements of the model that were relevant for our simulations.

Internal shocks develop within a relativistic outflow made up of multiple shells ejected with different Lorentz factors. Energy dissipation occurs when a faster shell catches up with a slower one, leading to an inelastic collision and the formation of a merged shell. Consistently with the notation adopted by K97, physical quantities associated with the rapid shell, the slow shell, and the merged shell are denoted by the subscripts r , s , and m , respectively.

¹ https://fermi.gsfc.nasa.gov/ssc/data/analysis/gbm/gbm_data_tools/gdt-docs/

Under the assumption of a completely inelastic collision, the Lorentz factor of the merged shell is given by

$$\gamma_m \simeq \sqrt{\frac{m_r \gamma_r + m_s \gamma_s}{m_r / \gamma_r + m_s / \gamma_s}}, \quad (1)$$

where m_r and m_s are the shell masses and $\gamma_r \gg 1$ and $\gamma_s \gg 1$ are the corresponding Lorentz factors. The internal energy generated during the collision was obtained from the difference between the kinetic energy of the system before and after the interaction:

$$E_{\text{int}} = m_r c^2 (\gamma_r - \gamma_m) + m_s c^2 (\gamma_s - \gamma_m). \quad (2)$$

Assuming that this internal energy is promptly radiated isotropically in the co-moving frame of the merged shell, the efficiency of conversion from kinetic to internal energy can be expressed as

$$\epsilon = 1 - \frac{(m_r + m_s) \gamma_m}{m_r \gamma_r + m_s \gamma_s}. \quad (3)$$

The observed emission from a single shell collision appears as a pulse whose temporal width δT is determined by several characteristic timescales, namely the radiative cooling time, the hydrodynamic time, and the angular spreading time. Since the internal energy is released via synchrotron and inverse Compton processes and electron cooling occurs on timescales much shorter than the hydrodynamic one, radiative cooling can be neglected in this context. The hydrodynamic timescale corresponds to the time required for the shocks generated during the collision to cross the shells. The interaction produces both a forward and a reverse shock, and the emission duration in the co-moving frame is approximately the time needed for the reverse shock to cross the rapid shell:

$$\delta t_e = \frac{l_r}{c(\beta_r - \beta_{rs})}, \quad (4)$$

where l_r is the width of the rapid shell, while β_r and β_{rs} denote the velocities of the rapid shell and of the reverse shock, respectively. Due to relativistic motion of the emitting region towards the observer, the corresponding emission timescale in the observer frame is reduced by a factor $1/2\gamma_m^2$.

If the collision is assumed to occur at $t = 0$ at a sufficiently large radius R , angular spreading effects play a role in shaping the observed pulse. In particular, when the initial separation L between the shells exceeds l_r , angular effects broaden the pulse and produce an asymmetric temporal profile characterised by a fast rise and a slower decay. In this case, the observed luminosity can be written as

$$\mathcal{L} = \begin{cases} 0 & (t \leq 0), \\ h \left[1 - \frac{1}{(1 + 2\gamma_m^2 ct/R)^2} \right] & (0 < t \leq \frac{\delta t_e}{2\gamma_m^2}), \\ h \left\{ \frac{1}{[1 + (2\gamma_m^2 t - \delta t_e)c/R]^2} - \frac{1}{(1 + 2\gamma_m^2 ct/R)^2} \right\} & (t > \frac{\delta t_e}{2\gamma_m^2}), \end{cases} \quad (5)$$

where $h = E_{\text{int}} 2\gamma_m^2 / \delta t_e$. Following the shock crossing, the two initial shells merge into a thinner shell of width

$$l_m = l_s \frac{\beta_{fs} - \beta_m}{\beta_{fs} - \beta_s} + l_r \frac{\beta_m - \beta_{rs}}{\beta_r - \beta_{rs}}, \quad (6)$$

with β_s and β_{fs} representing the velocities of the slower shell and of the forward shock, respectively. Both β_{fs} and β_{rs} are calculated following Sari & Piran (1995) and K97.

In a realistic outflow, a large number of shell collisions occur. Each interaction produces a shock that converts a fraction of the kinetic energy into thermal energy, which is subsequently radiated through synchrotron and inverse Compton emission, giving rise to a pulse described by Eq. (5). The observed LC is then obtained as the superposition of the pulses produced by all individual collisions.

The outflow is modelled as a wind of n shells, where n is sampled from a generalised Zipf distribution $p(n)^2$ (Mandelbrot 1953):

$$p(n) = \frac{K}{(n + S)^\alpha}, \quad (2 \leq n \leq n_{\text{max}}), \quad (7)$$

where K is a normalisation constant, n_{max} is the maximum value (in our model, it was set to a few hundred), and two free parameters, S (called ‘shift’) and α . This law is found to describe many ranking distributions in the context of complex systems (De Marzo et al. 2021). In physics, for instance, it describes the distribution of the size of earthquakes of a given geographical region as a function of their ranking, which is sorted according to the magnitude, the so-called Gutenberg-Richter law (Gutenberg & Richter 1944). The choice of Eq. (7) followed a number of alternative and less successful attempts. In particular, a heavy-tailed distribution turned out to be a key requirement.

For a given GRB, once n is drawn from Eq. (7), the corresponding emission times are drawn from a common exponential probability density function (PDF):

$$f(t) = \frac{1}{\tau} \exp\left(-\frac{t}{\tau}\right), \quad (8)$$

where τ is different for each GRB and results from sampling a log-normal distribution: $\log \tau \sim \mathcal{N}(\mu_\tau, \sigma_\tau^2)$, where μ_τ and σ_τ are treated as free parameters by the model. Equation (8) implies a temporally declining shell emission rate, which differs from the constant one assumed by K97. This is one of the two major changes adopted in the present model with respect to the original formulation of K97. The rationale behind Eq. (8) follows the simple idea that each shell has the same independent and constant probability per unit time of being emitted. Namely, the resulting distribution is precisely Eq. (8), similarly to the nuclear decay of a given number of atoms of a radioactive element. This simple process was first proposed by Guidorzi et al. (2015) to explain the observed waiting time distribution of gamma-ray peaks and GRB early X-ray flares. The property of a declining shell emission rate agrees with the picture of a central engine that behaves similar to an initially out-of-equilibrium system and progressively evolves towards a more stable configuration.

Each shell is characterised by three quantities: the Lorentz factor γ_i , mass m_i , and width l_i , with $i = 1, \dots, n$. As shown by K97, the model’s efficiency in converting kinetic energy into internal energy increases with the ratio $\gamma_{\text{max}}/\gamma_{\text{min}}$. The second important change with respect to the K97 model concerns the way the Lorentz factors of the shells for the j th GRB are sampled: we first drew two values, $\gamma_{\text{min}}^{(j)}$ and $\gamma_{\text{max}}^{(j)}$, within the interval $[\gamma_{\text{min}}, \gamma_{\text{max}}]$ and then assumed that γ_i is distributed uniformly in logarithmic space between $\gamma_{\text{min}}^{(j)}$ and $\gamma_{\text{max}}^{(j)}$. This prescription introduces burst-to-burst variability in the Lorentz factor contrast while preserving a broad overall range set by γ_{min} and γ_{max} .

The shell mass was assumed to depend on the Lorentz factor according to $m_i = p\gamma_i^\eta$, with $\eta = -1$ and p a proportionality constant. This choice corresponds to shells carrying equal energy,

² Also known as Zipf-Mandelbrot distribution.

such that lighter shells are faster. Finally, to reduce the dimensionality of the parameter space, we assumed a common ratio $r = c\Delta t/l$, where Δt is the time interval between two consecutive shells and l is the width of the later shell, effectively linking the shell separation to their width.

Under these assumptions, the free parameters of the model are α , S , μ_τ , σ_τ , $\gamma_{\min}$, $\gamma_{\max}$, $\log p$, and $\log r$. The parameters are summarised in Table C.1, together with their exploration ranges and corresponding priors.

The dynamical evolution of the system was implemented following K97; additional details are provided in Appendix A. The final temporal profile was obtained as the superposition of the pulses produced by the individual two-shell interactions described by Eq. (5). This represents the theoretical LC expressed in terms of bolometric luminosity, as measured by an observer at rest in the GRB co-moving frame.

A redshift value z_{GRB} was then assigned to each simulated event from sampling the GRB rate $R(z)$:

$$R(z) = \frac{R_{\text{grb}}(z)}{1+z} \frac{dV(z)}{dz}, \quad (9)$$

where $R_{\text{grb}}(z)$ is the GRB formation rate taken from Ghirlanda & Salvaterra (2022, hereafter GS22),

$$R_{\text{grb}}(z) = R_{\text{grb}}(0) \frac{(1+z)^{p_{z,1}}}{1 + \left(\frac{1+z}{p_{z,2}}\right)^{p_{z,3}}}. \quad (10)$$

Here, $R_{\text{grb}}(0)$ is the local volumetric rate and $p_{z,1} = 3.33$, $p_{z,2} = 3.42$, and $p_{z,3} = 6.21$. The co-moving volume element is given by

$$\frac{dV(z)}{dz} = \frac{4\pi c d_L^2(z)}{H(z)(1+z)^2}, \quad (11)$$

where $d_L(z)$ is the luminosity distance and $H(z) = H_0[\Omega_M(1+z)^3 + \Omega_\Lambda + (1 - \Omega_M - \Omega_\Lambda)(1+z)^2]^{1/2}$. The $(1+z)$ term in the denominator of Eq. (9) accounts for the cosmic dilation in the observed rate. The redshift was randomly drawn from Eq. (9), once it was properly normalised as a PDF.

The assigned redshift z_{GRB} was used to compute the luminosity distance and convert the bolometric luminosity LC into an observed flux LC within the instrumental energy passband. The rest-frame photon spectrum $N(E)$ was assumed to follow the Band function (Band et al. 1993), characterised by the intrinsic peak energy $E_{p,i}$ and the low- and high-energy spectral indices α and β . The value of $E_{p,i}$ was derived from the $E_{p,i}-L_{\gamma,\text{iso}}$ correlation reported in GS22. The spectral indices were randomly drawn from the corresponding observed distributions (see Appendix B for details). The k -correction factor was then computed as

$$k_{\text{flux}} = \frac{\int_{E_1(1+z_{\text{GRB}})}^{E_2(1+z_{\text{GRB}})} EN(E)dE}{\int_1^{10^4} EN(E)dE}, \quad (12)$$

where E_1 and E_2 define the limits of the instrumental energy passband. This expression holds under the assumption of no spectral evolution.

The flux LC, F , was subsequently converted into counts (integrated over 64 ms) or count rates, R , depending on the instrument, following the methodology described in M25. For completeness, we briefly summarise the procedure here. The conversion was expressed through a factor $k = \log_{10}(F/R)$, where F is the flux and R represents counts for BATSE, count

Table 1. Best-fitting parameters of the IS model for the *Swift*/BAT, *Fermi*/GBM, and *CGRO*/BATSE samples.

Parameter	<i>Swift</i> /BAT (15–150 keV)	<i>Fermi</i> /GBM (8–1000 keV)	<i>CGRO</i> /BATSE (25–2000 keV)
α	$2.06^{+0.01}_{-0.01}$	$2.25^{+0.02}_{-0.05}$	$2.26^{+0.20}_{-0.06}$
S	$1.36^{+0.56}_{-0.17}$	$2.41^{+0.31}_{-0.33}$	$2.45^{+0.48}_{-0.62}$
μ_τ	$3.53^{+0.01}_{-0.11}$	$3.14^{+0.14}_{-0.13}$	$2.76^{+0.10}_{-0.08}$
σ_τ	$0.41^{+0.36}_{-0.28}$	$0.25^{+0.19}_{-0.13}$	$0.30^{+0.31}_{-0.07}$
$\gamma_{\min}$	30^{+6}_{-3}	23^{+7}_{-10}	33^{+7}_{-10}
$\gamma_{\max}$	776^{+179}_{-214}	549^{+263}_{-195}	832^{+59}_{-140}
p (10^{32} g)	$5.2^{+0.1}_{-0.9}$	$5.4^{+0.3}_{-1.2}$	$3.9^{+0.2}_{-0.3}$
r	$3.1^{+1.1}_{-0.6}$	$7.2^{+1.9}_{-1.6}$	$4.6^{+2.5}_{-1.1}$
Loss (Train best)	0.29	0.64	1.29
Loss (Train avg.)	0.92	1.44	2.43
Loss (Test)	0.35	1.21	2.17
Loss (Test: $\langle F/F_p \rangle$)	0.18	0.36	0.30
Loss (Test: $\langle (F/F_p)^3 \rangle$)	0.12	0.13	0.27
Loss (Test: $\langle \text{ACF} \rangle$)	0.60	1.86	0.78
Loss (Test: $T_{20\%}$)	1.20	2.12	1.67
Loss (Test: S/N)	0.00	0.00	0.00
Loss (Test: N_p)	0.00	2.78	10.0

Notes. Parameter values correspond to the medians of the distributions in the final GA generation, with uncertainties given by the 5th and 95th percentiles. ‘Train best’ and ‘Train avg.’ denote the minimum and mean loss values in the final generation, respectively. The test set consists of an independent sample of 5000 simulated LCs. The last rows report the individual contributions of each metric to the total test loss.

rate per fully illuminated detector for an equivalent on-axis source in the case of *Swift*/BAT, and counts per detector for *Fermi*/GBM. The conversion factors were derived using the fluence values reported in the official GRB catalogues: the fourth BATSE catalogue (Paciesas et al. 1999), the third *Swift*/BAT catalogue (Lien et al. 2016), and the fourth *Fermi*/GBM catalogue (von Kienlin et al. 2020). The corresponding counts were obtained by integrating the signal over the 7σ time interval, defined as the interval between the first and last bins in which the signal exceeds the background by at least 7σ .

As a final step, background contributions were added to the simulated LCs. Following M25, constant background rates of $2.9 \text{ counts s}^{-1}$ and $39.4 \text{ counts s}^{-1}$ were adopted for BATSE and GBM, respectively. These values correspond to the median measured error rates of the respective real LCs. The final LCs were generated as Poisson realisations with expectation value equal to the sum of the noise-free LC and the background level. The background was subsequently subtracted. For *Swift*/BAT, each bin was instead sampled from a Gaussian distribution with mean equal to the noise-free LC and standard deviation extracted from the distribution of the real LC uncertainties.

Optimisation of the model parameters was carried out using a GA, following the same general framework adopted in M25. Further details are provided in Appendix C.

4. Results

4.1. Optimal parameter values

The optimised model parameters for the *Swift*, *Fermi*, and BATSE samples – defined as the median values of the final-generation population – are listed in Table 1. To provide a robust estimate of the model performance, the relative loss values were




Fig. 1. Distributions of the six metrics for the real (blue) and simulated (red) LCs of the *Swift*/BAT dataset. The simulated LCs were obtained from the optimised parameter set (Table 1). *Top left:* Average peak-aligned post-peak normalised time profile and corresponding root-mean-square (rms) deviation of the individual profiles, $F_{\text{rms}} \equiv [\langle (F/F_p)^2 \rangle - \langle F/F_p \rangle^2]^{1/2}$. *Top centre:* Average peak-aligned third-moment profile. *Top right:* ACF. *Bottom left:* Distribution of the $T_{20\%}$ duration. *Bottom centre:* S/N distribution. *Bottom right:* Distribution of the number of peaks per GRB. As in M25, the top-row profiles were smoothed using a Savitzky–Golay filter to reduce the effects of Poisson noise.

recomputed using an independent set of 5000 simulated LCs for each dataset. Figure 1 compares the distributions of the six metrics for real and simulated LCs of the *Swift*/BAT sample. Examples of LCs are illustrated in Fig. D.1.

Overall, the optimised solutions work perfectly for all three datasets, as indicated by the test loss. Only for BATSE, the model provides a poor result on the sixth metric (distribution of the number of peaks per GRB), while the remaining five metrics are satisfactorily fulfilled. All parameters have very similar values within uncertainties across the different datasets. Notably, μ_τ slightly decreases passing from the softest to the hardest energy passband. This is possibly driven by the fact that real LCs tend to last longer in the softer energy channels than in the harder ones.

4.2. Properties of the simulated GRB populations

We derived the distributions of several fundamental quantities from the optimised IS models and compared them with observations. As far as such intrinsic properties as $L_{\gamma,\text{iso}}$ and $E_{\gamma,\text{iso}}$ are concerned, the real distributions are based on the overall sample of GRBs with measured redshift, which is primarily given by *Swift*/BAT, and then *Fermi*/GBM, *Konus*/WIND and *BeppoSAX*. Since the observed sample includes only GRBs with measured redshift, it is subject to selection effects and is not a complete flux-limited sample. The comparisons shown below should therefore be regarded as consistency checks rather than formal fits. When available, $L_{\gamma,\text{iso}}$ and $E_{\gamma,\text{iso}}$ values were preferentially taken from the *Konus*/WIND catalogues (Tsvetkova et al. 2017, 2021) and *Fermi*/GBM (von Kienlin et al. 2020) because of their

broad energy band. The comparison between the expected and observed distributions contributes to further test the ability of the optimised IS model to account for the observed properties and, at the same time, to gain clues on the way IS-based inner engines work.

To enable a direct comparison that includes faint GRBs, we relaxed the S/N threshold and generated a new set of 5000 LCs for each dataset. We then simulated a detection process to determine whether a synthetic GRB would have been detected. Specifically, we applied a well-calibrated peak search algorithm, FAST-MEPSA (Maistrello et al. 2026): when at least one peak with $S/N \geq 5$ was identified, the GRB was considered detected. FAST-MEPSA is an optimised version of the MEPSA algorithm, originally developed to detect peaks in uniformly sampled time series affected by uncorrelated Gaussian noise (Guidorzi 2015). It has proven effective for large-scale analyses, particularly when the detection of faint events is required. The detection process implemented with FAST-MEPSA emulates a rate trigger logic. As a consequence, it is not sensitive to very long and faint events that could instead be detected through image-trigger techniques, such as those employed by *Swift*/BAT or any other imaging instrument. Nevertheless, our goal was to assess the capability of the model to reproduce the properties of the observed population of GRBs. Developing a detailed image-trigger simulator lies beyond the scope of this work.

Figures 2 and 3 show the comparisons between simulated and observed distributions of $L_{\gamma,\text{iso}}$ and of $E_{\gamma,\text{iso}}$. All $L_{\gamma,\text{iso}}$ distributions peak in the range 10^{52} – 10^{53} erg s $^{-1}$, whereas the $E_{\gamma,\text{iso}}$ distributions are centred around 10^{53} erg. This overall




Fig. 2. Bolometric, isotropic-equivalent peak luminosity distributions of the GRB populations simulated for BATSE (solid red line), *Swift*/BAT (shaded blue area), and *Fermi*/GBM (shaded orange area). The corresponding distribution of real GRBs with known redshift (solid black line) is mostly due to *Swift*, then *Fermi*, Konus/*WIND*, and *BeppoSAX*. The simulated distributions are rescaled to match the total number of observed events.




Fig. 3. Bolometric, isotropic-equivalent gamma-ray released energy distributions of the GRB populations simulated for BATSE (solid red line), *Swift*/BAT (shaded blue area), and *Fermi*/GBM (shaded orange area). The corresponding distribution of real GRBs with known redshift (solid black line) is mostly due to *Swift*, then *Fermi*, Konus/*WIND*, and *BeppoSAX*. The simulated distributions are rescaled to match the total number of observed events.

agreement indicates that the model successfully reproduces the characteristic luminosity and energy ranges observed in real GRBs. Concerning $E_{\gamma,iso}$, the real distribution tends to have more low-energy GRBs ($E_{\gamma,iso} \lesssim 10^{51}$ erg s⁻¹). The redshift distributions are shown in Fig. 4. The simulated samples display a broad maximum around $z \sim 2-3$, consistent with the observed distribution of GRBs with known z . The relative scarcity of events at $z \lesssim 1$ in the simulations is attributable to the adopted detection criteria. In particular, the observed low- z population includes low-luminosity GRBs (*ll*-GRBs), which preferentially triggered the imaging criteria of *Swift*/BAT.

Figures 5 and 6 display the predicted cumulative fluence and peak photon flux distributions of the *Swift*/BAT sample compared with the observed ones³ in the corresponding 15–150 keV passband. For reference, the solid black line indicates the canonical $-3/2$ slope expected for the $\log N - \log S$ distribution in a Euclidean universe. The simulated distributions closely follow the observed behaviour especially in the bright tail. Overall, these comparisons demonstrate that the model produces simu-




Fig. 4. Simulated redshift distributions inferred from the IS model optimised with BATSE (red line), *Swift* (shaded blue area), and *Fermi* samples (shaded orange area). The distribution of real GRBs with known z (solid black line) is predominantly due to *Swift*, then *Fermi*/GBM, Konus/*WIND*, and *BeppoSAX*. The simulated distributions are rescaled to match the total number of observed events.




Fig. 5. Simulated (shaded) and real (solid line) 15–150 keV fluence normalised survival distributions of the GRB population seen with *Swift*/BAT. The dashed line indicates the canonical $-3/2$ slope of the $\log N - \log S$ GRB distribution for a Euclidean universe.

lated populations consistent with the main observational properties of long GRBs.

4.3. Luminosity function

Figure 7 illustrates the luminosity function $dR_{\text{grb}}/dL_{\gamma,iso}$, i.e. the GRB volumetric rate as a function of $L_{\gamma,iso}$, evaluated at $z = 0$ assuming no redshift evolution for the *Swift*/BAT and *Fermi*/GBM samples simulated with the corresponding optimised IS models. Concerning the *Swift* sample, we count 505 GRBs detected by BAT with a 15–150-keV peak photon flux ≥ 2.6 ph cm⁻² s⁻¹⁴ over a time span of 21.1 yr. Given an average duty cycle of 0.75 and a field of view of 1.4 sr, the full-sky rate is 286 GRBs yr⁻¹. The fraction of simulated *Swift*/BAT GRBs with peak photon flux in excess of 2.6 ph cm⁻² s⁻¹ and $L_{\gamma,iso} > 10^{47}$ erg s⁻¹ amounts to $\sim 16\%$ of the total. Consequently, our IS model, optimised on BAT data, predicts ~ 1770 GRBs yr⁻¹ with $L_{\gamma,iso} > 10^{47}$ erg s⁻¹. By demanding that the integral of Eq. (9) matches this predicted rate, we determine $R_{\text{grb}}(0) = 0.13$ Gpc⁻³ yr⁻¹, in agreement with previous estimates (Salvaterra et al. 2012) and somewhat lower than others

³ https://swift.gsfc.nasa.gov/archive/grb_table/

⁴ This threshold is well above the BAT sensitivity of ~ 0.3 ph cm⁻² s⁻¹ (Lien et al. 2016).




Fig. 6. Simulated (shaded) and real (solid line) 15–150 keV peak-photon-flux-normalised survival distributions of the GRB population seen with *Swift*/BAT. The peak photon flux refers to LCs with a bin time of 1 s.




Fig. 7. Local isotropic-equivalent luminosity function obtained from simulated populations based on the optimised IS model calibrated on *Swift* and *Fermi* data. Also shown are the analogous quantities derived by Wanderman & Piran (2010) (solid black line) and GS22, the latter down-scaled by a factor of $f_b^{-1} = 100$ (dashed purple line), as the original value was corrected for collimation.

(Howell et al. 2014; Sun et al. 2015; Palmerio & Daigne 2021). For the *Fermi* sample, adopting a peak photon flux threshold of $5 \text{ ph cm}^{-2} \text{ s}^{-1}$ in the 10–1000 keV band, together with an average duty cycle of 0.55 and a field of view of 8.8 sr, our model predicts a local rate of $0.15 \text{ Gpc}^{-3} \text{ yr}^{-1}$. The un-normalised luminosity function, given by the differential distribution of the simulated $L_{\gamma,\text{iso}}$ values, was then fitted with a broken power law:

$$\frac{dN}{dL_{\gamma,\text{iso}}} \propto \begin{cases} L_{\gamma,\text{iso}}^{-a_1} & L_{\gamma,\text{iso}} \leq L_{\gamma,\text{iso},b}, \\ L_{\gamma,\text{iso},b}^{a_2-a_1} L_{\gamma,\text{iso}}^{-a_2} & L_{\gamma,\text{iso}} > L_{\gamma,\text{iso},b}, \end{cases} \quad (13)$$

where $L_{\gamma,\text{iso},b}$ is the break luminosity and a_1 and a_2 the power law indices at low- and high-luminosities, respectively. The best-fitting parameter values are reported in Table 2. We obtained the local luminosity function by normalising the simulated distribution and the corresponding best-fit broken power law by requiring that the integral matches $R_{\text{grb}}(0)$.

In Fig. 7, we compare our luminosity function with two other examples: Wanderman & Piran (2010) considered $L_{\gamma,\text{iso}}$ and found $R_{\text{grb}}(0) = 1.3^{+0.6}_{-0.7} \text{ Gpc}^{-3} \text{ yr}^{-1}$, whereas GS22 considered L_{γ} , which was corrected for collimation. In the latter case, we divided it by a factor of $f_b^{-1} = 100$, where $f_b = (1 - \cos \theta_j)$ is the collimation-correction factor, assuming a typical θ_j of a few degrees. Our luminosity function is broadly consistent with

Table 2. Best-fitting parameters of the broken power-law model (Eq. (13)) describing the luminosity function of the simulated $L_{\gamma,\text{iso}}$ for the *Swift*/BAT and *Fermi*/GBM samples.

Parameter	<i>Swift</i> /BAT	<i>Fermi</i> /GBM
$\log L_{\gamma,\text{iso},b}$	$51.75^{+0.07}_{-0.05}$	$52.14^{+0.05}_{-0.06}$
a_1	$0.65^{+0.02}_{-0.02}$	$0.70^{+0.01}_{-0.01}$
a_2	$1.46^{+0.03}_{-0.02}$	$1.55^{+0.03}_{-0.03}$

Notes. Quoted uncertainties correspond to 90% credible intervals derived from the posterior distributions obtained through Markov chain Monte Carlo (MCMC) simulations.

these estimates, exhibiting a similar break luminosity, while the inferred slopes lie between the two reference models.

5. Discussion

Several observed properties originally motivated the formulation of the IS model, most notably the lack of temporal pulse broadening within multi-peaked bursts (Fenimore et al. 1999; Ramirez-Ruiz & Fenimore 2000). Early IS-model simulations successfully reproduced the observed average power density spectral slope (Beloborodov et al. 1998, 2000; Spada et al. 2000) and qualitatively captured the diversity of LCs (K97; Daigne & Mochkovitch 1998; Ramirez-Ruiz & Fenimore 2000). However, because exploring the model complex, multi-dimensional parameter space was computationally prohibitive at the time, these early studies relied largely on fiducial values. Today, the availability of extensive catalogues containing thousands of GRBs, hundreds of which have confirmed redshifts, combined with modern high-performance computing and sophisticated machine learning, allows for a rigorous quantitative optimisation of the IS model. By requiring the simultaneous reproduction of multiple observed properties, this work represents the first comprehensive effort to constrain the IS model parameters in this manner.

Despite the relative simplicity of the K97 model, which assumes uniform kinetic energy across all shells and total conversion of dissipated energy into gamma rays, it successfully reproduces the diverse properties and distributions characteristic of observed LCs. Specifically, the predicted fluence and peak flux distributions are consistent with *Swift* observations (Figs. 5 and 6). One notable exception is the poorer agreement in the BATSE peak-number distribution. This may partly reflect differences in instrumental response, energy band, background treatment, or peak-identification efficiency and deserves a dedicated investigation.

Leveraging the GRB formation rate from GS22, we constrained several intrinsic properties, finding that the predicted distributions for both $E_{\gamma,\text{iso}}$ and $L_{\gamma,\text{iso}}$ of the detectable population align with observations. By accounting for the entire predicted population (independent of experimental detection limits), we derived the luminosity function, defined as the volumetric rate as a function of $L_{\gamma,\text{iso}}$ for $L_{\gamma,\text{iso}} > 10^{47} \text{ erg s}^{-1}$. Assuming no evolution with redshift beyond a normalisation that scales with the GRB formation rate, Fig. 7 presents the local luminosity function ($dR_{\text{grb}}/dL_{\gamma,\text{iso}}$ at $z = 0$). This result is well modelled by a broken power law, similar to that of GS22, though it exhibits a shallower slope at low luminosities. This discrepancy implies a lower fraction of *ll*-GRBs within the framework of the IS model, suggesting that a significant portion of *ll*-GRBs may have a distinct origin (Campana et al. 2006;

Liang et al. 2007; Virgili et al. 2009; Nakar & Sari 2012; Nakar 2015; Dong et al. 2023; Irwin & Hotokezaka 2025). Alternatively, if *ll*-GRBs are interpreted as standard GRBs viewed off-axis (Granot et al. 2005; Guidorzi et al. 2009; Pescalli et al. 2015; Salafia et al. 2016), our model’s inability to account for their abundance remains consistent, as it implicitly assumes on-axis observers. The broken power-law form is not imposed explicitly, but emerges from the combined distributions of the shell properties and collision histories selected by the optimisation. Its detailed physical origin cannot be uniquely identified within the present analysis and would require a further study.

Furthermore, we gained insight into intrinsic distributions that, while less accessible through direct observation, provide a detailed characterisation of the central engine activity: the number of shells per GRB and their rest-frame emission times.

- Number of shells (n): The number of shells per GRB follows a generalised Zipf distribution (Eq. 7). Unlike other discrete distributions such as the Poisson, which failed to fit the data, the Zipf-Mandelbrot distribution is heavy-tailed. It is ubiquitous in physics and in many other fields (socio-economic and ecological systems), as it describes many rank-size distributions (see De Marzo et al. 2021 and references therein). It incorporates a shift term (S in Eq. 7) representing the number of shells below which the probability levels off. In our datasets, S ranges between 1 and 3. The power-law index α^5 of the probability mass function (PMF) of Eq. (7) for all the datasets is approximately 2, which is the same as the Gutenberg-Richter law that describes the frequency–size distribution of tectonic earthquakes of any given geographical location⁶. Just as earthquake frequency scales with magnitude, the frequency of GRBs scales with the number of shells n . This power-law nature may indicate self-organised critical (SOC) behaviour, a concept previously proposed for GRB γ -ray pulses and X-ray flares (Wang & Dai 2013; Lyu et al. 2020; Li et al. 2023; Wei 2023; Maccary et al. 2024a).
- Shell emission times: For any given GRB, the shell emission times follow a negative exponential distribution with an e-folding time τ . This distribution is the hallmark of a simple stochastic process where each shell has a constant, independent probability per unit time of being ejected. The values of τ are sampled from a log-normal distribution (see Table 1 for mean and standard deviation). Such a process has been suggested to explain the power-law nature of collective waiting times in γ -ray pulses and X-ray flares (Guidorzi et al. 2015). If the GRB engine is powered by a hyper-accreting disc, this could imply that the disc is composed of fragments, each with an equal probability per unit time of being accreted and launching a corresponding shell. Theoretically, such disc fragmentation has been identified as a potential driver of central engine activity (Dall’Osso et al. 2017; Shahamat et al. 2021; Lerner et al. 2026).

Despite these successes, the present implementation involves several simplifying assumptions. Most importantly, the model does not explicitly include the microphysics of relativistic shocks. The relatively low radiative efficiency is a well-known challenge for standard IS models (e.g. Kumar 1999). Introduc-

ing microphysical parameters such as ϵ_e , ϵ_B , and the fraction of accelerated electrons, ξ , would allow the radiation spectrum and its characteristic peak energy to be calculated self-consistently. Requiring the spectral peak to lie within the observed range of approximately 0.1–1 MeV could modify the shell properties and other dynamical parameters inferred from the present optimisation, including those reported in Table 1. Although the overall dynamical dissipation efficiency of the present model is typically of the order of 10%, this quantity should be clearly distinguished from the prompt gamma-ray efficiency, whose self-consistent determination requires modelling how the dissipated energy is transferred to radiating particles and converted into gamma rays. Such a treatment is beyond the scope of the present work and is left for future investigation.

The model also neglects the relationship between photon arrival times and observed energies arising from latitude-dependent Doppler boosting. Owing to this simplified treatment, spectral evolution is not inherently captured; the dependence of the simulated LCs on the detector energy passband is mediated only through cosmological redshift and time dilation. Nevertheless, the primary objective of this study was to evaluate the ability of the IS model to reproduce the temporal and morphological diversity of GRB LCs, thereby constraining the intrinsic properties of the central engine.

The current framework can be extended to incorporate both shock microphysics and Doppler-related effects. A further development could involve a more realistic collision prescription, or an alternative simplified description of dissipation in an irregular outflow (e.g. Kobayashi & Sari 2001), in which internal energy that is not radiated is converted back into relative bulk motion and dissipated again in subsequent interactions. Such energy recycling could mitigate the reduction in cumulative radiative efficiency when only a fraction of the shock-generated energy is radiated in each collision. By contrast, the present prescription treats collisions as fully inelastic and assumes that all of the resulting internal energy is radiated promptly.

6. Conclusions

Although the IS model remains one of the most established frameworks for GRB prompt emission, a rigorous optimisation of its parameters to quantitatively reproduce the properties of observed LCs has remained elusive. This work addresses that gap by utilising three independent catalogues comprising several thousand bursts from *Swift*/BAT, *Fermi*/GBM, *CGRO*/BATSE, and optimising the model parameters via a GA. The model successfully accounts for all datasets, yielding parameter values that are broadly consistent across the different catalogues.

Our previous work introduced a stochastic avalanche model (Bazzanini et al. 2024; M25) to reproduce GRB LC distributions through GA optimisation. Based on the runaway pulse mechanism proposed by Stern & Svensson (1996), that model demonstrated that an avalanche process could underpin multi-peaked GRB structures, though it remained agnostic to the specific physics. The current study advances this effort by replacing the stochastic toy model with a physical IS framework. Combined with the redshift-dependent formation rate from GS22, this physical approach allows the GA to constrain the fundamental parameters governing the central engine’s dynamics.

Among its various insights, the model predicts a luminosity function that closely resembles the established broken power law, with a characteristic break near $10^{52} \text{ erg s}^{-1}$. The primary discrepancy lies in the lower predicted volumetric rate of *ll*-GRBs ($L_{\gamma, \text{iso}} < 10^{50} \text{ erg s}^{-1}$). We conclude that the current

⁵ The power-law index α of the probability mass function should not be confused with the index of the rank-size distribution $\gamma = 1/(\alpha - 1) = 1$ for $\alpha = 2$ (De Marzo et al. 2021).

⁶ In the literature, the Gutenberg-Richter law is often described as a survival distribution, $\log N(\geq M) = a - bM$, where M is magnitude and $N(\geq M)$ is the number of earthquakes with magnitudes equal or greater than M , with $b = \alpha - 1$. The b -value for tectonic earthquakes is around 1 (that is $\alpha = 2$) (Roberts et al. 2015).

formulation of the IS model cannot easily account for the inferred abundance of *ll*-GRBs; this suggests a distinct origin for these events, such as shock-breakout radiation or a different geometric configuration (e.g. standard GRBs viewed off-axis).

Future advancements could follow several complementary paths. These include: (i) modelling the energy dependence of the LC by accounting for high-latitude emission, where curvature effects lead to softer spectra due to reduced Doppler boosting; (ii) incorporating the micro-physics of relativistic shocks, assuming synchrotron and inverse Compton radiation, to constrain the micro-physical parameters; (iii) extending this framework to other extensive and upcoming catalogues, including *Insight-HXMT* (Hard X-ray Modulation Telescope; Song et al. 2022), *Konus/WIND* (Tsvetkova et al. 2017, 2021), and *SVOM* (Space-based multi-band astronomical Variable Objects Monitor; Atteia et al. 2022); and (iv) integrating additional metrics that capture the energy-dependent evolution of the LCs.

Finally, the optimised IS model serves as a predictive framework for the populations to be detected by future GRB missions. As a physically grounded tool, the IS model can effectively simulate the LC morphologies of many different GRBs at any redshift. These simulations will be instrumental in testing proposed on-board trigger algorithms. The source code used in this work is publicly available in a GitHub repository⁷.

Acknowledgements. We acknowledge the reviewer for their useful insights and comments. M. M. expresses his gratitude to the Astrophysics Research Institute of the Liverpool John Moores University for their warm hospitality during his three-month stay.

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *ApJ*, 848, L12
- Amati, L. 2006, *MNRAS*, 372, 233
- Atteia, J. L., Cordier, B., & Wei, J. 2022, *Int. J. Mod. Phys. D*, 31, 2230008
- Band, D., Matteson, J., Ford, L., et al. 1993, *ApJ*, 413, 281
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *Space Sci. Rev.*, 120, 143
- Bazzanini, L., Ferro, L., Guidorzi, C., et al. 2024, *A&A*, 689, A266
- Beloborodov, A. M., & Mészáros, P. 2017, *Space Sci. Rev.*, 207, 87
- Beloborodov, A. M., Stern, B. E., & Svensson, R. 1998, *ApJ*, 508, L25
- Beloborodov, A. M., Stern, B. E., & Svensson, R. 2000, *ApJ*, 535, 158
- Campana, S., Mangano, V., Blustin, A. J., et al. 2006, *Nature*, 442, 1008
- Daigne, F., & Mochkovitch, R. 1998, *MNRAS*, 296, 275
- Dall’Osso, S., Perna, R., Tanaka, T. L., & Margutti, R. 2017, *MNRAS*, 464, 4399
- De Marzo, G., Zaccaria, A., & Pietronero, L. 2021, *Phys. Rev. Res.*, 3, 013084
- Della Valle, M., Chincarini, G., Panagia, N., et al. 2006, *Nature*, 444, 1050
- Dong, X. F., Zhang, Z. B., Li, Q. M., Huang, Y. F., & Bian, K. 2023, *ApJ*, 958, 37
- Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nature*, 340, 126
- Fenimore, E. E., Ramirez-Ruiz, E., & Wu, B. 1999, *ApJ*, 518, L73
- Ghirlanda, G., & Salvaterra, R. 2022, *ApJ*, 932, 10
- Granot, J., Ramirez-Ruiz, E., & Perna, R. 2005, *ApJ*, 630, 1003
- Guidorzi, C. 2015, *Astron. Comput.*, 10, 54
- Guidorzi, C., Clemens, C., Kobayashi, S., et al. 2009, *A&A*, 499, 439
- Guidorzi, C., Dichiara, S., Frontera, F., et al. 2015, *ApJ*, 801, 57
- Guidorzi, C., Sartori, M., Maccary, R., et al. 2024, *A&A*, 685, A34
- Gutenberg, B., & Richter, C. F. 1944, *Bulletin of the Seismological Society of America*, 34, 185
- Howell, E. J., Coward, D. M., Stratta, G., Gendre, B., & Zhou, H. 2014, *MNRAS*, 444, 15
- Irwin, C. M., & Hotokezaka, K. 2025, *MNRAS*, 542, 1269
- Jin, Z.-P., Li, X., Cano, Z., et al. 2015, *ApJ*, 811, L22
- Kobayashi, S., & Sari, R. 2001, *ApJ*, 551, 934
- Kobayashi, S., Piran, T., & Sari, R. 1997, *ApJ*, 490, 92
- Kumar, P. 1999, *ApJ*, 523, L113
- Lerner, Y., Stone, N. C., & Ofengeim, D. D. 2026, *MNRAS*, 545, staf1835
- Levan, A. J., Malesani, D. B., Gompertz, B. P., et al. 2023, *Nat. Astron.*, 7, 976
- Li, X.-J., Zhang, W.-L., Yi, S.-X., Yang, Y.-P., & Li, J.-L. 2023, *ApJS*, 265, 56
- Liang, E., Zhang, B., Virgili, F., & Dai, Z. G. 2007, *ApJ*, 662, 1111
- Lien, A., Sakamoto, T., Barthelmy, S. D., et al. 2016, *ApJ*, 829, 7
- Lyu, F., Li, Y.-P., Hou, S.-J., et al. 2020, *Front. Phys.*, 16, 14501
- Maccary, R., Guidorzi, C., Amati, L., et al. 2024a, *ApJ*, 965, 72
- Maccary, R., Maistrello, M., Guidorzi, C., et al. 2024b, *A&A*, 688, L8
- Maccary, R., Guidorzi, C., Amati, L., et al. 2026, *A&A*, 708, L18
- MacFadyen, A. I., & Woosley, S. E. 1999, *ApJ*, 524, 262
- Maistrello, M., Ferro, L., Bazzanini, L., Maccary, R., & Guidorzi, C. 2025, *A&A*, 697, A76
- Maistrello, M., Maccary, R., & Guidorzi, C. 2026, *Astron. Comput.*, 55, 101040
- Mandelbrot, B. B. 1953, *Communication Theory* (London: Butterworths), 486
- Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, 702, 791
- Nakar, E. 2015, *ApJ*, 807, 172
- Nakar, E., & Sari, R. 2012, *ApJ*, 747, 88
- Paciesas, W. S., Meegan, C. A., Pendleton, G. N., et al. 1999, *ApJS*, 122, 465
- Paczynski, B. 1986, *ApJ*, 308, L43
- Paczynski, B. 1998, *ApJ*, 494, L45
- Palmerio, J. T., & Daigne, F. 2021, *A&A*, 649, A166
- Pe’er, A. 2015, *Adv. Astron.*, 2015, 22
- Pescalli, A., Ghirlanda, G., Salafia, O. S., et al. 2015, *MNRAS*, 447, 1911
- Poolakkil, S., Preece, R., Fletcher, C., et al. 2021, *ApJ*, 913, 60
- Popham, R., Woosley, S. E., & Fryer, C. 1999, *ApJ*, 518, 356
- Ramirez-Ruiz, E., & Fenimore, E. E. 2000, *ApJ*, 539, 712
- Rastinejad, J. C., Gompertz, B. P., Levan, A. J., et al. 2022, *Nature*, 612, 223
- Rees, M. J., & Meszaros, P. 1994, *ApJ*, 430, L93
- Roberts, N. S., Bell, A. F., & Main, I. G. 2015, *J. Volcanol. Geotherm. Res.*, 308, 127
- Salafia, O. S., Ghisellini, G., Pescalli, A., Ghirlanda, G., & Nappo, F. 2016, *MNRAS*, 461, 3607
- Salvaterra, R., Campana, S., Vergani, S. D., et al. 2012, *ApJ*, 749, 68
- Sari, R., & Piran, T. 1995, *ApJ*, 455, L143
- Shahamat, N., Abbassi, S., & Liu, T. 2021, *MNRAS*, 508, 6068
- Song, X.-Y., Xiong, S.-L., Zhang, S.-N., et al. 2022, *ApJS*, 259, 46
- Spada, M., Panaitescu, A., & Mészáros, P. 2000, *ApJ*, 537, 824
- Stern, B. E., & Svensson, R. 1996, *ApJ*, 469, L109
- Stratta, G., Nicuesa Guelbenzu, A. M., Klose, S., et al. 2025, *ApJ*, 979, 159
- Sun, H., Zhang, B., & Li, Z. 2015, *ApJ*, 812, 33
- Troja, E., Fryer, C. L., O’Connor, B., et al. 2022, *Nature*, 612, 228
- Tsvetkova, A., Frederiks, D., Golenetskii, S., et al. 2017, *ApJ*, 850, 161
- Tsvetkova, A., Frederiks, D., Svinkin, D., et al. 2021, *ApJ*, 908, 83
- Usov, V. V. 1992, *Nature*, 357, 472
- Virgili, F. J., Liang, E.-W., & Zhang, B. 2009, *MNRAS*, 392, 91
- von Kienlin, A., Meegan, C. A., Paciesas, W. S., et al. 2020, *ApJ*, 893, 46
- Wanderman, D., & Piran, T. 2010, *MNRAS*, 406, 1944
- Wang, F. Y., & Dai, Z. G. 2013, *Nat. Phys.*, 9, 465
- Wei, J.-J. 2023, *Phys. Rev. Res.*, 5, 013019
- Woosley, S. E. 1993, *ApJ*, 405, 273
- Yang, J., Ai, S., Zhang, B. B., et al. 2022, *Nature*, 612, 232
- Yonetoku, D., Murakami, T., Nakamura, T., et al. 2004, *ApJ*, 609, 935
- Zhang, B., & Yan, H. 2011, *ApJ*, 726, 90

⁷ https://github.com/mmanuele99/internal_shock_grbs

Appendix A: Multiple-shell interaction

Here we summarise the numerical procedure used to follow the dynamical evolution of multiple shells.

We consider an outflow composed of N shells, labelled by an index $i = 1, \dots, N$, according to their emission order from the central engine. Within this framework, the outermost shell corresponds to the first one emitted, while the innermost shell is the most recently emitted. The shell-ejection process is described as a stochastic Poisson process with a time-dependent rate, as introduced in Sect. 3, resulting in a higher concentration of shells in the outer regions of the outflow.

Each shell, emitted at a time t_i , is characterised by four quantities: the Lorentz factor γ_i , the mass $m_i = p\gamma_i^{-1}$, the initial radial position $R_i = c\beta_i t_i$, and the width l_i . The separation from the previously emitted shell L_i is given by $L_i = R_i - R_{i+1} - l_{i+1}$, where the position of the last-emitted shell is set to $R_N = 0$.

The dynamical evolution of the shell train proceeds as follows. Let j denote the index of the j th collision, occurring at time t_j and radius $R_c(t_j)$. Given the shell velocities β_i and separations L_i at time t_{j-1} , the collision time for each adjacent pair satisfying $\beta_{i+1} > \beta_i$ is $\delta t_{i,i+1} = L_i/[c(\beta_{i+1} - \beta_i)]$. Defining $\delta t_j = \min(\delta t_{i,i+1})$, the j th collision occurs at $t_j = t_{j-1} + \delta t_j$. Let s and r denote the indices of the slow and rapid shells involved in the collision at time t_j . The collision radius is $R_c(t_j) = R_r(t_j) = R_r(t_{j-1}) + c\beta_r \delta t_j$. The radiation produced in this interaction is detected by an observer located at a distance R_0 from the central engine at time $t_{\text{obs},j} = [R_0 - R_c(t_j)]/c + t_j$, where R_0 is taken to be equal to the total radial extent of the outflow. This time corresponds to the onset of the pulse described by Eq. (5).

After the collision, each shell advances from its previous position $R_i(t_{j-1})$ to $R_i(t_j) = R_i(t_{j-1}) + c\beta_i \delta t_j$. All shells retain their Lorentz factor, mass, and width, except for the s and r shells, which merge into a single shell labelled by index r , with properties determined as described in Sect. 3. The s shell is then removed from the system.

The above procedure is iterated until either all shells have merged into a single shell or the shells become ordered in such a way that their Lorentz factors increase monotonically outward, preventing further collisions.

Appendix B: Spectral parameters

The spectral parameters α , β , and $E_{p,i}$ of the Band function (Band et al. 1993), used to compute the flux correction factor k_{flux} described in Sect. 3, are assigned as follows.

The values of the spectral indices α and β were retrieved from the publicly available Fermi GBM Burst Catalog⁸. We downloaded all GRBs detected between 14 July 2014 and 14 October 2025, and discarded entries that do not satisfy the following constraints (Poolakkil et al. 2021):

- A : positive relative error < 0.44 and negative relative error < 0.83 ;
- α : positive error < 0.37 and negative error < 0.51 ;
- β : positive error < 1.0 and negative error < 0.53 ;
- E_p : positive relative error < 0.35 and negative relative error < 0.43 ;

where A denotes the spectral normalisation. This selection yields a total of 990 sets of spectral parameters, which we refer to as the good sample.

For each simulated GRB, the values of α and β are randomly drawn from the good sample. The intrinsic peak energy $E_{p,i}$ is assigned differently. Simply setting $E_{p,i} = E_p(1 + z_{\text{GRB}})$, with E_p randomly extracted from the good sample, would neglect the observed $E_{p,i}$ - L_{iso} correlation (Amati 2006; Yonetoku et al. 2004), where L_{iso} is the bolometric isotropic-equivalent peak luminosity.

We therefore proceed as follows. Starting from the value of L_{iso} derived from the theoretical LC, we compute the corresponding $E_{p,i}$ using the $E_{p,i}$ - L_{iso} relation reported in GS22, including the intrinsic scatter of the correlation.

The final values of α , β , and $E_{p,i}$ obtained through this procedure are those assigned to each simulated GRB.

Appendix C: Genetic algorithm and model parameter optimisation

Genetic algorithms are machine-learning algorithms inspired by evolutionary processes, in which a population of candidate solutions evolves over successive generations towards an optimal configuration, as measured by a predefined set of metrics. Within this framework, each candidate solution, or ‘individual’, is represented by a set of eight model parameters (Table C.1), analogous to a genetic chromosome, with each parameter corresponding to a ‘gene’. The evolution of the population is driven by the relative fitness of individuals, which determines their likelihood of contributing to subsequent generations. Evolutionary processes, such as crossover between parent individuals and random mutations of individual genes, are implemented to promote exploration of the parameter space and prevent premature convergence.

Each individual encodes a specific combination of model parameter values, sampled within the ranges reported in Table C.1. The optimisation proceeds through a fixed sequence of steps, identical to those described in M25, which we summarise here for completeness:

- the algorithm was run for 30 generations, each consisting of a population of $N_{\text{pop}} = 2000$ individuals.
- For each individual, $N_{\text{GRB}} = 2000$ synthetic LCs were generated using the corresponding parameter set. The same selection criteria applied to the real data (Sect. 2.1) were applied, and simulated LCs failing to meet these requirements were discarded.
- The metrics described in Sect. 2.2 were then computed for the selected simulated LCs and compared with the corresponding metrics derived from the real datasets. The comparison was quantified through an L_2 loss function (see M25 for details). The final loss assigned to each individual was defined as the average of the five metric losses, while the fitness score was taken as the inverse of this quantity.
- At each generation, individuals were ranked according to their loss values. New offsprings were produced by combining the genes of the top 15% fittest individuals. In addition, random mutations were introduced by resampling individual genes from their allowed ranges, with a mutation probability of 4% per gene.

Appendix D: Examples of light curves

Figure D.1 shows a few examples of simulated *Swift*/BAT GRB LCs obtained with the optimised IS model. Each panel reports the redshift z , the bolometric isotropic-equivalent peak luminosity $L_{\gamma,\text{iso}}$, the isotropic-equivalent gamma-ray energy $E_{\gamma,\text{iso}}$, and the peak photon flux P of the corresponding LC.

⁸ <https://heasarc.gsfc.nasa.gov/w3browse/fermi/fermigbrst.html>




Fig. D.1. Examples of simulated *Swift*/BAT GRB LCs obtained with the optimised IS model. The bolometric isotropic-equivalent peak luminosity $L_{\gamma, \text{iso}}$ and the isotropic-equivalent gamma-ray energy $E_{\gamma, \text{iso}}$ are expressed in erg s^{-1} and erg , respectively, while the peak photon flux P is given in $\text{ph s}^{-1} \text{cm}^{-2}$. For each LC, the bin time is 330 ms and corresponds to the shortest detection timescale obtained with FAST-MEPSA, a value that incidentally happens to be common to all four LCs.

Table C.1. Parameter space explored in the model optimisation.

Parameter	Lower bound	Upper bound	Prior
α	2	4	Uniform
S	1	4	Uniform
μ_{τ}	0	4	Uniform
σ_{τ}	0.1	2	Uniform
γ_{min}	10	250	Log-uniform
γ_{max}	350	1000	Log-uniform
$p(\text{g})$	10^{31}	10^{33}	Log-uniform
r	2	20	Log-uniform

Notes. The lower and upper bounds define the parameter ranges explored during the optimisation, while the priors indicate the sampling distributions adopted in the GA.