

# Methods and results of an automatic analysis of a complete sample of *Swift*-XRT observations of GRBs

P. A. Evans,<sup>1\*</sup> A. P. Beardmore,<sup>1</sup> K. L. Page,<sup>1</sup> J. P. Osborne,<sup>1</sup> P. T. O’Brien,<sup>1</sup>  
 R. Willingale,<sup>1</sup> R. L. C. Starling,<sup>1</sup> D. N. Burrows,<sup>2</sup> O. Godet,<sup>1</sup> L. Vetere,<sup>2</sup> J. Racusin,<sup>2</sup>  
 M. R. Goad,<sup>1</sup> K. Wiersema,<sup>1</sup> L. Angelini,<sup>3</sup> M. Capalbi,<sup>4</sup> G. Chincarini,<sup>5,6</sup> N. Gehrels,<sup>3</sup>  
 J. A. Kennea,<sup>2</sup> R. Margutti,<sup>5,6</sup> D. C. Morris,<sup>3,7</sup> C. J. Mountford,<sup>1</sup> C. Pagani,<sup>1</sup> M. Perri,<sup>4</sup>  
 P. Romano<sup>8</sup> and N. Tanvir<sup>1</sup>

<sup>1</sup>*X-ray and Observational Astronomy Group, Department of Physics and Astronomy, University of Leicester, Leicester LE1 7RH*

<sup>2</sup>*Department of Astronomy and Astrophysics, Pennsylvania State University, 525 Davey Lab, University Park, PA 16802, USA*

<sup>3</sup>*NASA/Goddard Space Flight Center, Greenbelt, MD 20771, USA*

<sup>4</sup>*ASI Science Data Center, ASDC c/o ESRIN, via G. Galilei, 00044 Frascati, Italy*

<sup>5</sup>*INAF, Osservatorio Astronomico di Brera Via E. Bianchi 46, I-23807, Merate (LC), Italy*

<sup>6</sup>*Università degli Studi di Milano-Bicocca, Dipartimento di Fisica, Piazza della Scienza 3, 20126, Milano, Italy*

<sup>7</sup>*Center for Nuclear Studies, Department of Physics, The George Washington University, Washington, DC 20052, USA*

<sup>8</sup>*INAF, Istituto di Astrofisica Spaziale e Fisica Cosmica, Via U. La Malfa 153, I-90146 Palermo, Italy*

Accepted 2009 April 15. Received 2009 April 15; in original form 2008 December 12

## ABSTRACT

We present a homogeneous X-ray analysis of all 318 gamma-ray bursts detected by the X-ray telescope (XRT) on the *Swift* satellite up to 2008 July 23; this represents the largest sample of X-ray GRB data published to date. In Sections 2–3, we detail the methods which the *Swift*-XRT team has developed to produce the enhanced positions, light curves, hardness ratios and spectra presented in this paper. Software using these methods continues to create such products for all new GRBs observed by the *Swift*-XRT. We also detail web-based tools allowing users to create these products for any object observed by the XRT, not just GRBs. In Sections 4–6, we present the results of our analysis of GRBs, including probability distribution functions of the temporal and spectral properties of the sample. We demonstrate evidence for a consistent underlying behaviour which can produce a range of light-curve morphologies, and attempt to interpret this behaviour in the framework of external forward shock emission. We find several difficulties, in particular that reconciliation of our data with the forward shock model requires energy injection to continue for days to weeks.

**Key words:** methods: data analysis – catalogues – gamma-rays: bursts.

## 1 INTRODUCTION

The *Swift* satellite (Gehrels et al. 2004) has revolutionized our understanding of gamma-ray bursts (GRBs), and the X-ray telescope (XRT; Burrows et al. 2005) has played a major role in this process. For example, XRT data have shown that the early X-ray light curve is much more complex than originally thought, often containing a shallow-decay ‘plateau’ phase interpreted variously as energy injection in a forward shock (Zhang et al. 2006), reverse shock emission (Uhm & Beloborodoc 2007) and dust-scattering of the prompt GRB emission (Shao & Dai 2007; although this has recently been ruled out by Shen et al. 2009). The high impact of the XRT has been made possible by *Swift*’s unique rapid slewing capability: the large

field-of-view Burst Alert Telescope (BAT; Barthelmy et al. 2005a) detects GRBs, and *Swift* then automatically repoints itself so that the narrow-field XRT and UV/Optical Telescope (UVOT; Roming et al. 2005) can observe the new burst within ~60–90 s of the trigger. The XRT has detected ~95 per cent of *Swift*-BAT-triggered GRBs, most of those within minutes of the trigger, usually providing the most accurate rapidly available localization of the GRB and enabling early follow up with ground telescopes. *Swift*’s X-ray afterglow light curves and spectra, starting at typically <100 s, have been crucial to the discoveries made by *Swift* (see O’Brien et al. 2006; Zhang 2007; Willingale et al. 2007, for reviews of the impact of *Swift* on GRB science).

Because GRBs fade rapidly, follow-up observers need to make quick decisions about whether or not to invest observing time on a given burst. This is especially true of potentially unusual bursts, such as high-redshift candidates or under-luminous GRBs. Thus,

\*E-mail: pae9@star.le.ac.uk

it is desirable for GRB data to be rapidly available and analysed quickly, reliably and ideally in a uniform manner. *Swift* data are downlinked many times per day and are immediately processed by the Swift Data Center (SDC) at Goddard Space Flight Center and made available to the public via the SDC and data centres in the UK and Italy, minutes to hours after the downlink. We (the *Swift*-XRT team) have developed software which, when data arrive at the UK Swift Science Data Centre (UKSSDC), automatically determines the best possible XRT GRB position and builds X-ray light curves and spectra; the results are then published on the internet. This provides a homogeneously generated catalogue of data, which we detail and discuss in this paper (Section 4), following a description of the software.

There are two types of GRB follow-up data telemetered from *Swift*: TDRSS and Malindi data (described below). Initially, our software only worked with Malindi data; however in 2008 February we modified our ground-based software to work with TDRSS data as well (Evans et al. 2008a); this provides better positions, positions of fainter GRBs, and more reliable light curves and spectra than those produced on board (although these should not be used for scientific analysis).

### 1.1 TDRSS data

When *Swift* first detects and observes a GRB, some data are immediately telemetered to the ground via NASA's Tracking and Data Relay Satellite System (TDRSS). Among these 'TDRSS data' are the XRT Single Pixel Event Report ('SPER') data. *Swift*'s XRT selects between Windowed Timing (WT; high time resolution but only 1D spatial information) and Photon Counting (PC; lower time resolution but full spatial information) modes automatically, based on the count rate in the central portion of the CCD (Hill et al. 2005). When the XRT enters PC mode during the first look at a new GRB, event lists containing every single-pixel (grade 0) event above 0.55 keV detected within the central  $200 \times 200$  pixel region, the SPER data, are delivered to the XRT team via the Gamma-ray Burst Co-ordinates Network (GCN) system every 2 minutes until *Swift* slews away from the burst (up to  $\sim 2$  ks after the trigger). It must be noted that while we have made our products as reliable as possible, SPER data are not fully calibrated and the light curves and spectra are intended as quick-look products; they should not be used for scientific analysis.

### 1.2 Malindi data

'Malindi data' are available several hours after a GRB trigger and comprise the full observation data set. Nine or ten times per day, *Swift* passes over the Malindi ground station in Kenya and downlinks the data buffered on board. These are then processed at the SDC and delivered to archives in the US, UK and Italy, typically 90–120 min after the Malindi downlink. *Swift* also observes GRBs which did not trigger the BAT, but are uploaded as Targets of Opportunity (ToOs). For these bursts only Malindi data are telemetered.

Malindi data are grouped into *observations*, usually one per day, each with a unique ObsID. A single observation may contain many *snapshots*; that is, pointings toward the source, since *Swift* is in a low Earth orbit and thus its targets get occulted by the Earth once per orbit. The terminology *observations* and *snapshots* is standard *Swift* parlance,<sup>1</sup> and will be used throughout this paper.

### 1.3 GRB afterglow models

In this paper, we introduce tools to produce high-precision positions, light curves and spectra from XRT data, and describe the automatic application of these tools to GRBs. This gives us a catalogue of results for all 318 GRBs detected by the *Swift*-XRT up to GRB 080723B, and we discuss the implications for afterglow science from this analysis. While our data set can be used to test any models for GRB emission, we do this in the context of the fireball model (e.g. Rees & Mészáros 1994; Sari, Piran & Narayan 1998), which is the current consensus model.

In this model, the GRB progenitor launches highly relativistic jets of material in a series of shells, of differing bulk Lorentz factors. Internal collisions (i.e. within the jet) between shells cause shocks which radiate the GRB 'prompt emission'. The X-ray data presented in this paper may contain the tail of this prompt emission, however it is thought to arise predominantly from the afterglow. This is emission from an external shock which forms where the jet is decelerated by the circumburst medium, and which propagates into that medium, cooling by synchrotron radiation as it does so (see Piran 2005 for a comprehensive review of the fireball model).

### 1.4 Layout of the paper

This paper is organized as follows.

Section 2 (pp 2–10): XRT automatic analysis tools.

Section 2.1 (pp 3–4): Spectra.

Section 2.2 (pp 4–7): Positions.

Section 2.3 (pp 7–8): Light curves and hardness ratios.

Section 2.4 (pp 8–10): Automatic light-curve fitting.

Section 3 (pp 10–11): Web tools to analyse any XRT source.

Section 4 (pp 11–21): Results of XRT GRB analysis.

Section 5 (pp 21–23): A canonical X-ray light curve?

Section 6 (pp 23–24): Understanding the X-ray afterglow.

Throughout this paper, errors quoted in all tables are at the 90 per cent confidence level. Error bars on data in all figures are  $1\sigma$  uncertainties.

## 2 AUTOMATED DATA PRODUCTS

We have developed software to produce three types of data products: 'enhanced' positions, light curves and spectra. The spectra were announced in Evans et al. (2008b), and are presented in Section 2.1. The method of enhancing XRT positions has been previously documented (Goad et al. 2007, hereafter G07); however we have improved the algorithm, resulting in a factor of  $\sim 2$  improvement in precision (Section 2.2). The light-curve code has been published by Evans et al. (2007) and only minor modifications are described here (Section 2.3) along with details of new functionality which has been made available to the user. We also describe automatic light-curve fitting in Section 2.4.

For GRBs which trigger the BAT, these products are created automatically while GRBs observed as ToOs by *Swift* must be manually registered for automatic analysis. This is usually done at the time of the ToO upload, so the data products are available as rapidly as for BAT-detected bursts. Before any of these products are produced, the XRT data are reprocessed at the UKSSDC using the latest release of the XRTPIPELINE tool.<sup>2</sup> This may differ from the version used at the SDC to create the cleaned event lists available from the quick-look

<sup>2</sup> Part of the XRT software, distributed with the HEASOFT package: <http://heasarc.gsfc.nasa.gov/heasoft/>.

<sup>1</sup> <http://heasarc.gsfc.nasa.gov/docs/swift/archive/archiveguide1/node3.html>

and archive sites. When a new version of the *Swift* software or calibration is released, we carry out some tests to confirm our product-generation code works reliably with the new release and then switch to using the latest version. We do not, however, reprocess earlier GRBs with the new release of the software.

## 2.1 Spectra

To create a spectrum from Malindi data, an image is formed from the first available PC mode event list, and `XIMAGE` is used to identify any sources. The brightest source within the BAT error circle (or equivalent for bursts detected by other missions) is assumed to be the GRB, and the sky-coordinate point-spread-function (PSF) fit routine developed for the position enhancement (Section 2.2) is used to determine the source position in the XRT astrometric frame. We use only observations which begin within 12 h of the first one. Each observation is subdivided into snapshots, and these may be further subdivided into times where pile up – is or is not an issue. To identify intervals affected by pile up, we first search for times where the count rate within a 30 pixel radius circular region centred on the source is above  $0.6 \text{ counts s}^{-1}$  in PC mode or  $150 \text{ counts s}^{-1}$  in WT mode. In PC mode, we then obtain the PSF profile of the source and compare it to the calibrated, non-piled-up PSF (Moretti et al. 2005). This indicates not only whether the source is piled up, but also the radius out to which pile up is a factor,  $R_p$ . If the source is piled up, we use an annular extraction region to obtain source data, with an inner radius  $R_p$ . For WT mode data, we assume the data are piled up whenever the count rate exceeds  $150 \text{ counts s}^{-1}$  and use a box annulus extraction region, where the inner radius is that necessary to keep the measured count rate below this level. For alternative methods of identifying and eliminating pile up, see Vaughan et al. (2006) or Romano et al. (2006). Where the data are not piled up, we use a circular (rectangular) extraction region for PC (WT) mode. For each snapshot, we determine the mean source count rate and use this to choose the radius of the source extraction region, according to Table 1; this maximizes the signal-to-noise ratio of the spectrum and is identical to the method used for light curves.

Once the data have been divided into time intervals – each with a source extraction region – we generate a source spectrum, source event list and full frame event list for each one. Using the full frame event lists, we create an exposure map and (using this and the source spectrum) an Ancillary Response File (ARF) per interval. We then combine the source event lists from each time interval using `EXTRACTOR` to get a single source spectrum. We similarly combine the ARFs, using the `ADDARF` tool; each ARF is weighted according

to the proportion of counts in the total source spectrum which came from this time interval. This is not the same as extracting a single spectrum, exposure map and ARF from an event list spanning multiple snapshots, since in that case the weighting of the snapshots is determined by the exposure map (hence exposure in each orbit), but GRBs are not of constant brightness. The ARFs must be correctly weighted since the proximity of the source to the bad columns on the CCD, and hence the effective detector area, will change from one snapshot to the next. That weighting by counts is the correct approach is readily demonstrated: the true number of counts from the source ( $C_t$ ) is simply the measured counts ( $C_m$ ) multiplied by some correction factor ( $Q$ , implicit in the ARF) which reflects counts lost to (for example) the bad columns. If there are multiple snapshots, the true total number of counts is  $C_t = \sum(C_m Q) = \sum C_m \times \frac{\sum(C_m Q)}{\sum C_m}$ , i.e. the overall correction factor (or ARF) is the count-weighted mean of the individual ARFs.

The `BACKSCAL` keyword must be set for the source spectrum, for use by `XSPEC` (Arnaud 1996). Since this can be different for each time interval included (the source region size is variable), this is taken as the weighted mean of the `BACKSCAL` values from the individual source spectra extracted per time interval, weighted according to the number of counts in those spectra.

A background spectrum is also produced. In WT mode, background data are extracted from the entire window, excluding a 120-pixel (283 arcsec) wide box centred on the source. In PC mode, the background region is an annulus with an inner radius of 60 pixels (142 arcsec) and an outer radius of 110 pixels (260 arcsec); if this extends beyond the edge of the detector window, it is shifted accordingly; the inner circle of course remains centred on the source. To create the background spectrum, we do not subdivide the data more finely than the observations. For each observation, we identify any sources in the background region using `XIMAGE` and exclude those areas from the extraction region. The individual observations' background spectra are then combined as for the source spectra.

The spectra thus produced were extensively verified by the *Swift*-XRT team. Our test procedure consisted of producing spectra both manually and using this software. These were both fitted with the same model and the best-fitting parameters compared. We did this for more than 30 spectra, and in every case, the difference between the parameters from the automatic and manual spectra was much less than the uncertainties on those parameters. We therefore conclude that the automatic generation of spectra is reliable.

Spectra are automatically created for each new GRB observed by the XRT, and are updated as new data are received. The results are posted online at [http://www.swift.ac.uk/xrt\\_spectra](http://www.swift.ac.uk/xrt_spectra) in postscript and GIF format. We also provide a tar archive for download which contains the source and background spectra and the ARF file necessary for users to fit the data themselves.<sup>3</sup> The results of automatic spectral fitting (below) are available from the same website.

Since the creation of an XRT spectrum is useful for any target observed by *Swift*, not just GRBs, we have created a tool to allow users to build spectra using our software for any object observed by the XRT (see Section 3).

**Table 1.** Source extraction radii used for given PC-mode count rates.  $R$  is the measured, uncorrected count rate. This table is reproduced from Evans et al. (2007). Values are given in XRT pixels and arcseconds: one XRT pixel corresponds to 2.36 arcsec.

Count rate $R$ (counts $\text{s}^{-1}$ )	Source radius in pixels (arcsec)
$R > 0.5$	30 (70'8)
$0.1 < R \leq 0.5$	25 (59'0)
$0.05 < R \leq 0.1$	20 (47'2)
$0.01 < R \leq 0.05$	15 (35'4)
$0.005 < R \leq 0.01$	12 (28'3)
$0.001 < R \leq 0.005$	9 (21'2)
$0.0005 < R \leq 0.001$	7 (16'5)
$R \leq 0.0005$	5 (11'8)

### 2.1.1 Automatic spectral fitting

After creating the spectra, our software automatically models them with an absorbed power law. Fitting is performed using the

<sup>3</sup> Users also require a Response Matrix File (RMF) file, which can be found in the Calibration Data base (CALDB; <http://swift.gsfc.nasa.gov/docs/heasarc/caldb/swift/>).

$C$ -statistic in XSPEC 12; we first apply a *group min 1* command in GRPPHA (necessary for XSPEC to correctly calculate the  $C$ -statistic; Arnaud, private communication). An RMF must also be used to fit the data; the XRTMKARF task used to create the ARF files when compiling the spectrum selects the appropriate RMF from the CALDB, and we supply this file to XSPEC. For a review of the XRT spectral response, see Godet et al. (2009).

The NHFTOOL<sup>4</sup> is used to determine the Galactic column density in the direction of the burst using the map from Kalberla et al. (2005), and the spectrum is fitted in with the model *phabs\*phabs\*pow*. The first absorption component is frozen at the Galactic value and the second is free to vary. The abundances are fixed at those from Anders & Grevesse (1989). A TCL script (SHAKEFIT) developed by Simon Vaughan (Hurkett et al. 2008; see their section 3.2.2), which uses the ERROR command to detect and recover from local minima, is then used to find the true global best fit. If a spectroscopic redshift has been reported in the literature, the XRT team can supply this, and the second absorption component in the model is replaced with a *zphabs* component, with the redshift frozen at the reported value. The 90 per cent confidence intervals of each free parameter are found using the ERROR command in XSPEC which steps the parameter of interest and repeats the fit until the  $C$ -statistic has worsened by 2.706 compared to the best-fitting value. The observed and unabsorbed 0.3–10 keV flux for the model are also obtained from XSPEC, and the former of these is used to determine the count-rate to flux conversion factor for the GRB. This is then automatically applied to the count-rate light curve to produce flux units version.

Very occasionally (<1 per cent of the time), XSPEC finds a local minimum of the  $C$ -statistic rather than the best fit. This is usually immediately obvious from the plot presented on the web pages, or because the best-fitting values are unusual (the probability distributions in Section 4 give a quantitative definition of ‘unusual’). In this case, a member of the XRT team will determine fitting parameters and supply these to the software, which will then repeat the fit, using these parameters as the initial values. When new data arrive and the spectrum is updated, the manually supplied values will again be used as the initial values for the fit. All of the results presented in this paper have been verified by visual inspection, and the few with poor fits corrected.

The results of these fits are presented in a table on the web page for each GRB, accessible via <http://www.swift.ac.uk/xrt.spectra>.

### 2.1.2 Time-resolved spectra

Only time-averaged GRB spectra are produced automatically, however the software allows for arbitrary time intervals to be specified, over which the spectra can be compiled: on the web page presenting the spectra for each GRB is an option to ‘create time-sliced spectra’. Following this link, users can specify times of interest and their spectra will be built. If a member of the XRT team decides that a particular time-resolved spectrum should be made available along with the time-averaged spectra for a GRB, they will add that spectrum to the main spectrum results page of the burst.

### 2.1.3 Application to SPER data

Spectra can also be obtained from the rapidly available SPER data using this software. Note, however, that the SPER data have not been fully calibrated and SPER spectra should be used as a quick-look

product, rather than for a scientific analysis. Additional processing of the SPER data is necessary to prepare them for spectral extraction. All SPER messages are combined into a single file, and a value of 100 is subtracted from the PHA (uncalibrated event energy) column (this was added on board to avoid negative values being obtained after bias subtraction). A ‘Grade’ column is then added (all events in SPER are grade 0). The tool XRTCALCPI is executed to create a PI (calibrated event energy) column from which spectra can be built. At this point, a spectrum can be extracted. For SPER data, we do not subdivide the data into intervals since they only cover a single snapshot. If the data are piled up at any point an annular source region is used for all of the SPER data. This means that we have only one source spectrum and ARF. Fitting is performed as for Malindi data (see Section 2.1.1). Because SPER data only include events above 0.55 keV, it is harder to constrain the absorbing column.

## 2.2 ‘Enhanced’ positions

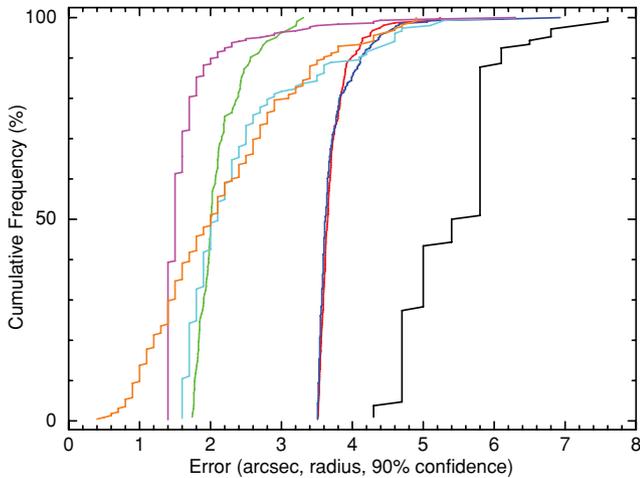
The technique of ‘enhancing’ XRT positions from Malindi data by determining the spacecraft attitude from images taken with the UVOT was first presented in G07. This technique produced positions with error radii typically 40 per cent smaller than the ‘standard’ positions determined using the spacecraft attitude information determined from the on-board star trackers. With all available data, the G07 positions were a little less precise than those determined by Butler (2007) using serendipitous X-ray sources to find the correction; but while the latter typically have large initial errors and only offer improvement over the ‘standard’ positions  $\sim 1$  day after the trigger, the G07 positions were available within hours of a trigger.

We have made significant revisions to the enhancement process, reducing the error radii by a further 30–50 per cent. These positions are still available within hours of a trigger, and now in  $\sim 75$  per cent of cases give better precision than those of Butler et al. (2007), and are the most accurate XRT positions available. Furthermore, we have developed a version of this algorithm applicable to the SPER data, reducing the error radius of the prompt localizations by up to 60 per cent. These positions are typically available 10–20 min after a GRB trigger. The relative precisions of the different XRT positions are shown in Fig. 1.

Although we concentrate here on the improvements made since G07, it is necessary to briefly summarize the algorithm used there, to give the context for the improvements (for full details, see G07).

To produce an enhanced XRT position using the G07 method, the available data are first split into ‘overlaps’ – times of simultaneous XRT (PC mode) and UVOT ( $v$ -band) data. For each such overlap, a detector co-ordinate X-ray image is produced and the GRB localized therein. This is done using a PSF fit which corrects for the effects of the bad columns on the CCD. The position is then transformed into an equivalent position in UVOT detector co-ordinates, and thence into UVOT sky co-ordinates, using the attitude information from the star trackers on board *Swift*. An image of the UVOT field of view is also constructed, with the sky co-ordinates calculated using the same attitude information. Serendipitous sources in the UVOT image are matched to the USNO-B1 catalogue, giving the quaternion needed to correct the image’s attitude. This quaternion is then applied to the previously calculated XRT GRB position in UVOT sky co-ordinates, to give the ‘enhanced’ XRT position for that overlap. This process is performed for every available overlap, the weighted mean of all overlaps is then calculated and systematic errors are added to give the enhanced XRT position of the GRB. Since G07, this process has been improved in the following ways.

<sup>4</sup> <http://heasarc.gsfc.nasa.gov/heasoft/ftools>



**Figure 1.** The cumulative frequency of the 90 per cent confidence error radius for XRT GRB positions determined using different techniques. Black, on-board positions; red, positions from SPER data derived using a PSF fit; green, enhanced SPER positions; blue, ‘refined’ positions obtained from Malindi data using `XRTCENTROID`; cyan, enhanced Malindi positions using the G07 algorithm; magenta, the enhanced Malindi positions using the algorithm presented here; orange, positions from Butler et al. (2007). For the latter, we have removed positions with errors  $>5$  arcsec before calculating the distribution since these occur when there is insufficient data to correct the position and bias the plot towards larger uncertainties than are found when positions are successfully improved.

(1) The XRT-UVOT map, originally only derived for the UVOT  $v$  filter, has been extended to all UVOT filters. We have found, however, that limiting the software to use just  $v$ ,  $b$  and white filters gave the best results.<sup>5</sup>

(2) Multiple observations can now be used, whereas originally only a single one was used (recall that a single observation can still contain many snapshots). We do not include all available observations as the X-ray afterglow will generally be too faint to detect in a single overlap after  $\sim 12$ – $24$  h and the inclusion of more observations slows the enhancement process significantly. Instead, we use any observation which began within 12 h of the first. Since *Swift*’s automatic response to a new GRB is a 60-ks observation, in many cases this still means only the first observation is used and this change to the code then has no impact. It is not uncommon, however, for the automatic observations to be interrupted, either manually [to avoid overheating the XRT; Kennea et al. (2005)] or automatically (i.e. *Swift* detects another GRB), in which case multiple observations are suitable for use in the enhancement process.

(3) We have substantially re-written the PSF fitting code used to determine the position of the object on the XRT. The  $C$ -statistic was previously calculated by comparing the number of counts expected under the PSF with those in the data (Cash 1978). It is now calculated by comparing the value of each pixel in the image with that predicted by the model PSF. Additionally, the fit was previously performed in XRT detector co-ordinates [so that the positions of the bad columns on the CCD caused by the micrometeoroid impact (Abbey et al. 2005) were known], but it is now performed in sky co-ordinates, using exposure maps to compensate for bad columns or hot pixels.

These changes improve the enhanced positions for several reasons.

<sup>5</sup> This is probably because the USNO-B1 catalogue is an optical catalogue so matching UV images to it is error prone.

(i) Although *Swift* does not remain perfectly steady while on-source, the star trackers provide precise *relative* attitude information. Thus, while the detector-coordinate image may be distorted by spacecraft motion, the sky image will not be, and the PSF fit in this co-ordinate system will be more precise.

(ii) UVOT data are corrected for the spacecraft movement before being downlinked. Every photon detected is shifted to the position it would have had if the spacecraft had not moved since the image exposure began. As a result, the effective time of the UVOT image is the start of the exposure. In contrast, a detector-coordinate XRT image has no such correction and its effective time is the mean photon arrival time within the image. Thus, the aspect solution determined from the UVOT image is actually for a different time, and possibly slightly different attitude, than that at which the GRB position was measured on the XRT detector. This was a contributor to the empirically derived systematic error in G07 (see their section 6). XRT sky-coordinate images are built using the sky-coordinate position of each event in the event lists which were calculated on an event-by-event basis using the spacecraft attitude information, and are thus analogous to the UVOT sky images. The sky co-ordinate position of an object is therefore stable through a snapshot, allowing the `POINTXFORM` tool to be used to convert the measured sky-coordinate position into XRT detector co-ordinates at the start of the UVOT exposure, so the aspect solution and corrected position can be made contemporaneous.

(iii) Because the GRB position in XRT sky co-ordinates is stable within a snapshot, we no longer need strictly simultaneous XRT and UVOT data. Instead, we use the longest XRT exposure possible without using the same data in multiple overlaps, up to a maximum duration of an entire snapshot. For example, if a snapshot contains two UVOT images, one in the  $v$  filter and one in white, the XRT data will be split in two – the division occurring between the UVOT exposures. An XRT position will be found for each part of the snapshot and enhanced using the approximately contemporaneous UVOT data. On the other hand, if the snapshot contains  $v$  and  $uvm2$  filter UVOT images, the XRT position will be determined using the entire snapshot of data, and then enhanced using the  $v$  image (we do not use the UV filter images to enhance positions; see above). This means that overlaps can be used where the XRT source is fainter than was possible in the previous version.

(4) We have supplemented the improvement to the PSF fitting code by deducing PSF profiles of piled up sources, in addition to the profile in the CALDB (Moretti et al. 2005). For each XRT image, we perform the PSF fit six times; once with each PSF profile. The fit with the lowest  $C$ -statistic is then assumed to give the most accurate position of the GRB.

To determine the piled-up PSF profiles, we identified objects observed by *Swift* in PC mode which were piled up, away from the bad columns on the detector and approximately constant in brightness. We then obtained the PSF profile using `XIMAGE` and modelled it with a King–Gaussian function; the latter component reflecting the counts lost to pile up. The centre of the Gaussian profile was frozen at zero (i.e. the centre of the PSF). The PSF profiles thus deduced are given in Table 2, defined as  $\text{King}(R_c, \beta) - N_{\text{gaus}} \text{Gaus}(\sigma)$ , where  $N_{\text{gaus}}$  represents the relative normalization of the Gaussian component compared to the King component.

(5) *Swift* data are normally delivered to the data centres when the data from all three instruments have been processed. However, during the first observation of a new GRB special ObsIDs are created which contain just the XRT or UVOT data. These ObsIDs – which end with ‘991’ for XRT or ‘992’ for UVOT – are usually available

**Table 2.** PSF profiles for piled-up sources of different count rates. The profile is defined as  $\text{King}(R_c, \beta) - N_{\text{gaus}}\text{Gaus}(\sigma)$ .

Count rate (counts s <sup>-1</sup> )	King $R_c$ (arcsec)	King $\beta$	$N_{\text{gaus}}$	Gaus $\sigma$ (arcsec)
0.90	4.654	3.321	0.522	3.019
1.36	6.578	3.627	0.504	4.139
2.59	6.620	3.606	0.634	3.990
5.15	12.790	3.505	1.039	6.411
8.58	19.880	3.889	1.245	6.682

**Table 3.** The peak and FWHM of the PDFs plotted in this paper, where the PDFs have been approximated as Gaussians.  $a-d$  refer to the different classes of GRB illustrated in Fig. 9, e.g.  $\alpha^b_{\text{steep}}$  refer to the decay index of the steep-decay phase of those light curves which begin steep and then flatten. Note: the underlying distributions are not always Gaussian, or are poorly sampled so the values herein are indicative.

Parameter	Peak of PDF	FWHM
$\alpha$	1.1	1.5
$\beta$	0.98	0.66
$\log(N_{\text{H}})$ cm <sup>-2</sup>	21.3	1.2
Counts-to-flux factor (observed)	$3.8 \times 10^{-11}$ (erg cm <sup>-2</sup> ct <sup>-1</sup> )	$1.4 \times 10^{-11}$ (erg cm <sup>-2</sup> ct <sup>-1</sup> )
$\alpha^a_{\text{plat}}$	0.32	0.79
$\alpha^a_{\text{norm}}$	1.2	0.5
$\beta^a_{\text{plat}}$	1.0	0.53
$\beta^a_{\text{norm}}$	1.1	0.65
$\log(T_{\text{plat,start}})$	2.5	0.7
$\log(T_{\text{plat,end}})$	4.0	1.2
$\alpha^b_{\text{steep}}$	~2.7	~2.1
$\alpha^b_{\text{shallow}}$	0.82	0.46
$\beta^b_{\text{steep}}$	~1.1	~1.3
$\beta^b_{\text{shallow}}$	1.21	0.68
$\log(T^b_{\text{break}})$	2.8	0.75
$\alpha^c_{\text{steep}}$	~1.2	~0.9
$\alpha^c_{\text{shallow}}$	~0.7	~0.8
$\beta^c_{\text{steep}}$	1.1	0.6
$\beta^c_{\text{shallow}}$	0.98	0.38
$\log(T^c_{\text{break}})$	~4.0	~2.8
$\alpha^d_{\text{steep}}$	1.1	0.7
$\beta^d_{\text{steep}}$	1.1	0.6

before the instrument-combined data set. We have modified our software to use these special ObsIDs, rather than waiting for the ‘all-in-one’ ObsID data. This means that the positions are available up to half an hour earlier than previously.

After implementing these changes, we ran the code for every GRB observed by *Swift*, obtaining positions for 83 per cent. Following G07, we compared these positions to the UVOT positions for those GRBs which were detected by that instrument (taken from the GCN circulars). For each burst with an enhanced XRT position and a UVOT position, we calculated the angular distance between these positions,  $R$ . We also calculated the combined 90 per cent confidence error in the enhanced XRT and UVOT positions  $E_t = \sqrt{(E_{\text{XRT,enh}}^2 + E_{\text{UVOT}}^2)}$ . If  $R/E_t \leq 1$  the positions agree. If we included only the statistical errors and the systematic uncertainty arising from the

XRT-UVOT map (1.3 arcsec), we found that fewer than 90 per cent of enhanced positions agreed with the UVOT positions, as G07 did. However, we needed to increase our systematic only marginally, to 1.36 arcsec, to achieve 90 per cent agreement, whereas in G07 we needed a systematic of 1.5 arcsec. Our error radii are now  $\leq 1.5$  arcsec 50 per cent of the time, and  $\leq 2.0$  arcsec 90 per cent of the time. As Fig. 1 shows, these represent the best available XRT positions for the majority of GRBs.

The enhanced positions of all GRBs observed by *Swift* are available online at [http://www.swift.ac.uk/xrt\\_positions](http://www.swift.ac.uk/xrt_positions). When a new GRB is observed by *Swift*, an enhanced position is automatically produced and added to this page as soon as the necessary data are available. A GCN circular is also dispatched to advise the community of the position, since in the majority of cases this represents the best available position at this time and its rapid dissemination helps follow-up astronomers. If no XRT counterpart was identified in the TDRSS data, this circular is not automatically dispatched. Under such circumstances, the software enhances the position of the brightest source within the BAT error circle; experience has shown that this is not always the GRB. Positions will none the less be posted online even in these circumstances. As *Swift* continues to observe the new GRB, the position is updated with each new delivery of data. These updates appear at the above URL, but are not disseminated by other means.

Enhanced positions cannot be produced for all GRBs. If the burst is too faint to be seen by XRT in a single snapshot, the position cannot be found. Also, sometimes the `UVOTSKYCORR` tool is unable to find sufficient matches between UVOT field stars and USNO-B1 entries to return an aspect solution. Full details of less common reasons why a position is not produced were given in G07. To date, 83 per cent of GRBs detected by the XRT have enhanced positions. If we consider only BAT-triggered GRBs and exclude those where the UVOT was not in operation, this number rises to 90 per cent.

Since many *Swift* non-GRB observations use the XRT to determine the position of a newly discovered (or known, but poorly localized) object, we have developed a web tool to apply the position enhancement to user-selected non-GRB objects (see Section 3 for details).

### 2.2.1 Application to SPER data

We can also apply the enhancement software to the SPER data. This was not previously possible because the XRT and UVOT data had to be simultaneous – for the limited TDRSS data this is often not the case, or requires the XRT SPER data to be filtered on time; given that SPER data contain substantially fewer events than Malindi data, this made determining the XRT position unreliable. Since the modification to the PSF fit (point 3 above) removed the simultaneity requirement, we can enhance positions using TDRSS data. The process is almost identical to the Malindi data process described above, with the following differences.

(i) Since we always have just one snapshot, the XRT position is determined once, using all available data, and corrected multiple times; the statistical uncertainty in the XRT position is thus only applied once, at the end of the process.

(ii) We use the UVOT source lists telemetered via TDRSS, rather than the UVOT images, to determine the aspect solution.

Point (ii) arises because the TDRSS source lists are calculated on board from the full field-of-view image, however only a sub-image is downlinked; thus the on-board source list provides more

matches to the USNO-B1 catalogue and a solution can be found more frequently. The source localization performed on board is less sophisticated than is done on the ground, increasing the systematic uncertainty needed to achieve 90 per cent agreement between our ‘enhanced SPER’ positions and the UVOT positions to 1.7 arcsec, compared to 1.36 arcsec for Malindi data. In 2008 October, the filter sequence followed by UVOT in its initial burst response was changed to use the *U* filter for the second exposure rather than the White filter (the *U* filter has a lower background rate and is thus expected to find more afterglows than the White filter). We do not currently have sufficient data to evaluate the effect of this change on the enhanced SPER positions directly, however using historical Malindi data and *U*-filter images taken later in the observations, we determined that our systematic error needed to increase from 1.7 to 1.9 arcsec if the *U*-filter images were to be used. On the other hand, not using the *U*-filter TDRSS data would mean most enhanced SPER positions would be based on only one UVOT aspect solution and consequently have errors typically 0.2 arcsec larger than hitherto. They would also be produced longer after the burst as they would need to wait for the third UVOT exposure. Thus, on 2008 November 13, we modified the SPER position enhancement to use the *U*-filter images and a 1.9 arcsec systematic error.

### 2.3 Light curves and hardness ratios

The software we used to generate automatic XRT light curves and hardness ratios of GRBs was previously presented in Evans et al. (2007). This software has received minor revision, which we describe briefly below, along with a significant increase in functionality (Evans et al. 2008b).

#### 2.3.1 Modifications to the software

(1) Originally, the `XRTCENTROID` tool was used to determine the source position in XRT co-ordinates. When the bad columns on the CCD intersected the source PSF, this could result in an inaccurate position. This in turn could cause the pile-up or bad column correction factors to be incorrect. We have altered the software to use the PSF fit instead, since this gives accurate positions despite the bad columns. This made small changes to a few light curves.

(2) In the original version of the software, the final bin was always plotted as an upper limit if it contained fewer than 15 counts (i.e. the errors could not be considered Gaussian). In many cases, the source is clearly detected at these times and an upper limit is inappropriate. In 2007 June, we modified the software such that, if the final bin contains fewer than 15 counts, the Bayesian approach of Kraft, Burrows & Nousek (1991) is used to determine whether the source is detected at the  $3\text{-}\sigma$  level. If it is the case, the Bayesian method is then used to determine the  $1\text{-}\sigma$  confidence interval on the count rate, and this is plotted as a point on the light curve. Otherwise, an upper limit is plotted as previously. Subsequent data deliveries are still included in this final bin until it contains at least 15 counts, at which point the errors on the bin are calculated using Gaussian statistics and a new bin is begun.

(3) In 2008 July, we fixed a minor bug in the software, which occasionally caused Good Time Intervals (GTIs; times during which XRT was collecting data) containing no events and lying between light-curve bins to be ignored. This fix made almost no difference to the light curves.

(4) We made two other minor changes in 2008 July, of a cosmetic nature. First, we amended the definition of the actual ‘time’ value

of a bin to be 10 to the power of the mean of the logarithms of the event times within the bin, rather than the linear mean. This makes light-curve plots reflect more accurately the distribution of counts in a bin.<sup>6</sup> Secondly, at the end of an observing snapshot, any events which are not yet sufficient to form a light-curve bin are either appended to the previous bin or carried forward to the next bin, whichever maximizes the fractional exposure. Previously, these events could not be appended to the previous bin if this meant spanning gaps in observations. This change reduces the number of low fractional exposure bins in light curves.

(5) Now that we are producing automatic spectra of GRBs (Section 2.1), the counts-to-flux conversion factor determined from this is automatically used to create a flux-unit light curve. We still use a single conversion for the entire light curve. This conversion factor is taken from the PC mode data unless there are fewer than 200 events in the PC spectrum and more than 200 in the WT spectrum, in which case the WT-mode conversion factor is used.

(6) Four event lists are also now available for download: the source and background event lists in WT and PC mode. As well as containing the events used in the light curve, these contain the GTIs (which are used in the light-curve fitting; Section 2.4.2) and the columns SRCRAD – the radius of the source extraction region, and PUPRAD – the radius of the data excluded to counter pile up. The former is in units of XRT pixels and the latter in arcseconds, as required by the light-curve software. One XRT pixel corresponds to 2.36 arcsec.

(7) Evans et al. (2007) stipulated that, if  $C$  is the number of counts needed to complete a bin in the main light curve, there should be  $2C$  counts in each band of the hardness ratio to complete a bin. This has been relaxed to  $C$  counts in each bin, giving significantly better time resolution in the hardness ratio.

As with the positions, we have also produced a web-based tool to create light curves for non-GRB objects. This tool allows ‘conventional’ binning, i.e. bins of fixed duration, and is described in Section 3.

#### 2.3.2 User-defined data binning

In order to fully automate the light-curve creation, the binning criteria defined in Evans et al. (2007) are applied to all light curves, and these give a useful, valid representation of the XRT data; but not necessarily the ‘best’ representation.

We have therefore produced a web-based tool to allow users to change the binning criteria for a GRB. On the results page for each GRB,<sup>7</sup> there is a link entitled ‘rebin this GRB’. Following this link, the user can specify the minimum number of counts per bin in each XRT mode and whether this is to be used for all bins or is ‘dynamic’ (i.e. it varies with source brightness). On this page, there is a link to the ‘advanced’ rebinning interface, which allows the user to choose which event grades and energy bands are used in light-curve creation. The bands used for the hardness ratio can also be adjusted here.

<sup>6</sup> For example, a bin with four events at  $2 \times 10^4$  s and a single event at  $2 \times 10^6$  s has a ‘time’ of  $3.2 \times 10^4$  s using the new method, but  $4.2 \times 10^5$  s using the standard mean.

<sup>7</sup> [http://www.swift.ac.uk/xrt\\_curves](http://www.swift.ac.uk/xrt_curves)

### 2.3.3 SPER light curves

This software can be easily applied to SPER data to produce light curves within minutes of a trigger. These are of greater reliability than the light curves telemetered via TDRSS, as the SPER light curves are background subtracted and binned with the same method as used on Malindi data (Evans et al. 2007). However, SPER data were originally designed purely for source detection and localization, and thus do not contain GTI information; we are forced therefore to assume that there is no dead-time in the light curve. Since SPER data only cover the first snapshot, this is a safe assumption unless the XRT switches back into WT mode.<sup>8</sup>

## 2.4 Automatic light-curve fitting

We have produced software to automatically fit broken power-law models to the light curves. Note that this is not currently routinely applied to new data. The fitting was performed using the least-squared approach implemented via the MINUIT2 minimization routines produced at CERN,<sup>9</sup> using the  $\chi^2$  statistic. Only light curves containing at least three bins were fitted (i.e. only fits with at least 1 degree of freedom were attempted).

The procedure can be described as a four-step process:

- (i) Identify ‘deviations’ in the light curve.
- (ii) Ignore the times of flares, and fit a series of power laws to the light curve.
- (iii) Use the F-test to determine the best, justifiable fit.
- (iv) Check the results and repeat manually if necessary.

These steps are discussed in detail below. Steps (i)–(iii) are automated, however in Step (iv) a human can decide that the automatic results were incorrect, and repeat those steps manually.

### 2.4.1 Identifying ‘deviations’

For this paper, we aim to characterize the light curves in terms of power-law decays (see Nousek et al. 2006). However, many light curves show deviations from such behaviour which must first be removed. The most common deviations are flares, but other phenomena, such as the slow, curved rise seen in GRB 060218 (Campana et al. 2006), must also be ignored since they cannot be sensibly modelled as power-law decays. To ensure that readers can compare model predictions with the *Swift* data presented here, we list in Table 4 the times which were ignored from the fits. This should not be considered a statistical sample of flares since it includes other phenomena; also the times of the flare are based on when the count rate in the light curve begins to rise (see below) rather than on fitting a flare model to the light curve. For studies of flares in XRT light curves, see, for example, Falcone et al. (2007), Chincarini et al. (2007), Kocevski, Butler & Bloom (2007).

Our automatic script to identify deviations first searches the light curve for any interval where the count rate rises for at least two consecutive bins. If this leads to an increase in count rate of at least  $2\sigma$ , a possible deviation is deemed to have started.

To determine when this potential deviation ends, we look for a shallowing of the decay. We first identify the peak of the potential

**Table 4.** Times identified as deviating from power-law decays, and excluded from the light-curve fits (see Section 2.4.1 for details).

GRB	Times ignored from fit (s since trigger)
GRB 050406	115–430
GRB 050410	45000–100000
GRB 050502B	298–1680, 24800–120000
GRB 050607	223–428
GRB 050712	179–960
GRB 050713A	96–152, 157–203
GRB 050714B	237–486
GRB 050724	220–350, 15831–103776
GRB 050726	0–195, 200–320
GRB 050730	191–791, 72464–119752
GRB 050801	160–700, 3000–6000
GRB 050814	1161–1627, 1813–2718
GRB 050820A	171–260
GRB 050822	115–193, 218–760
GRB 050904	308–600, 1190–1378, 2900–55000
GRB 050908	298–477
GRB 050915A	131–180
GRB 050916	6501–29640
GRB 050922B	612–1517
GRB 051117A	150–200, 300–700, 800–2000
GRB 051227	101–121
GRB 060111A	79–515
GRB 060115	292–564
GRB 060124	250–1200,
GRB 060202	154–163, 35137–52020, 110145–127545
GRB 060204B	104–138, 246–413
GRB 060206	1776–11829
GRB 060210	103–612
GRB 060218	159–3000
GRB 060223A	800–2000
GRB 060312	66–190, 475–817
GRB 060319	198–474, 740–1861
GRB 060413	490–950
GRB 060418	107–215
GRB 060510B	732–1007
GRB 060512	147–380
GRB 060526	208–647
GRB 060602B	120–300
GRB 060604	118–237
GRB 060607A	90–400
GRB 060707	180–400
GRB 060712	250–400
GRB 060714	107–190
GRB 060719	170–400
GRB 060729	165–188
GRB 060814	118–294
GRB 060904A	160–397, 572–1120
GRB 060904B	91–726
GRB 060926	290–1000
GRB 060929	280–1100
GRB 061121	66–80
GRB 061202	126–200
GRB 070107	250–547, 1245–1661, 64099–116566
GRB 070110	30000–80000
GRB 070129	218–1139
GRB 070318	182–429
GRB 070330	142–374
GRB 070419B	87–118, 180–1000
GRB 070517	700–2000
GRB 070518	89–227
GRB 070520B	120–380

<sup>8</sup> This rarely occurs unless the GRB emits a strong flare, in which case this is obvious from the SPER light curve.

<sup>9</sup> <http://project-mathlibs.web.cern.ch/project-mathlibs/sw/Minuit2/html/index.html>

Table 4 – continued

GRB	Times ignored from fit (s since trigger)
GRB 070616	480–550, 714–995
GRB 070621	135–220
GRB 070704	257–658
GRB 070721B	124–141, 187–858
GRB 070724A	90–150
GRB 071021	148–300, 1000–9000
GRB 071031	109–700
GRB 071112C	523–726
GRB 071118	189–225, 321–900
GRB 071122	398–498
GRB 080210	165–379
GRB 080212	169–491
GRB 080229A	100–156
GRB 080310	130–980
GRB 080319D	264–640
GRB 080320	245–500, 650–1000
GRB 080325	158–400
GRB 080506	364–1200
GRB 080604	18000–100000
GRB 080607	116–210, 3000–7000

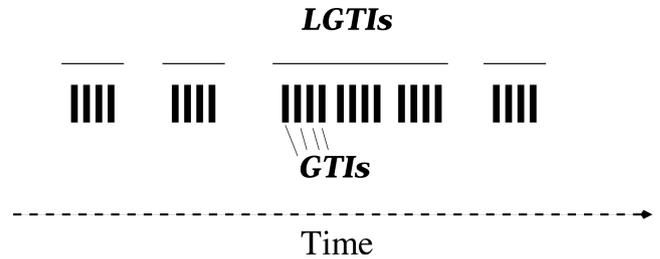
deviation as the highest intensity bin before the count rate systematically drops again. Beyond this peak, the software steps through each bin, calculating two decay indices for each bin. If the current bin is bin  $n$ , these indices are:  $\alpha_1$ : the index of the decay from the peak of the deviation to bin  $n$ , and  $\alpha_2$ : the index of the decay from bin  $n - 1$  to bin  $n$ . A deviation is deemed to have ended if at least one of the following conditions is met: (1) there is a gap of at least 1 ks between light-curve bins, and prior to the gap the time was  $<2$  ks since the trigger (i.e. they cannot begin in the first *Swift* snapshot and extend into the second), or (2) two out of three consecutive bins are found with  $\alpha_2 < \alpha_1$  and both  $\alpha_1$  and  $\alpha_2$  decrease from one bin to the next.

The first condition exists because occasionally a flare starts towards the end of the first *Swift* observing snapshot, and during the gap in observations the flare ends and the light curve enters its final decay; thus the end of the flare using the second test is never found. Chincarini et al. (2007) found that flares tend to have duration: mid-point ratio of  $\sim 0.1$ , so we allow later time flares to span snapshots.

Any time identified above is only confirmed as a genuine deviation, and hence ignored for power-law fitting, if it contains fewer than 10 changes in XRT mode (so is not an artefact of mode switching), contains at least 2 per cent of its component bins before the peak and has a ‘significance’ (defined below) of 1.8. If the potential deviation occurs in the first snapshot (i.e. within 2 ks of the trigger), the ‘significance’ must be at least 3 for it to be confirmed. These numbers were arrived at by trial and improvement to minimize the numbers of false positives and negatives. The significance of a deviation is defined as the peak count rate minus the pre-deviation count rate, divided by the errors in these values added in quadrature.

#### 2.4.2 Fitting power laws

Once the deviations have been identified and removed, we fitted the remaining count-rate light curve with power-law segments separated by zero to five breaks. Note that these breaks are not smoothed. Upper limits were excluded from the fit. To fit the light curves cor-



**Figure 2.** An illustration of the LGTI concept. The vertical rectangles represent individual GTIs. The lines above show the LGTIs – groups of GTIs with less than 30 s of dead-time between consecutive intervals.

rectly is non-trivial because of two interconnected factors: the bins have a finite duration, and many bins have a fractional exposure of less than 1. The normal way to calculate  $\chi^2$  – comparing model and data at the ‘time’ of the bin – does not account for the finite (and sometimes very large) duration of the bin and the evolution of the light curve through this. Instead, one should compare the number of counts detected in each bin to the number predicted by the model. Determining the number of counts predicted by the model in a given bin is non-trivial. The simplest approach (Method 1) is to integrate the model across the light-curve bin. However, if the fractional exposure in the bin is not unity this technique fails, since the model has been integrated over a longer time interval than that during which counts were being collected. Renormalizing the model by the fractional exposure is not a valid solution: this assumes that the source count-rate is the same during the dead-times and live-times within the light-curve bin; for GRBs, which fade, this is clearly untrue.

A better fitting technique (Method 2) is to use the GTI information available with the light curves<sup>10</sup> and to integrate the model across the times during which the XRT was collecting data. Unfortunately, it has not been uncommon for the XRT to ‘mode switch’ (to toggle rapidly between PC and WT modes). This creates a large number of short GTIs. Integrating the model over each of these GTIs dramatically slows down the fit, especially for well-observed bursts, rendering this method impractical. Although recent software and operational changes mean that mode-switching is now a rare occurrence, it has been sufficiently common in the past that to fit all XRT light curves using Method 2 would take several months.

We thus tried a compromise (Method 3), whereby within a light-curve bin we group GTIs together into ‘long Good Time Intervals’ (LGTIs). An LGTI is defined as a cluster of GTIs with less than 30 s of dead-time between consecutive GTIs. This is illustrated in Fig. 2. The model value for a given bin is found by integrating the model across each LGTI, multiplying by the fractional exposure in the LGTI to correct for dead-time and then summing this integration for all LGTIs within the light-curve bin. This method, like Method 1, contains a simplifying assumption; in this case, that during the dead-time within an LGTI the model can be assumed to be constant and at its mean value for the LGTI. This is a much more defensible assumption from Method 3 than that for Method 1: the dead-time within an LGTI is generally very short compared to the duration of the LGTI, whereas the dead-time in a bin can be a substantial fraction of the bin duration. To check whether this assumption affects the fitted models, we chose six GRBs containing bins of non-unity fractional exposure, and modelled their light curves using both Methods 2 and 3. We found that almost all the fitted parameters and

<sup>10</sup> This is contained in the source event lists, available from the *Swift* light-curve repository.

errors agreed to at least three significant figures; where they only agreed to two, the parameter errors were large.

For each light curve, our automatic script first fitted an unbroken power law and then added breaks and refitted with up to a maximum of five breaks (fewer if there was less than one degree of freedom before this). Although break time is a free parameter, it was necessary to estimate the time at which a break would most improve the  $\chi^2$  before adding it, to reduce the likelihood of the fit converging on a local minimum. To achieve this, the software compared the data with the previous model and identified time intervals where the data lie systematically above or below the model. It calculated the  $\chi^2$  contribution from each such interval and added the break at the end of the interval with the greatest  $\chi^2$  contribution, or the start of the interval if it extended to the end of the observation.

#### 2.4.3 F-test

For each broken power-law fit performed in the previous step,  $\chi^2$  is compared to that from the fit with one fewer break, and an F-test used to determine whether the break is significant. We define a break as significant if the F-test returns a probability of the  $\chi^2$  improvement as  $<0.3$  per cent. Note that we do not interpret this quantitatively as confirming the break at the  $3\text{-}\sigma$  level, rather we use it as a convenient means of determining how many breaks to use. Even if the break is not deemed significant in this way, the software still adds a further break; sometimes the ‘true’ best fit requires two breaks, but the  $\chi^2$  improvement from a no-break fit to a one-break fit is not significant. The fit with the most breaks which is deemed significant by the test above is taken as the ‘best’ fit.

#### 2.4.4 Human intervention

For each GRB, we checked the results of the automated steps above. In 23 per cent of cases, flares were misidentified; occasionally genuine flares were missed, but most of the failures were false positives. In these cases, we manually defined the times to be excluded and re-ran the automatic power-law fitting.

In  $\sim 5$  per cent of cases, visual inspection suggested that the use of the F-test had not identified the true best fit. Sometimes this was because one of the fits had found a local minimum of  $\chi^2$  rather than the best fit. In these cases, we manually adjusted the parameters and refitted until the true best fit was found and re-performed the F-test. In other cases ( $\sim 1$  per cent of light curves), the F-test deemed a break necessary only at the 90–99 per cent level, however, knowing that light curves often show a ‘steep-shallow-steep’ behaviour (Nousek et al. 2006; Zhang et al. 2006) allowed us to confirm that the break was genuine – an example of this is given in Fig. 3, along with several examples of good automatic fits.

### 3 NON-GRBS

So far, we have only discussed GRBs, since it is for these that we create products automatically, and present results below. However, *Swift* regularly observes other types of source, with the XRT data often being crucial to the science goals of the observation. We have thus adapted the three tools discussed in Section 2 for use with non-GRB sources, and created a web interface to execute them on-demand. This is available at [http://www.swift.ac.uk/user\\_products](http://www.swift.ac.uk/user_products). In this section, we briefly detail the differences between the ‘non-GRB’ versions of the software and those described in Section 2.

While these tools are public, we refer users to the usage policy at the end of this paper.

#### 3.1 Spectra

The spectrum creation and fitting software for non-GRBs are almost identical to that used for GRBs, except that the user can choose which observations are used to form the spectrum. The user can also specify up to four time intervals rather than creating a single time-averaged spectrum (for GRBs this is possible as well, but only after the average spectrum has been automatically produced). As with GRBs, the spectra are automatically fitted with absorbed power-law models, and the spectral files are provided for download so users can fit other models and interact with the data as required.

#### 3.2 Positions

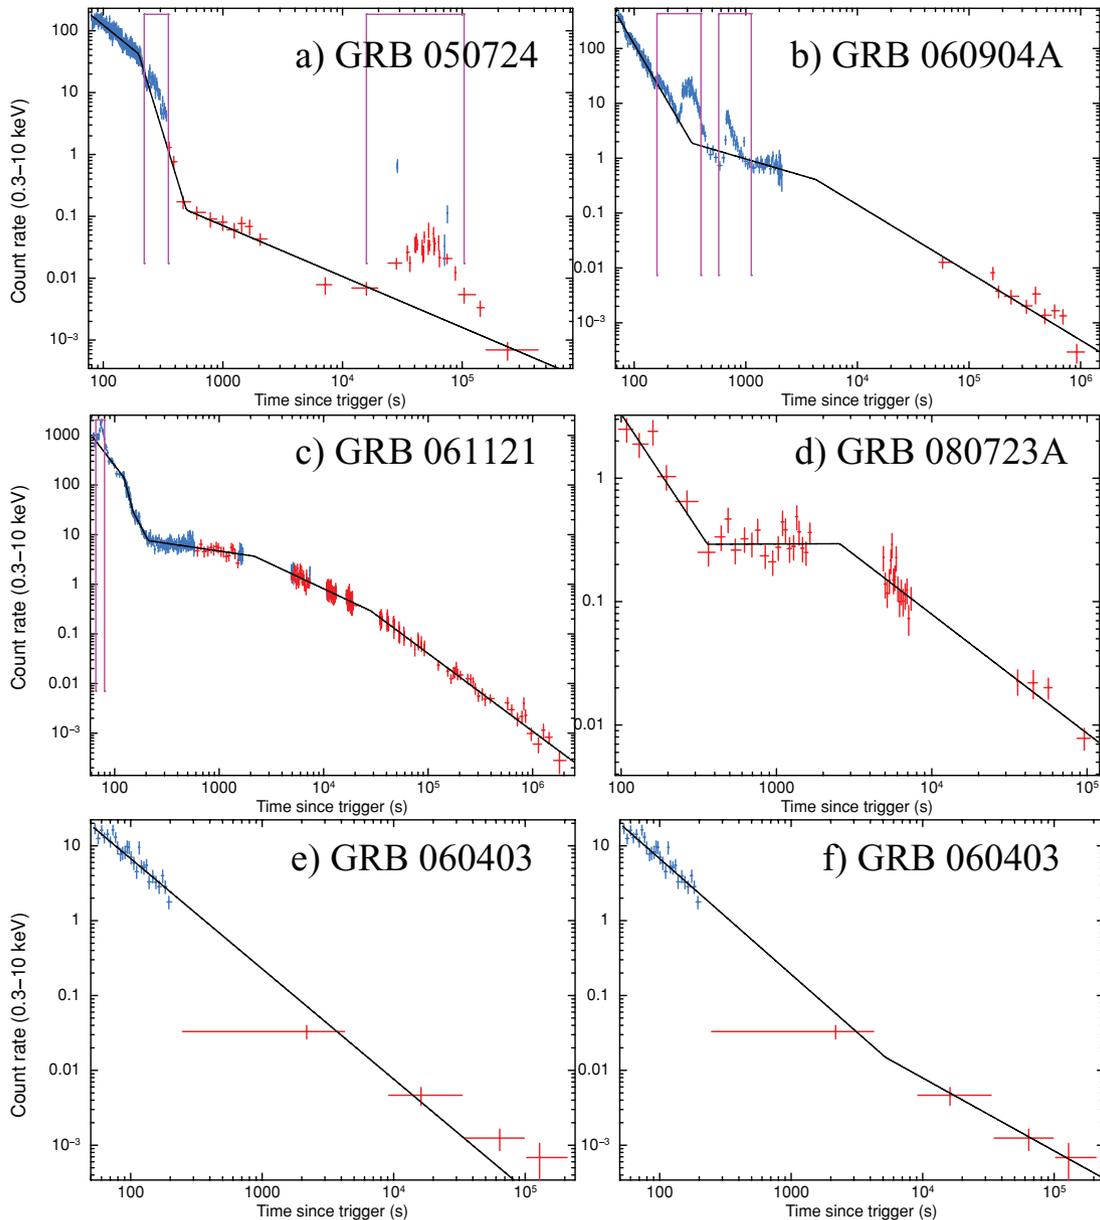
As noted in Section 2.2, if UVOT images taken in the UV filters are used to enhance XRT positions, we tend to find larger errors, thus for GRBs we use only the *v*, *b* and white filters. However, many non-GRB observations are performed using the ‘filter of the day’ (to prolong the life of *Swift*’s filter wheel), which is usually one of the UV filters. Thus, the non-GRB position enhancement tool works in two passes: first it tries to find observations containing PC mode XRT data and UVOT images obtained in the optical filters; if successful, it uses these data to enhance the position. If no such observations can be found, it reverts instead to using the UV filters, and the systematic error is accordingly increased from 1.36 to 1.9 arcsec.

For GRBs, only observations which begin within 12 h of the first one are included; this is because the GRB is generally too faint to detect after this point. Although this is not necessarily true for non-GRB sources, this behaviour is kept by default because using more observations increases the length of time taken to produce the position, and our GRB experience shows that most positions produced with this selection criteria are limited by systematics, not statistics. This can be changed by the user, or the user can explicitly state which observation(s) should be used. In the case of an object which was monitored for some time in quiescence by *Swift* before undergoing some outburst, it is particularly recommended that users specify the ObsID to use, since the brighter data from the outburst would not be included by default but are more likely to give a good position (unless the outburst pushes XRT into WT mode).

#### 3.3 Light curves

For most non-GRB objects, the binning method used for GRBs – defining bins by the number of counts they contain – is not ideal. We have thus produced a version of the software which bins in a conventional way: the user specifies the duration of the bins (the method used for binning GRB light curves is, however, still available for non-GRB sources; alternatively users can choose to produce one bin per snapshot, or per observation). After a gap in the data (e.g. between snapshots), a new bin begins at the start of the next GTI which is not necessarily an integer number of bin-widths after the last bin ended. There are several caveats about this binning method.

(1) The software uses Gaussian statistics to calculate the uncertainty after background subtraction. If there are fewer than 15 counts in a bin, this may not be accurate. It is the user’s responsibility to



**Figure 3.** Examples of automatically fitted light curves. WT data are shown in blue and PC data in red. The purple boxes mark times identified as deviations from power-law behaviour and excluded from the fit. Panels (a)–(d) show cases where no human intervention was needed. Panels (e) and (f) illustrate a case where we overruled the automatic determination of the ‘best’ model. The break added in (f) is only significant at the 96 per cent level according to an F-test, so the automatic software suggested (e) as the best fit. However, given that we know that GRB light curves tend to flatten in the first few kiloseconds, we believe that the automatic fit with a break shown in (f) is the most appropriate fit.

choose a bin size which ensures sufficient counts per bin. A warning is given if any bins contain fewer than 15 counts.

(2) Sometimes the XRT enters WT mode for reasons other than the source being bright. This can produce spurious light-curve points with large error bars (both are invalid), and disrupt the scaling of the plots. By default, any WT mode bin with fewer than 15 counts is assumed to be spurious and is not included in the light curve produced. The web interface allows the user to change the minimum number of WT counts necessary for a valid bin – setting it to zero will include all WT data points.

(3) The last bin in a snapshot may have a low fractional exposure, in which case any statistical fluctuations in the data will be exaggerated. We recommend that users check any such points and consider

rejecting any point with a low fractional exposure; an option to do this automatically is provided by the interface.

When the fixed bin-width binning method is used, an OGIP compliant FITS file is produced containing the light curve, in addition to the standard products created for GRBs.

## 4 RESULTS

In Section 2, we presented details of how enhanced positions, light curves and spectra are produced. These methods have been applied to every GRB detected by the XRT, and will run automatically on all new GRBs. The results are posted online and are available via the following websites:

**Table 5.** Enhanced XRT positions of GRBs observed by *Swift*. The complete table is available in the online version of this paper. The positions are also available online at [http://www.swift.ac.uk/xrt\\_positions](http://www.swift.ac.uk/xrt_positions), which is updated automatically for each new GRB.

GRB	RA (J2000.0)	Dec. (J2000.0)	Error*
GRB 050124	12 51 30.55	+13 02 39.6	1.6
GRB 050128	14 38 17.73	−34 45 55.2	1.7
GRB 050219A	11 05 39.01	−40 41 03.1	2.1
GRB 050219B	05 25 16.05	−57 45 29.3	1.5
GRB 050223	18 05 32.38	−62 28 21.9	4.3
GRB 050315	20 25 54.20	−42 36 01.5	1.5
GRB 050318	03 18 50.99	−46 23 45.0	1.4
GRB 050319	10 16 47.88	+43 32 55.3	1.4

...

\*Arcsec, 90 per cent confidence radius.

- (i) Index: [http://www.swift.ac.uk/xrt\\_products](http://www.swift.ac.uk/xrt_products)
- (ii) Positions: [http://www.swift.ac.uk/xrt\\_positions](http://www.swift.ac.uk/xrt_positions)
- (iii) Light curves: [http://www.swift.ac.uk/xrt\\_curves](http://www.swift.ac.uk/xrt_curves)
- (iv) Spectra: [http://www.swift.ac.uk/xrt\\_spectra](http://www.swift.ac.uk/xrt_spectra)

For SPER data, the results are available via: <http://www.swift.ac.uk/sper>. Each page contains detailed documentation, including a log of any changes made after publication of this paper. All of these pages are interlinked and an index of these results, and those from the BAT, is available via the GCN ground analysis web page at: [http://gcn.gsfc.nasa.gov/swift\\_gnd\\_ana.html](http://gcn.gsfc.nasa.gov/swift_gnd_ana.html).

In Tables 5–8, we list the enhanced positions, best-fitting light-curve parameters and best-fitting spectral fit results for all GRBs observed by the XRT up to GRB 080723B.<sup>11</sup> In Table 6, we also give the *Swift* target ID and the BAT  $T_{90}$  (from the *Swift* Data Table<sup>12</sup>) for reference. Although our automated processing only produces time-averaged spectra, in this compilation of results it is interesting to consider possible spectral evolution. So, for each GRB with a break in its light curve we extracted and fitted spectra for each light-curve segment (delimited by the breaks). These results are presented in Tables 9 and 10.<sup>13</sup> The tables are also available online, through the Virtual Observatory ([ivo://uk.ac.le.star.swift/dsa\\_grb\\_aux/SwiftXRTGRBCat](http://ivo://uk.ac.le.star.swift/dsa_grb_aux/SwiftXRTGRBCat)) and via CDS (<http://cdsweb.u-strasbg.fr/cgi-bin/qcat?J/MNRAS/>).

#### 4.1 Validation of SPER results

The analysis of SPER data is intended to give users an indication of a GRB’s properties extremely rapidly; they are not intended for scientific analysis. Since these data are in every case superceded by Malindi data, we do not list the SPER results in this paper; however we demonstrate their veracity and the limits thereof.

As with the enhanced Malindi positions, we determined the offset between the enhanced SPER position and the UVOT position of every GRB with both of these positions, and confirmed that they agreed 90 per cent of the time, i.e. the enhanced SPER 90 per cent confidence error radius is correctly calibrated.

<sup>11</sup> In order to fit the table within an A4 page, it is necessary to tabulate the decay indices and break times separately. The online data contain these together in a single file.

<sup>12</sup> [http://heasarc.gsfc.nasa.gov/docs/swift/archive/grb\\_table.html/](http://heasarc.gsfc.nasa.gov/docs/swift/archive/grb_table.html/)

<sup>13</sup> We have separated  $\beta$  and  $N_H$  for the paper, however these are in a single table online.

We chose not to fit the SPER light curves as they typically have few bins. Instead, we compared SPER and Malindi curves by eye, an example is given in Fig. 4. We found good agreement between SPER and Malindi light curves.

To test the spectra, we created spectra from Malindi data covering the same time region as the SPER data, and fitted them. We then compared the column density, spectral index and observed flux between these fits and those from the SPER spectra. The first two parameters were in good agreement, however the fluxes only agreed within their 90 per cent errors 70 per cent of the time. This discrepancy probably reflects the lack of GTI information for SPER data, and may also suggest that using the covariance matrix from the fit to estimate the flux errors (as we do) underestimates the errors.

#### 4.2 ‘Malindi’ data

The large volume of uniformly analysed Malindi data presented in Tables 6–10 (i.e. temporal and spectral analyses) allows us to consider the sample properties of the GRBs observed by *Swift*’s XRT.

We have created probability distribution functions (PDFs) for the temporal and spectral indices and temporal break times. We prefer these to histograms as the latter neglect uncertainties on the parameters; a PDF accounts for them and so gives a better representation of the data. The sample PDFs were obtained by treating the PDF of any single parameter as two halves of normal distributions with widths set by the measured uncertainties, each half-normalized so as to form a continuous function, as done by Starling et al. (2008), i.e.

$$P(x|\bar{x}, \sigma_1, \sigma_2) = \frac{\sqrt{2}}{\sqrt{\pi}(\sigma_1 + \sigma_2)} \begin{cases} e^{-(x-\bar{x})^2/2\sigma_1^2} & (x \leq \bar{x}) \\ e^{-(x-\bar{x})^2/2\sigma_2^2} & (x > \bar{x}), \end{cases} \quad (1)$$

where the 1- $\sigma$  errors on each parameter are taken as the calculated 90 per cent confidence error divided by 1.6. We then created overall PDFs of the temporal index ( $\alpha$ ) and spectral index ( $\beta$ ) by summing the PDFs of each individual  $\alpha$  or  $\beta$  parameter and dividing the merged PDF by the number of contributing values. The peak and full width at half-maximum (FWHM) of the various PDFs are given in Table 3; these were calculated by fitting Gaussians to the PDFs; note that many distributions are clearly more complex than a simple Gaussian, in which case the values in Table 3 should be taken as indicative, rather than precise.

The  $\alpha$  PDF is given in Fig. 5, and shows a fairly tight distribution of values. There are a total of 665 values suggesting that the steep drops in probability around  $\alpha = 0.5, 1.5$  are real. This is perhaps not surprising, the ‘canonical’ X-ray light curve (Nousek et al. 2006; Zhang et al. 2006) contains four phases, and as discussed below there are several other light-curve morphologies observed, with one or more distinct phases. Each phase has its own  $\alpha$  distribution corresponding to the peaks in Fig. 5 (see Section 5 for more details).

The PDF of observer-frame break times is given in Fig. 6; we lack the redshift information necessary to translate to the rest frame for most bursts. Since GRB light curves span many decades, the probability density is defined here as probability per unit  $\log(\text{time})$ , rather than per unit time. The two peaks around  $\sim 1\text{--}300$  s and  $10^4$  s arise from the ‘canonical’ light curves and reflect the most common start and end times of the plateau phase; however there is significant probability of a break at all times between  $\sim 100$  and  $10^5$  s. This is likely the result of two effects: the redshift distribution of GRBs, and an intrinsic scatter of GRB light-curve morphology. There are also selection effects which may affect this distribution: to tightly

**Table 6.** Power-law decay indices from light-curve fits. A positive value of  $\alpha$  indicates a decay. The break times between indices are given in Table 7. For bursts where the light curve straddles  $T_0 + 11$  hr, the model flux at 11 h is given (in  $\text{erg cm}^{-2} \text{s}^{-1}$ ). The *Swift* target ID and BAT  $T_{90}$  (taken from the Swift Data Table) are given for reference. The complete table is available in the online version of this paper.

GRB	Target ID	$T_{90}$	$F_{11}$	$\alpha_1$	$\alpha_2$	$\alpha_3$	$\alpha_4$	$\alpha_5$	$\alpha_6$
GRB 041223	00100585	109.1		$1.90^{+0.40}_{-0.39}$					
GRB 050124	00103647	4.0	$1.12\text{e-}12$	$1.50^{+0.15}_{-0.15}$					
GRB 050126	00103780	24.8	$1.22\text{e-}13$	$2.57^{+0.99}_{-0.46}$	$0.94^{+0.24}_{-0.31}$				
GRB 050128	00103906	19.2	$2.25\text{e-}12$	$0.955^{+0.038}_{-0.108}$	$1.361^{+0.075}_{-0.059}$				
GRB 050215B	00106107	8.1	$3.10\text{e-}13$	$0.92^{+0.15}_{-0.13}$					
GRB 050219A	00106415	23.7	$8.73\text{e-}13$	$2.91^{+0.25}_{-0.25}$	$0.754^{+0.089}_{-0.083}$				
GRB 050219B	00106442	30.7	$3.53\text{e-}12$	$1.254^{+0.038}_{-0.037}$					
GRB 050223	00106709	22.5	$1.17\text{e-}13$	$0.89^{+0.27}_{-0.22}$					
GRB 050315	00111063	95.6	$4.24\text{e-}12$	$3.51^{+0.24}_{-0.21}$	$0.257^{+0.034}_{-0.038}$	$1.216^{+0.064}_{-0.063}$			
GRB 050318	00111529	32	$8.25\text{e-}13$	$1.21^{+0.13}_{-0.17}$	$1.86^{+0.15}_{-0.13}$				
GRB 050319	00111622	152.5	$3.26\text{e-}12$	$4.77^{+0.24}_{-1.05}$	$0.493^{+0.048}_{-0.049}$	$0.96^{+0.20}_{-0.12}$		$3^{+26}_{-1}$	
...									

**Table 7.** Times of the breaks in the light-curve fits. The power-law indices are given in Table 6. The complete table is available in the online version of this paper.

GRB	Obs times*	$t_{\text{break}, 1}$	$t_{\text{break}, 2}$	$t_{\text{break}, 3}$	$t_{\text{break}, 4}$	$t_{\text{break}, 5}$
GRB 041223	16.7–28.6 ks					
GRB 050124	11.1–4967.5 ks					
GRB 050126	128 s–93.2 ks	$495^{+540}_{-260}$				
GRB 050128	105 s–99.6 ks	$6492^{+2247}_{-3355}$				
GRB 050215B	5.8–3011.3 ks					
GRB 050219A	112 s–3154.9 ks	$339^{+53}_{-37}$				
GRB 050219B	3.2–3205.2 ks					
GRB 050223	2.9–1047.2 ks					
GRB 050315	83 s–948.3 ks	$464^{+51}_{-47}$	$(5.87^{+0.62}_{-0.63}) \times 10^4$			
GRB 050318	3.3–832.7 ks	$(1.21^{+1.09}_{-0.41}) \times 10^4$				
GRB 050319	234 s–2436.1 ks	$396^{+38}_{-33}$	$(1.51^{+1.10}_{-0.49}) \times 10^4$	$(4^{+3}_{-1}) \times 10^5$		
...						

\*Time range covered by the cleaned event lists. Zero is the BAT trigger time.

constrain a break requires good sampling of the decay on either side of the break, thus towards the end of the light-curve breaks are much harder to constrain and will have broader PDFs, or not be seen at all (Curran et al. 2008a; Racusin et al. 2009). Also, there is usually a gap in *Swift* observations between  $\sim 2$  and 4 ks after the trigger (while *Swift* is the far side of the Earth from the GRB) making it harder to tightly constrain breaks in this interval.

Turning to the spectra, the PDF of the spectral index ( $\beta$ ) is shown in Fig. 7 (panel a). Here, and in Tables 8 and 9, we give the spectral index  $\beta$  (i.e.  $F_\nu \propto E^{-\beta}$ ). Some authors prefer the photon index  $\Gamma$  [i.e.  $N_E(E) \propto E^{-\Gamma}$ ; this is the value used in the XSPEC power-law model]. These are very simply linked:  $\beta = \Gamma - 1$ .

In creating the spectra, we made no attempt to exclude times of flares, preferring to maximize the number of counts in the spectra. Previous studies (e.g. Falcone et al. 2007) and the X-ray hardness ratios on the XRT light-curve repository show that flares tend to have harder spectra than the underlying afterglow emission. Table 4 shows that 81 of the GRBs in our sample contained ‘deviations’ from power-law decays, many of which were flares. To determine whether this has biased our results, we regenerated the spectra

excluding the times of any deviations from power-law behaviour, as identified in the light-curve-fitting phase (Table 4). In panel (b) of Fig. 7, we show the PDF of the spectral index from these data (this includes all the GRBs with no deviations, as well as those where deviations were removed); this is almost identical to that in panel (a). In panel (c), we show the PDF of the change in  $\beta$  caused by ignoring the times of deviations (derived only from the 81 GRBs listed in Table 4). This shows the mean change to be  $0 \pm 0.2$ ; for comparison, the median uncertainty in  $\beta$  in Table 8 is  $\pm 0.16$ . We thus conclude that the presence of flares has a negligible effect on our time-averaged spectra.

Another factor which may affect the spectral index is the redshift ( $z$ ) of the burst; if the redshift used in the fit is incorrect and there is significant absorption in the GRB host galaxy, the absorption will not be correctly modelled – this is most notable around the neutral oxygen edge at 0.525 keV (rest frame). This in turn affects the spectral index. For example, for GRB 050904 ( $z = 6.29$ ; Cusumano et al. 2007), the fit reported in Table 8 has  $\beta = 0.927^{+0.048}_{-0.049}$ . If we set the absorption above the Galactic value to be at  $z = 0$ , we find  $\beta = 0.964^{+0.062}_{-0.060}$ .

**Table 8.** Best-fitting results for time-averaged spectra. Column densities are in units of  $10^{20} \text{ cm}^{-2}$ . WT and PC modes were fitted independently. The complete table is available in the online version of this paper.

GRB	Galactic $N_{\text{H}}$	Redshift	WT mode		PC mode	
			Intrinsic $N_{\text{H}}$	Spectral index ( $\beta$ )	Intrinsic $N_{\text{H}}$	Spectral index ( $\beta$ )
GRB 041218	27.4		$21.6^{+133.6}_{-21.6}$	$2.3^{+5.4}_{-2.0}$	$<56.18$	$0.5^{+1.7}_{-1.1}$
GRB 041223	10.1				$3.1^{+4.5}_{-3.1}$	$1.13^{+0.19}_{-0.16}$
GRB 050124	2.66		$0.96^{+8.00}_{-0.96}$	$0.54^{+0.35}_{-0.24}$	$<3.23$	$0.84^{+0.19}_{-0.12}$
GRB 050126	4.62	1.29			$10.9^{+65.5}_{-10.9}$	$1.15^{+0.64}_{-0.49}$
GRB 050128	5.19		$<125.82$	$0.41^{+1.80}_{-0.45}$	$0.91^{+102.46}_{-0.91}$	$0.40^{+0.91}_{-0.76}$
GRB 050215B	1.85				$<9.90$	$0.73^{+0.61}_{-0.30}$
GRB 050219A	9.50		$17.4^{+4.7}_{-4.4}$	$1.16^{+0.16}_{-0.15}$	$4.8^{+11.7}_{-4.8}$	$0.65^{+0.37}_{-0.32}$
GRB 050219B	2.98		$21.9^{+5.0}_{-4.6}$	$1.09^{+0.18}_{-0.17}$	$12.4^{+4.2}_{-4.9}$	$1.08^{+0.19}_{-0.20}$
GRB 050223	7.08	0.5915			$<16.16$	$0.89^{+0.47}_{-0.42}$
GRB 050315	3.69	1.949	$117.8^{+87.2}_{-64.0}$	$1.11^{+0.35}_{-0.29}$	$88.1^{+20.8}_{-10.5}$	$0.982^{+0.086}_{-0.057}$
GRB 050318	1.86	1.44	$15.3^{+137.9}_{-15.3}$	$1.48^{+3.13}_{-0.71}$	$9.9^{+7.2}_{-3.8}$	$0.989^{+0.074}_{-0.078}$
GRB 050319	1.26	3.24			$9.5^{+23.1}_{-9.5}$	$0.998^{+0.067}_{-0.047}$
...						

**Table 9.** Spectral indices for the time-resolved spectra.  $\beta_1$  corresponds to the time during which the decay followed  $\alpha_1$  in Table 6. Column densities for these spectra are given in Table 10. The complete table is available in the online version of this paper.

GRB	Mode	$\beta_1$	$\beta_2$	$\beta_3$	$\beta_4$	$\beta_5$	$\beta_6$
GRB 050126	PC	$1.29^{+0.20}_{-0.20}$	$1.04^{+0.52}_{-0.36}$				
GRB 050128	WT		$0.35^{+1.74}_{-0.63}$				
	PC	$0.98^{+0.13}_{-0.12}$	$0.966^{+0.088}_{-0.092}$				
GRB 050219A	WT	$1.08^{+0.18}_{-0.17}$	$1.09^{+0.29}_{-0.13}$				
	PC		$1.16^{+0.48}_{-0.53}$				
GRB 050315	WT		$1.09^{+0.34}_{-0.29}$				
	PC	$1.07^{+0.11}_{-0.10}$	$0.979^{+0.085}_{-0.059}$	$1.088^{+0.060}_{-0.104}$			
GRB 050318	WT	$1.51^{+3.51}_{-0.73}$					
	PC	$0.99^{+0.10}_{-0.10}$	$0.99^{+0.12}_{-0.12}$				
GRB 050319	WT			$1.19^{+0.53}_{-0.28}$			
	PC	$1.49^{+0.17}_{-0.16}$	$0.908^{+0.075}_{-0.069}$	$1.09^{+0.11}_{-0.14}$	$14.0^{+8.0}_{-4.7}$		
...							

In our default approach (Section 2.1), unless the redshift of the GRB has been spectroscopically determined, we use an unredshifted absorber to model the excess absorption, as well as the Galactic component. Since this is clearly incorrect, we tried refitting all these GRBs, using  $z = 2.23$ ,<sup>14</sup> for the excess absorption if the redshift was not known. In panel (d) of Fig. 7, we show the PDF of spectral indexes where  $z$  was taken from the literature, where available, and otherwise set to 2.23. Panel (e) shows the PDF of the change in  $\beta$  caused by assuming  $z = 2.23$  instead of  $z = 0$  for those GRBs with no spectroscopic redshift. The latter is  $0 \pm 0.2$ , suggesting that it is acceptable to assume no redshift if no spectroscopic determination has been made.

In Fig. 8, we present the PDF of the ‘excess’ absorption: that is, the value of the second absorption component in the fit (the first

being frozen at the Galactic value from Kalberla et al. 2005; see Section 2.1). The panels are as in Fig. 7.

#### 4.2.1 Time-resolved analysis

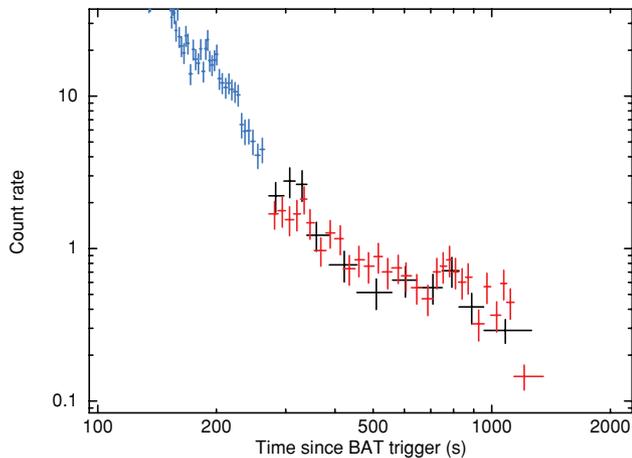
It is interesting to consider how the sample of *Swift* GRBs presented here compare to theoretical predictions. In this paper, we limit ourselves to comparison with the fireball model (e.g. Sari et al. 1998), introduced in Section 1.3. We do this using the ‘closure relationships’ which relate the temporal index ( $\alpha$ ) and the spectral energy index ( $\beta$ ) (see e.g. Zhang et al. 2006 for a table of such relationships). These are simplifications of the complete model; in particular they assume that the electron distribution  $p \sim 2.2$  and that the microphysical parameters such as the proportion of blastwave energy stored in magnetic fields are not evolving through the outburst.

Since the ‘canonical’ afterglow (Nousek et al. 2006; Zhang et al. 2006) contains four distinct phases – high-latitude emission (which

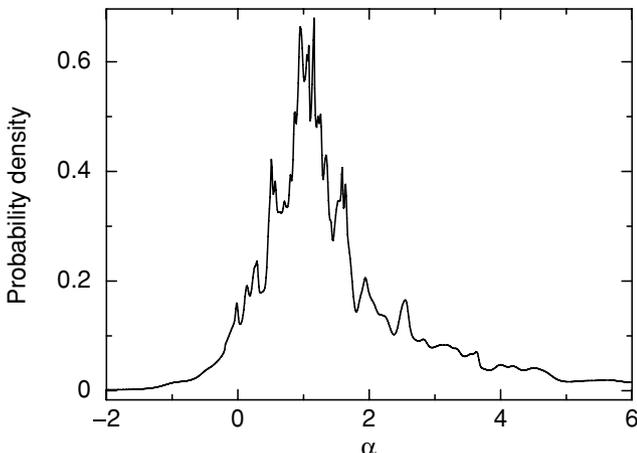
<sup>14</sup> The mean reported redshift for *Swift*-detected GRBs to date, calculated from Table 8.

**Table 10.** Column densities (in excess of the Galactic value) for the time-resolved spectra, in units of  $10^{20}$   $\text{cm}^{-2}$ .  $N_{\text{H},1}$  corresponds to the time during which the decay followed  $\alpha_1$  in Table 6. Spectral indices for these spectra are given in Table 9. The complete table is available in the online version of this paper.

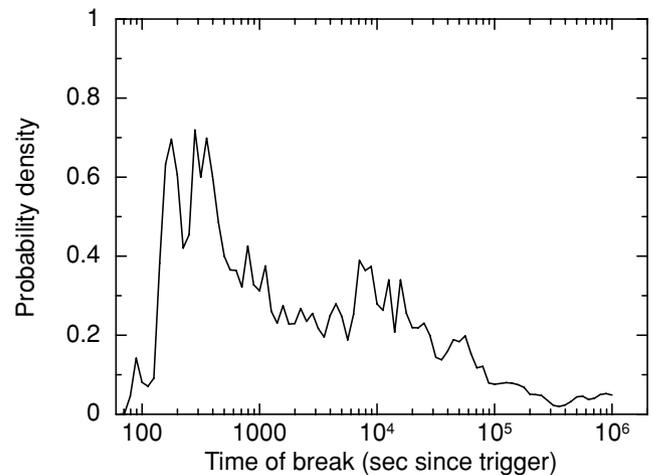
GRB	Mode	$N_{\text{H},1}$	$N_{\text{H},2}$	$N_{\text{H},3}$	$N_{\text{H},4}$	$N_{\text{H},5}$	$N_{\text{H},6}$
GRB 050126	PC	<10.43	$5.5^{+37.1}_{-5.5}$				
GRB 050128	WT		<123.27				
	PC	$6.3^{+3.0}_{-2.9}$	$4.0^{+1.6}_{-1.7}$				
GRB 050219A	WT	$15.3^{+5.2}_{-4.8}$	$15.2^{+8.9}_{-7.5}$				
	PC		$29.9^{+28.2}_{-20.7}$				
GRB 050315	WT		$112.5^{+85.2}_{-65.7}$				
	PC	$35.8^{+18.9}_{-17.8}$	$79.1^{+19.7}_{-17.1}$	$100.0^{+19.9}_{-14.3}$			
GRB 050318	WT	$15.9^{+84.8}_{-15.9}$					
	PC	$11.5^{+10.0}_{-9.6}$	$16.8^{+12.8}_{-10.8}$				
GRB 050319	WT			$97.3^{+167.0}_{-97.3}$			
	PC	<16.50	$18.5^{+28.0}_{-18.5}$	$33.6^{+41.7}_{-23.3}$	$(7.0^{+11.1}_{-2.1}) \times 10^3$		
...							



**Figure 4.** An example SPER light curve (black) plotted with the Malindi data light curve (red – PC and blue – WT). The SPER data clearly give a good representation of the true light curve.

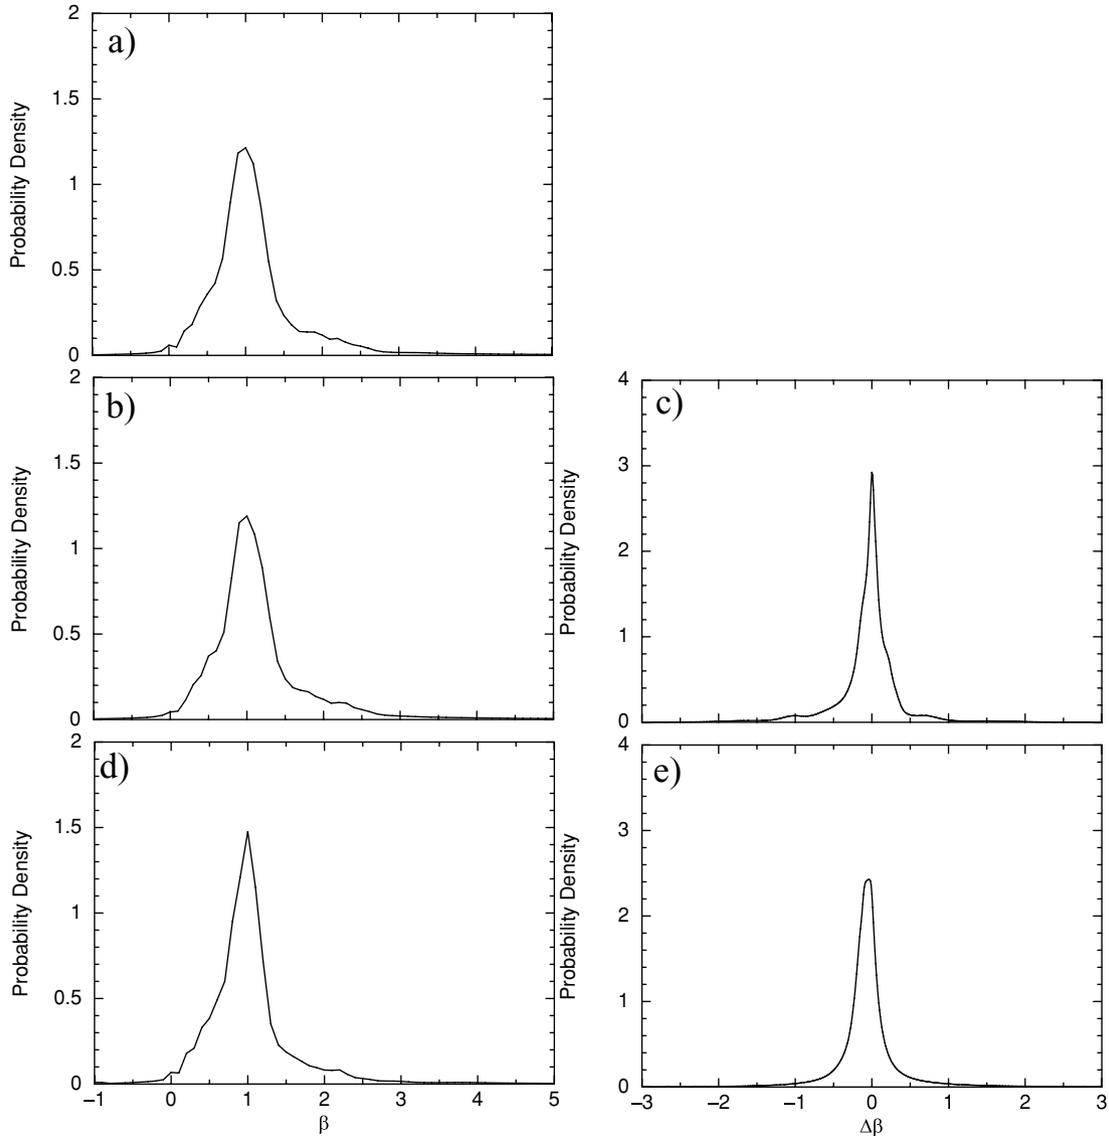


**Figure 5.** PDF of the light-curve power-law decay indices tabulated in Table 6.



**Figure 6.** PDF of the break times [probability per unit log(time)] between the power-law decay segments, tabulated in Table 7.

is not afterglow emission), a plateau, a ‘normal’ decay phase and post-jet-break decay – it is not sensible to compare all of the  $(\alpha, \beta)$  pairs we have derived *en masse*. Instead, we classified each light-curve segment to study the groups separately. To achieve this, we first classified each light curve either as ‘no break’, ‘one break’, ‘canonical’ or ‘oddball’. The first two are self-explanatory. For the latter two: any light curve with at least two breaks was considered canonical if it contained one shallowing break, with  $\Delta\alpha \leq -0.5$ , and a later steepening break, with  $\Delta\alpha \geq 0.5$ , and an oddball otherwise [note that the canonical light curve, as defined by Nousek et al. (2006), does not contain this quantitative definition; we provide it for homogeneity]. We manually checked these classifications and reclassified seven light curves from canonical to ‘oddball’ (e.g. GRB 060202 shows a steep-shallow-steep-shallow behaviour which is not canonical, but meets the criteria defined above). The list of light curves in each class is given in Table 11, and schematic diagrams of all classes except ‘oddball’ (which comprises a range of morphologies and cannot be shown schematically) are given in Fig. 9. The two types of singly broken decay morphologies will be referred to as ‘type b’ and ‘type c’ morphologies hereafter, for brevity.



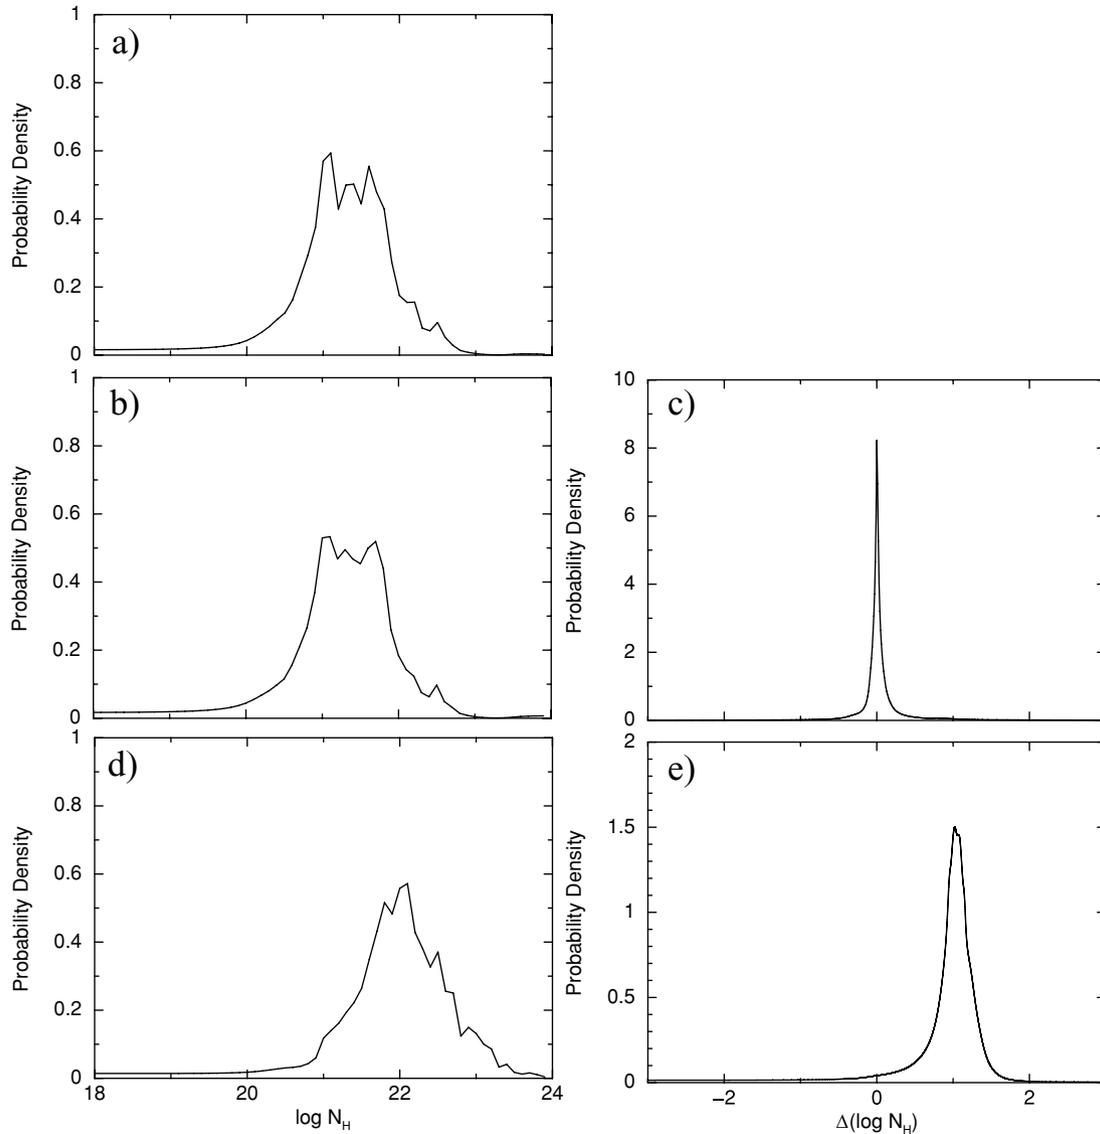
**Figure 7.** PDFs of the spectral index ( $\beta$ ) from the time-averaged spectra. Panel (a) shows the overall distribution. Panel (b) shows this when times of flaring are excluded. Panel (c) shows the PDF of  $\Delta\beta$  caused by removing flares. Panel (d) shows the distribution when bursts with unknown redshift are assumed to be at 2.23, rather than 0, and panel (e) shows the PDF of  $\Delta\beta$  resulting from this change.

For each canonical light curve, we defined any segment with a positive  $\alpha$  (i.e. decaying) before the break with  $\Delta\alpha \leq -0.5$  as belonging to the steep decay phase. Segments after this, but before the break with  $\Delta\alpha > 0.5$ , were identified as the plateau phase. The next segment was identified as the ‘normal’ decay phase and any subsequent decay segments were assumed to be post-jet-break decays. These classifications were again checked by eye, and a small number of segments reclassified accordingly.

In Fig. 10, we plot  $(\alpha, \beta)$  from Tables 6 and 9 for each of these four segments of the canonical afterglow (panels a–d; PDFs of these are given in Fig. 11). We also show the regions covered by the standard afterglow closure relationships (from Zhang & Mészáros 2004); the thick grey band shows the range allowed by the pre-jet-break emission in a slow cooling regime (i.e. the synchrotron peak frequency,  $\nu_m$ , is below the cooling frequency  $\nu_c$ ). The two dark grey lines map the closure relationships for the fast cooling regime ( $\nu_m > \nu_c$ ). In panel (d), we also show the region covered by the post-

jet-break closure relationships (blue band). The points in panels (a) and (b) appear uncorrelated with the closure relationships, however this is unsurprising. In the standard interpretation, the steep decay phase (a) is not afterglow emission, and in the plateau phase (b) energy is being injected into the afterglow. What is perhaps surprising is that many of the ‘normal decay’ phase points (c) do not lie within the range predicted by afterglow theory. The post-jet-break points (d) are the only ones which agree well with afterglow theory, although they agree much better with the pre-jet-break relationships than the post-break ones. We consider these facts in more detail below (Section 6).

For the bursts showing only one break, we show the  $(\alpha, \beta)$  values of the steeper decay in panel (e) of Fig. 10 and the shallower decay in panel (f). The black points are type-b curves and the red ones type c (cf. Fig. 9). Panel (g) shows the bursts with no breaks in their X-ray light curves. Overall, Fig. 10 shows similar results to those in Butler & Kocevski (2007b, fig. 1); however they had a



**Figure 8.** Probability density functions (PDFs) of the ‘excess’ column density, i.e. that above the Galactic value, from the time-average spectra. Panel (a) shows the overall distribution. Panel (b) shows this when times of flaring are excluded. Panel (c) shows the PDF of  $\Delta \log N_{\text{H}}$  caused by removing flares. Panel (d) shows the distribution when bursts with unknown redshift are assumed to be at 2.23, rather than 0, and panel (e) shows the PDF of  $\Delta \log N_{\text{H}}$  resulting from this change.

smaller sample and divided the light curves into segments based on a uniform time-slice rather than individual fits, thus our results are not directly comparable.

### 4.3 Other sample statistics

In the rest of this paper, we will concentrate on the collection of light curves presented in this catalogue. These data can also be combined with other data sets such as the BAT or UVOT catalogues (Sakamoto et al. 2008; Roming et al. 2009) or the online *Swift* Data Table.<sup>15</sup>

For example, Gehrels et al. (2008) compared  $E_{\text{iso}}$  with the X-ray luminosity 11 h after the trigger and found a correlation with a slope of  $\sim 1$ . They thus suggested that the radiative efficiency

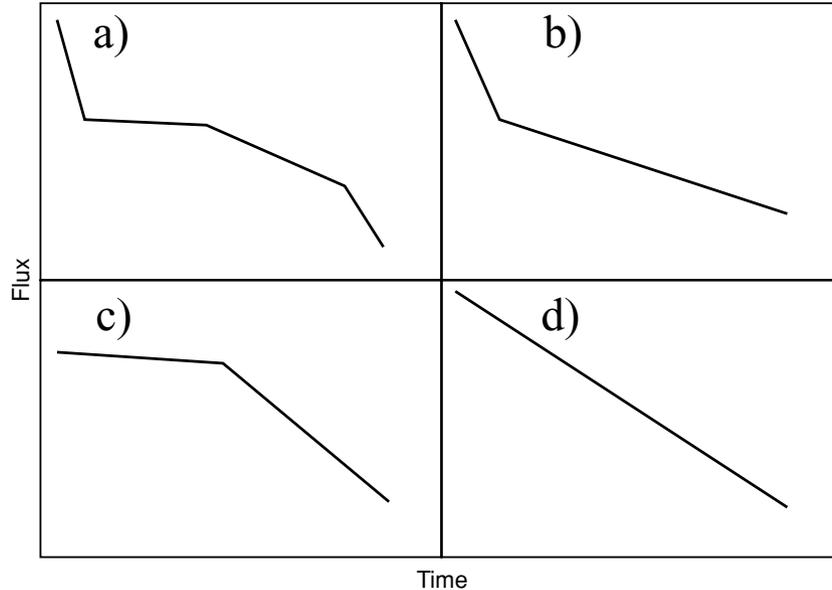
of the blastwave is similar from GRB to GRB. In Fig. 12, we show the BAT fluence plotted against the X-ray flux at 11 h (i.e. we have made no correction for distance, since only  $\sim 30$  per cent of the GRBs in our catalogue have known redshift) using the large sample of GRBs in this paper; the trend can still be seen.

As another example, Willingale et al. (2007) compared the fluence of their two emission components (prompt and afterglow), found a weak correlation and showed that the afterglow fluence never exceeds the prompt fluence. We have performed an analogous analysis, comparing the 15–150 keV BAT fluence from the Swift Data Table with the fluence of the plateau phase in the canonical light curves (Fig. 13). Note that Willingale et al. (2007) determined the fluence over a wide energy band whereas we have only considered the two distinct bands covered by the data (i.e. 0.3–10 keV for XRT and 15–150 keV for BAT). Our results are, however, consistent with those of Willingale et al. (2007).

<sup>15</sup> <http://heasarc.gsfc.nasa.gov/docs/swift/results/>

**Table 11.** Classification of GRB light curves. Bursts in bold are those for which a canonical light curve would be discernable (see text for details).

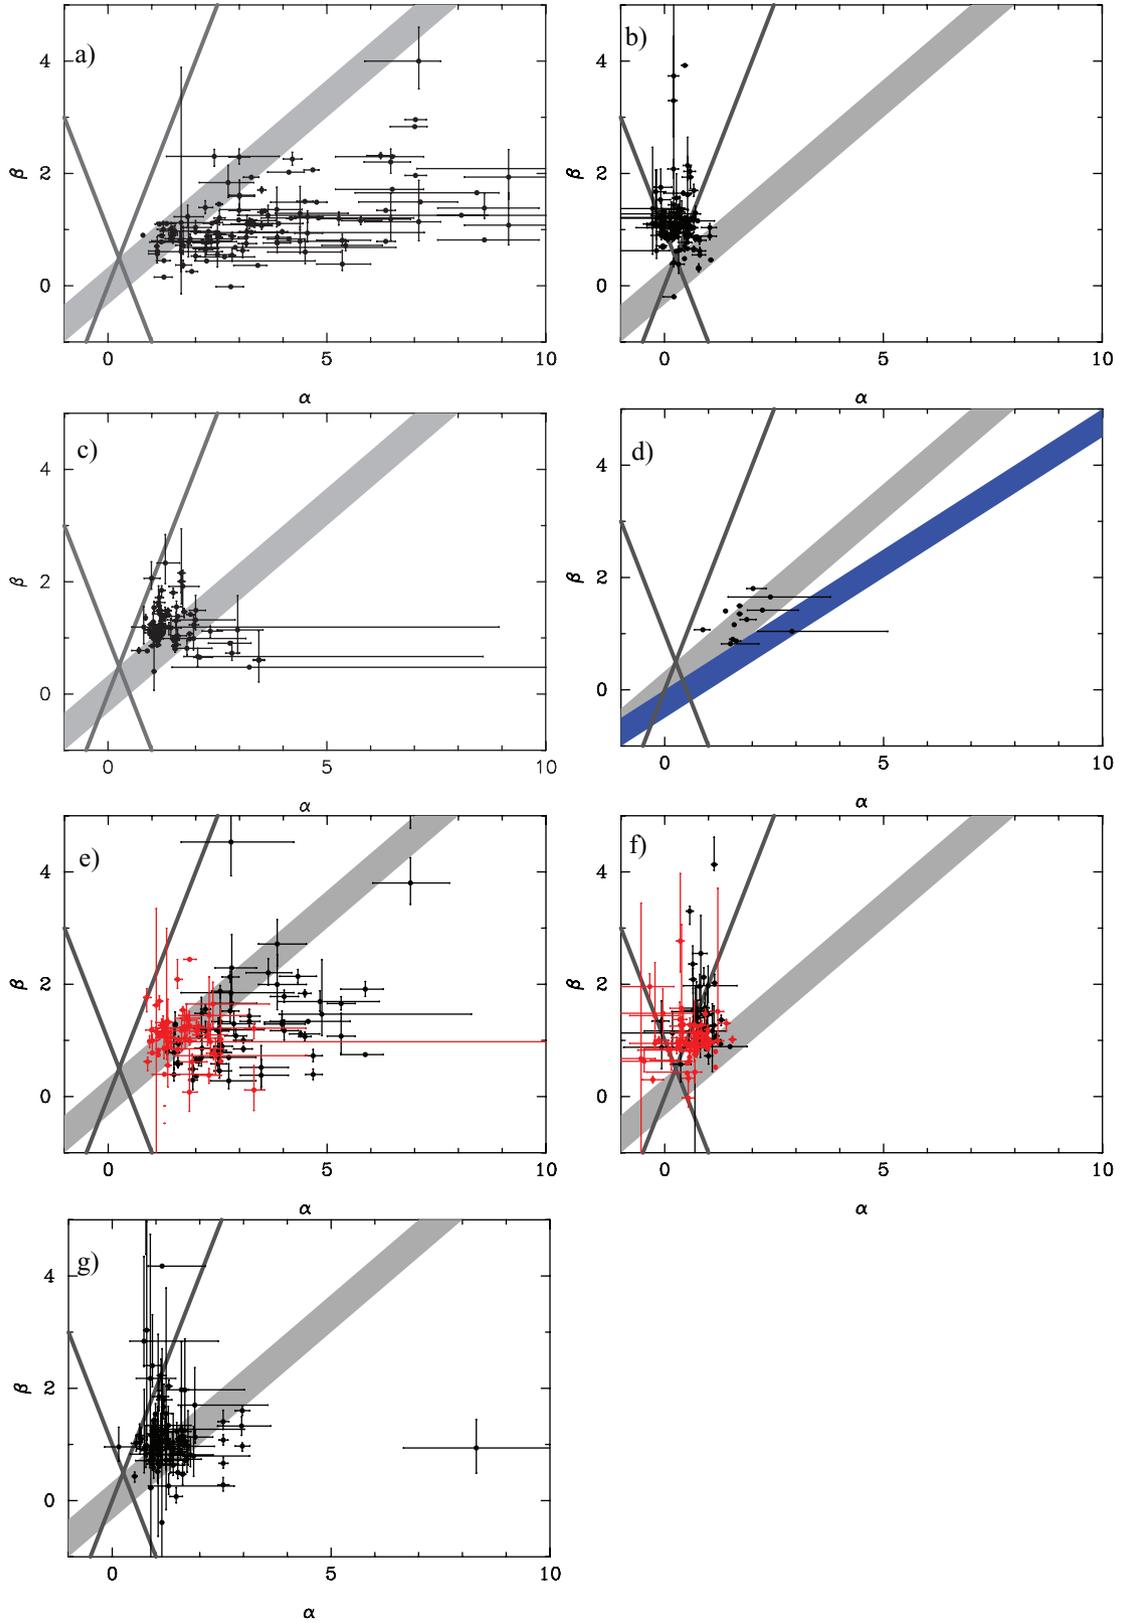
Type of burst	Bursts
No breaks	050408, 050520, 051021A, 051028, 051211B, 060121, 060123, 060805B, 060901, 060928 061025, 061122, 070309, 070724B, 070925, 071104, 080120, 080229B, 080405, 080507 080514B, 080603A, 080613A, 080723B, 041223, 050124, 050215B, 050219B, 050223, 050326 050412, 050509A, 050603, 050701, 050827, <b>050908</b> , <b>050915A</b> , 051006, 051111 051117B, 051210, 051221B, 060110, <b>060116</b> , 060203, 060323, 060421, 060505, 060515 060602A, 060602B, 060717, 060825, 060919, 060923B, 061006, 061110B, 061210, <b>061222B</b> 070227, <b>070318</b> , <b>070330</b> , 070411, 070506, 070509, 070531, 070611, 070612B, 070714A 070809, 070911, 070917, 071003, 071010B, 071028B, 071101, 071117, 071122, <b>080210</b> 080218B, 080319A, 080515, 080520, 080623, 080701, 080702A
One break	051022, 070125, 080625, 050126, 050128, 050219A, 050318, <b>050401</b> , 050406, 050410 050421, 050422, <b>050502B</b> , <b>050525A</b> , <b>050712</b> , <b>050714B</b> , <b>050721</b> , <b>050726</b> , <b>050730</b> , <b>050801</b> 050802, 050815, <b>050819</b> , 050824, <b>050904</b> , <b>050915B</b> , 050916, <b>050922C</b> , 051008, <b>051016A</b> 051021B, <b>051227</b> , <b>060111A</b> , <b>060115</b> , <b>060206</b> , <b>060210</b> , 060211B, <b>060218</b> , 060223A, <b>060312</b> <b>060319</b> , <b>060403</b> , <b>060427</b> , 060428B, 060507, <b>060512</b> , <b>060804</b> , 060805A, <b>060904B</b> , <b>060908</b> <b>060912A</b> , 060927, <b>060929</b> , 061002, 061019, 061028, 061102, 061126, <b>061201</b> , <b>070107</b> <b>070224</b> , <b>070419B</b> , <b>070517</b> , <b>070518</b> , <b>070520A</b> , 070520B, <b>070628</b> , <b>070704</b> , 070721A, <b>070724A</b> <b>070808</b> , 070810A, 071001, 071010A, 071011, <b>071020</b> , <b>071021</b> , 080129, <b>080207</b> , <b>080303</b> <b>080319D</b> , <b>080325</b> , 080409, 080413A, 080426, 080516, 080517, <b>080604</b> , <b>080703</b> , 080710
Canonical	070311, <b>050315</b> , 050319, <b>050416A</b> , <b>050607</b> , <b>050713A</b> , <b>050713B</b> , <b>050803</b> , <b>050814</b> , <b>050820A</b> <b>050822</b> , 050922B, <b>051016B</b> , <b>051109A</b> , <b>051109B</b> , <b>051221A</b> , <b>060105</b> , <b>060108</b> , <b>060109</b> , <b>060204B</b> <b>050826</b> <b>060219</b> , <b>060313</b> , <b>060413</b> , <b>060418</b> , <b>060428A</b> , <b>060502A</b> , <b>060510A</b> , <b>060526</b> , <b>060604</b> , <b>060605</b> <b>060607A</b> , <b>060614</b> , <b>060707</b> , <b>060708</b> , <b>060712</b> , <b>060714</b> , <b>060719</b> , <b>060729</b> , <b>060807</b> , <b>060814</b> <b>060904A</b> , <b>060906</b> , <b>060923A</b> , 060923C, <b>060926</b> , <b>061004</b> , <b>061021</b> , <b>061121</b> , <b>061202</b> , <b>061222A</b> <b>070129</b> , 070219, <b>070220</b> , <b>070306</b> , <b>070328</b> , <b>070420</b> , <b>070429A</b> , <b>070529</b> , <b>070714B</b> , <b>070721B</b> 070802, <b>071028A</b> , <b>071112C</b> , <b>071118</b> , <b>080205</b> , <b>080212</b> , <b>080229A</b> , <b>080310</b> , <b>080320</b> , 080328 <b>080330</b> , <b>080430</b> , 080602, <b>080707</b> , <b>080723A</b>
Oddball	050505, <b>050716</b> , <b>050717</b> , <b>050724</b> , <b>051001</b> , <b>051117A</b> , <b>060111B</b> , <b>060124</b> , <b>060202</b> , <b>060211A</b> <b>060306</b> , <b>060510B</b> , <b>060522</b> , 060801, <b>060813</b> , <b>061007</b> , <b>061110A</b> , <b>070103</b> , <b>070110</b> , <b>070208</b> <b>070223</b> , <b>070412</b> , <b>070419A</b> , <b>070508</b> , <b>070521</b> , <b>070616</b> , <b>070621</b> , <b>071025</b> , <b>071031</b> , <b>071227</b> 080123, <b>080307</b> , <b>080319B</b> , 080319C, 080411, <b>080413B</b> , <b>080503</b> , 080506, <b>080523</b> , 080603B <b>080605</b> , <b>080607</b> , <b>080613B</b> , <b>080714</b> , <b>080721</b>

**Figure 9.** Schematic diagrams of the different light-curve morphologies seen, excluding the ‘oddballs’. Panel (a) shows the so-called ‘canonical’ light curves. Panels (b)–(c) are those with one break, either flattening (b) or steepening (c). Panel (d) are those with no breaks.

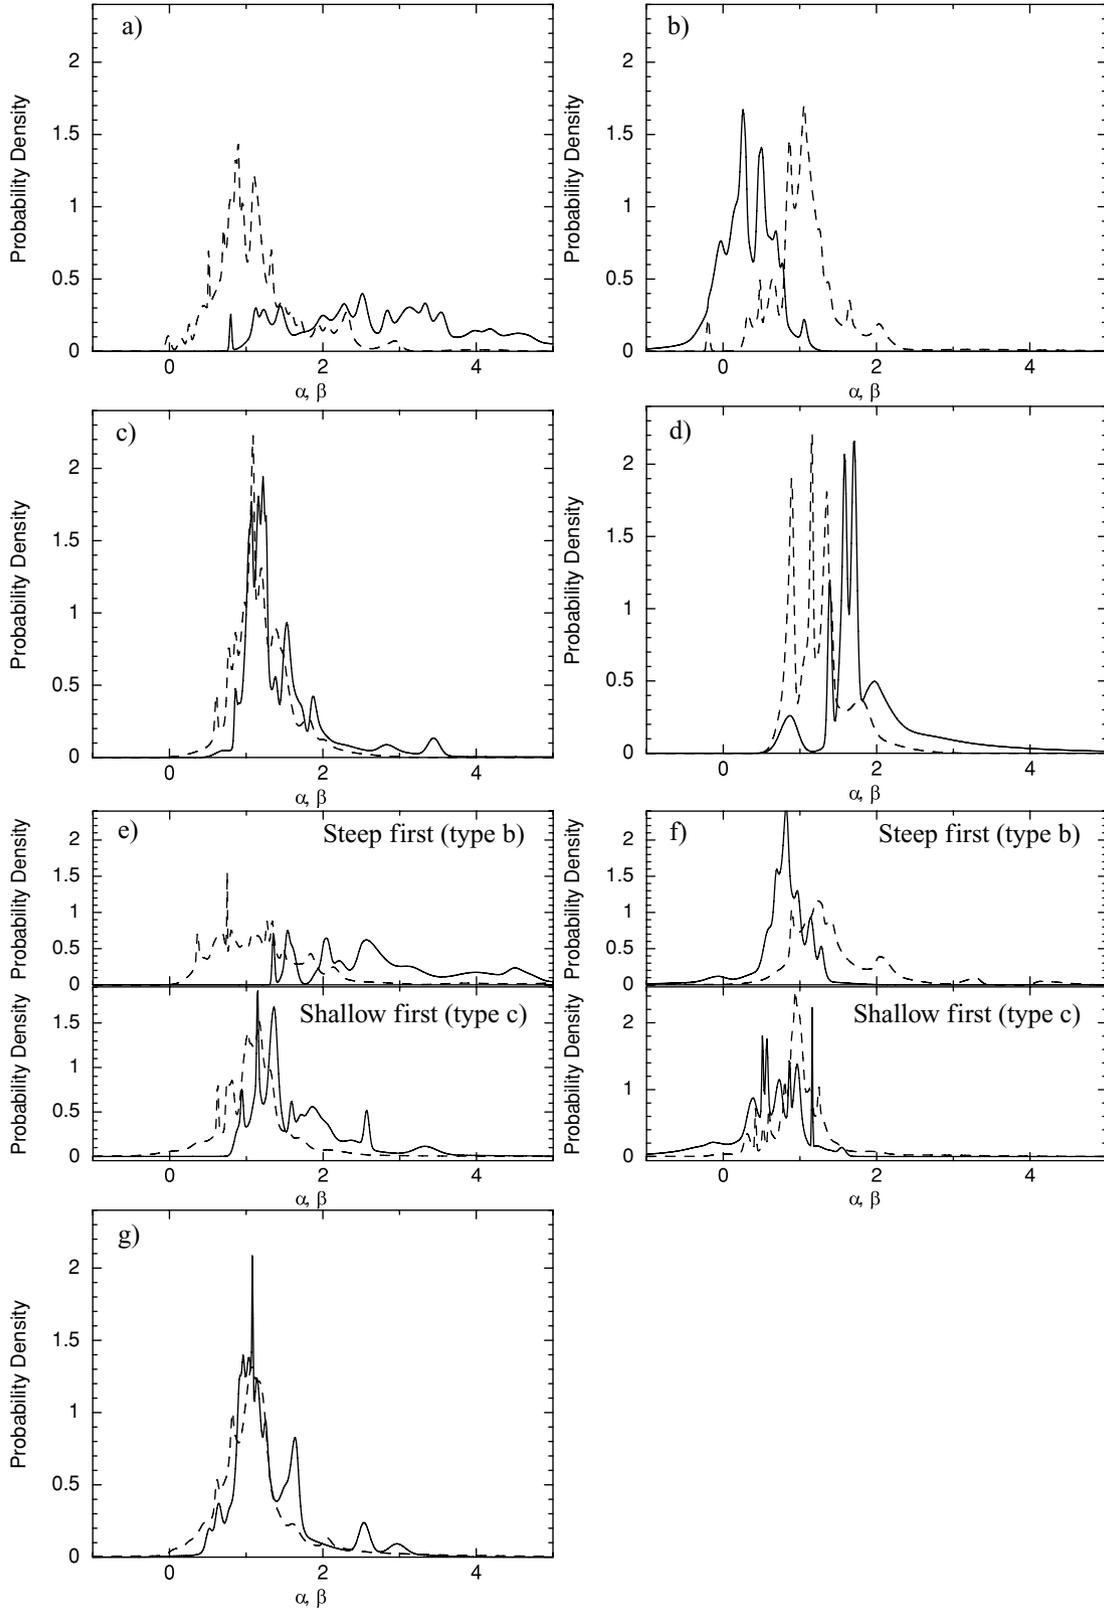
Considering the plateau further if it is caused by energy injection (see below), one may naïvely expect to see some relationship between  $T_{90}$  and the plateau duration, e.g. perhaps longer lived bursts also inject energy for longer. Combining our data with the

*Swift* Data Table, however, reveals no correlation between these two parameters.

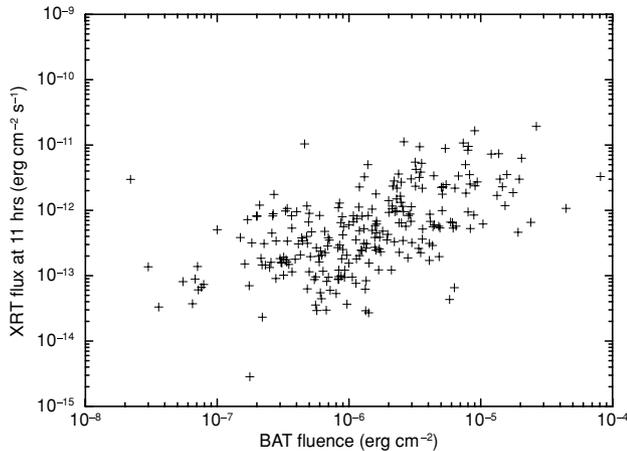
These are a few examples of the large-scale studies which our data set enables; to aid in such studies all of our light curves,



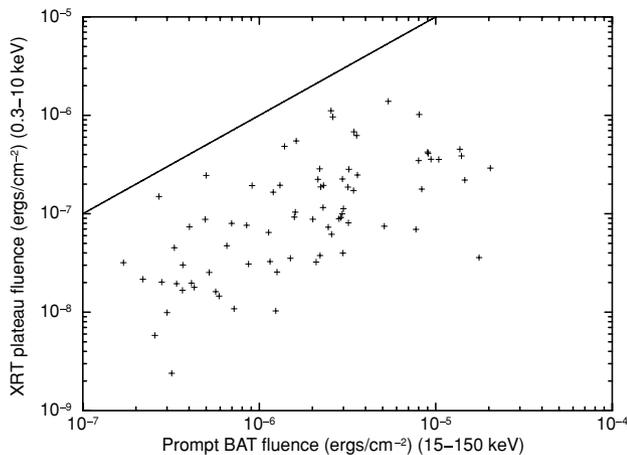
**Figure 10.** Spectral indices ( $\beta$ ) versus temporal indices ( $\alpha$ ) for different light-curve phases (see Fig. 9). Panels (a)–(d) are for the ‘canonical’ light curves and show the values from the steep decay, plateau, ‘normal’ and ‘late break’ phases. Panels (e)–(f) show the values from those light curves showing a single break; the steeper of the segments are plotted in (e) and the shallower in (f). The black and red points indicate types b and c light curves, respectively (see Fig. 9). Panel (g) shows the values for those light curves which do not contain a break. The grey bands mark the areas permitted by standard afterglow closure relationships; the narrow grey lines are for the fast-cooling regime. The blue band in panel (d) marks the range permitted by post-jet-break closure relationships.



**Figure 11.** PDFs of the temporal (solid lines) and spectral (dashed lines) indices ( $\alpha$  and  $\beta$ , respectively) for the different light-curve phases. Panels (a)–(d) are for the ‘canonical’ light curves and show the values from the steep decay, plateau, ‘normal’ and ‘late break’ phases. Panels (e)–(f) show the values from those light curves showing a single break; the steeper of the segments are plotted in (e) and the shallower in (f). For clarity, these plots have been split into two panes to separate the type b and c cases. Panel (g) shows the values for those light curves which do not contain a break.



**Figure 12.** The X-ray flux at 11 h post-trigger plotted against the BAT fluence. The correlation reported by Gehrels et al. (2008) is still present in our larger data set. Note that Gehrels differentiated between long and short bursts (the correlation being more obvious for the former) which we have not done.



**Figure 13.** The fluence of the X-ray plateau phase plotted against the BAT fluence of the prompt emission, for the canonical light curves. The line shows where the two quantities are equal. This plot is analogous to fig. 6 of Willingale et al. (2007) but with a larger sample, and confirms their findings.

positions and spectra are online and the tabulated data in this paper are available in machine readable format; details have already been given above in this section.

## 5 A CANONICAL LIGHT CURVE?

For the rest of this paper, we consider the sample of GRB light curves, and, specifically, the range of morphologies found, as demonstrated in Fig. 9. Nousek et al. (2006), Zhang et al. (2006) and Panaitescu et al. (2006) proposed that there is a ‘canonical’ X-ray GRB light curve, consisting of four power-law phases: a steep initial decay, a shallow plateau and then a ‘normal’ decay which is steeper than the plateau, but not as steep as the first segment. There may also be a fourth segment, post-jet-break decay. Panels (c) and (d) of Fig. 3 show light curves which conform to this behaviour. Nousek et al. noted that we do not see this exact behaviour in all GRB afterglows, and suggested that this is simply due to limited temporal coverage. O’Brien et al. (2006) meanwhile

showed that, for GRBs observed by *Swift*, the prompt emission seen by the BAT transitions smoothly into the emission seen by the XRT.

Willingale et al. (2007) interpreted the observed X-ray emission as the combination of two components, each following a simple exponential-to-power-law form. A late-time break in the power law may be added occasionally as well. Physically, the two exponential-to-power-law components were identified with the prompt GRB emission from internal shocks in the ejecta, and afterglow emission from an external shock in the circumburst medium. Under this model, not all GRBs exhibit all segments of the ‘canonical’ curve. For example, the afterglow component can be sufficiently weak compared to the prompt component that it is never seen, alternatively it can dominate from an early time.

In this paper, we have presented a sample of GRBs much bigger than those used by Nousek et al. or Willingale et al. (who used 27 and 107 bursts, respectively), and can thus consider the possibility of a unified afterglow model with more confidence.

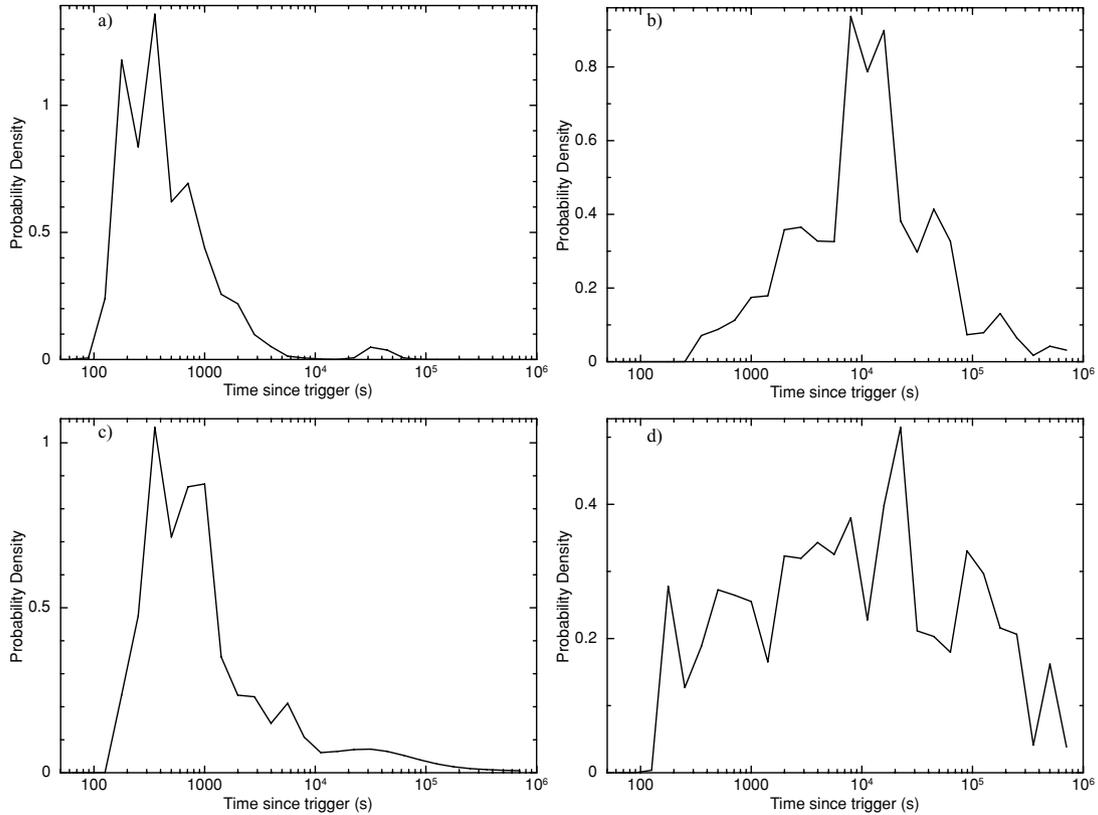
We defined a subset of the bursts presented in this paper, comprising only those for which we can reasonably expect to have seen the three phases of the canonical light curve. By inspecting the break times (Table 7) of the canonical light curves (Table 11), we defined such bursts as having XRT data beginning at  $T \leq T_0 + 200$  s and extending to at  $T \geq T_0 + 50000$  s. We also specified that the light curve should contain at least 20 bins. This gave a sample of 162 GRBs, which are shown in bold type in Table 11 [Note that previous studies (e.g. O’Brien et al. 2006; Willingale et al. 2007) did not create such subsets; their samples are analogous to our complete sample, not this subset]. Of these 162 bursts:

- (i) 7 (4 per cent) have no breaks.
- (ii) 49 (30 per cent) have one break (25 type b, 24 type c).
- (iii) 68 (42 per cent) are canonical.
- (iv) 38 (24 per cent) are oddballs.

This immediately shows that the ‘canonical’ light curve, while the most common morphology, actually occurs in less than half the GRBs in which it would be identifiable. Whether or not the underlying afterglow behaviour follows a single behaviour is readily testable with our data set, and we discuss the different light-curve morphologies in this context below. To aid this discussion, we show in Fig. 11 the  $\alpha$  and  $\beta$  PDFs for the various light-curve phases (the panels are as in Fig. 10), and in Fig. 14 PDFs of the break times for the different light-curve morphologies. These figures were built using all 318 GRBs presented in this paper, not just the subset defined above. Immediately, we see from this that the  $\beta$  values are roughly the same in each light-curve segment, although the steep-decay phase of the canonical curves and the steep portion of the one break, steep-to-shallow light curves have a wider range of values than the other phases (a similar conclusion was drawn by Butler & Kocevski 2007b). Note that we do not consider the ‘oddball’ bursts here since these need to be studied individually, whereas we are interested in the bulk properties of afterglows. Schematics of the morphologies discussed below are given in Fig. 9.

### 5.1 Light curves with no breaks

Although we have only seven GRBs with no breaks (Fig. 9, panel d) in our subsample, the distribution of  $\alpha$  values from these GRBs is inconsistent with the distribution from any phase of the canonical GRBs in Table 11 except for the normal decay phase. Even this consistency is low; Kolmogorov–Smirnov (K–S) test gives a 1.9 per cent probability that the  $\alpha$  values of these eight GRBs were drawn from the same sample as the normal decay phase of the canonical



**Figure 14.** PDFs of the break times for the canonical light curve and curves with one break. Panels (a) and (b) give the start and end times, respectively, of the plateau phase of canonical light curve. Panel (c) shows the break times for light curves which show a single, flattening break (type b), and panel (d) gives the break times for those which show a single, steepening break (type c).

GRBs (note that the K–S test is not necessarily believable with such a low number of values). This suggests that the no-break GRBs are consistent with the Willingale et al. (2007) model provided we are seeing only the power-law phase of the afterglow component, however the lack of a steep decay phase means that either the prompt emission must decay very rapidly or the afterglow must be bright enough to dominate from a very early time. Further, the lack of a plateau phase implies that energy injection does not dominate at any time; the GRBs in this subset (panel g of Fig. 10) show reasonable agreement with standard afterglow theory, supporting this idea. The outlier in that panel with  $\alpha \sim 8$  is GRB 051221B, and is a candidate ‘naked’ GRB (Willingale et al. 2007).

## 5.2 Light curves with one break: type b (shallowing decays)

Under the Nousek/Willingale models, type b light curves should correspond to the first two segments of a ‘canonical’ light curve, with the plateau phase ongoing when observations cease. In terms of Figs 10–11, this means that the black points (upper pane in Fig. 11) in panel (e) should come from the same parent population as those in panel (a), and those in panel (f) from the same population as those in panel (b). By eye, the first of these statements seems believable, and a K–S test gives a 17 per cent probability that the  $\alpha$  values of the two samples came from the same parent population. However, the  $\alpha$  values of the shallow decay in these light curves and the plateau phase of the canonical ones are completely different. From Fig. 11, one can see that the distribution in the upper pane of panel (f) ( $\bar{\alpha} = 0.85$ ) lies towards significantly higher  $\alpha$  than those in panel (b) ( $\bar{\alpha} = 0.34$ ); a K–S test gives a  $<0.1$  per cent chance that

the two come from the same parent population. Further, although the distribution of plateau start times (Fig. 14, panel a) looks similar to that of the type b break times (Fig. 14, panel c), the latter are shifted towards later times; a K–S test gives a 0.3 per cent chance that these share a common population.

It is still possible to reconcile the bursts with a single, steep to shallow break to the same behaviour as the canonical bursts, if the shallow phase is similar to the plateau phase, but the energy injection in these bursts is longer lived and at a lower rate than in the ‘canonical’ bursts. A rigorous investigation of this is beyond the scope of this paper, and will be tackled in a future publication. We do note, however, that if this is true, energy injection must continue at least to the end of the *Swift* observations, which in all but three of these cases is more than a day (often many days) post-trigger. Producing such long-lived energy injection at the necessary level, from the standard GRB progenitor models, is difficult; however X-ray flares have been seen  $>1$  day after the trigger (e.g. GRB 050502B, Falcone et al. 2006; GRB 080810, Page et al. in preparation; see Curran et al. 2008b for a discussion of late-time X-ray flares), implying that the central engine can still affect the afterglow on these time-scales.

## 5.3 Light curves with one break: type c (steepening decays)

Compared to the canonical light curve, type c light curves could correspond to the normal and post-jet-break phases of a GRB light curve. From an  $(\alpha, \beta)$  point of view, this is acceptable; K–S tests show  $>1$  per cent probability that the red points (lower pane) in panel (e) of Figs 10–11 come from the same population as those in

panel (c), and those in panel (f) come from the same population as panel (d). However, for this to be the case 7/25 (=28 per cent) of the ‘jet-breaks’ would have to occur within 1000 s of the GRB trigger, suggesting an extremely confined jet. Alternatively, the ‘jet-breaks’ in the canonical light curves may not be jet-breaks at all; this is suggested by panel (d) of Fig. 10, and we discuss this further in Section 6.

Instead of the above, the two phases of these GRBs could be identified with the plateau and normal decay phases of the canonical light curve. The distribution of plateau end times (Fig. 14, panel b) is similar to the (poorly sampled) distribution of type-c break times (Fig. 14, panel d), and a K–S test gives a 40 per cent probability that these represent the same population of times. However, a K–S test between the decay slopes of the shallow part of the type c light curves and the plateaux of canonical bursts gives a <0.1 per cent probability that these come from the same population. This does not definitively rule out this interpretation: if the afterglow dominates the X-ray light curve before the prompt component decays, the effect of energy injection may be less than in a canonical GRB, giving a steeper shallow decay slope, as seen in our data. To investigate, we obtained the BAT fluence from the Swift Data Table<sup>16</sup> for all of the canonical GRBs in Table 11, and for the one break, shallow-to-steep GRBs from our subsample. If the latter bursts have systematically lower fluence than the canonical GRBs, the above explanation holds. No such trend is seen, however.

## 6 UNDERSTANDING THE X-RAY AFTERGLOW

We have shown above that the different morphologies of GRB light curves are consistent with the two component model of Willingale et al. (2007); implying a consistent underlying behaviour (if not a canonical shape). We now consider what physical processes drive each of the phases obtainable from such a light curve. The large, homogeneously generated, data set in this paper is an ideal test bed for this. The usual explanation of the phases (e.g. Nousek et al. 2006; O’Brien et al. 2006; Zhang et al. 2006) is as follows.

- (i) Steep decay – high latitude prompt emission (internal shocks).
- (ii) Plateau – emission from a collimated external forward shock (afterglow) which is undergoing energy injection. The edge of the jet is not visible to the observer.
- (iii) Normal decay – emission from a collimated external forward shock with no energy injection. The edge of the jet is not visible to the observer.
- (iv) Post-jet-break – emission from a collimated external forward shock with no energy injection. The edge of the jet is visible to the observer.

To compare our data with theoretical predictions for the steep decay phase requires modelling of the BAT data, since  $\alpha$  in this regime is sensitive to  $T_0$ , which should be taken as the start time of the final pulse. This is beyond the scope of our XRT-data paper, however many other authors have confirmed that the steep decay phase is consistent with the expectation for high latitude emission (e.g. Tagliaferri et al. 2005; Barthelmy et al. 2005b; O’Brien et al. 2006; Goad et al. 2006; Liang et al. 2006; Willingale et al. 2007; Butler & Kocevski 2007a).

The plateau phase likewise is in good agreement with the above model. Zhang et al. (2006) give the closure relationships for the energy injection scenario, assuming the luminosity of the injecting

source  $L \propto t^{-q}$ , where  $q \leq 1$ . The lower edge of the solid grey band plotted in Fig. 10 corresponds to the most tolerant  $q = 1$  limit from such relationships [as does the blue band in panel (d) for the post-jet-break, energy injection relationships taken from Panaitescu et al. (2006)]. Points lying above this line are consistent with energy injection afterglow theory.

The normal decay phase is not in good agreement with the interpretation above. Panel (c) of Fig. 10 shows many of the points to be inconsistent with afterglow theory without energy injection. This disagreement continues into the post-jet-break phase (panel d), where the majority of points are consistent with the standard pre-jet-break models, but very few are consistent with post-jet-break theory. This suggests that the interpretation of the X-ray light curve as given above is incorrect unless energy injection continues for some time after the burst. (Note that we have not given an exhaustive study of jet-breaks as we include only breaks which occur after the first three ‘canonical’ phases. For a targeted study of potential jet-breaks in any light-curve morphology, see Racusin et al. 2009.)

Other models have been proposed instead to explain the observed X-ray light curves. For example, Ghisellini et al. (2007) suggest that the long-term X-ray emission is actually ‘late prompt’ emission, from internal shocks with lower bulk Lorentz factors than in the initial case. They use this to model X-ray and optical light curves with some success (Ghisellini et al. 2009). However, in such a model we may expect to see spectral evolution in the late–prompt emission analogous to that seen during the prompt emission; as Table 9, Fig. 11 and the hardness ratios on the XRT light-curve repository (Evans et al. 2007) show, there is very little spectral evolution seen in XRT data after the first few hundred seconds post-trigger (O’Brien et al., in preparation). Note that the Ghisellini model implicitly assumes that the late–prompt component does not spectrally evolve, so does fit the observed hardness ratios, however why it does not evolve is not clear. The dust-scattering model of Shao & Dai (2007) suffers from the same problem (Shen et al. 2009). de Pasquale et al. (2009) recently suggested that the end of the X-ray plateau could signify a jet-break, where energy injection is ongoing, subsequent ‘jet-breaks’ would then signify the end of energy injection. However, this is not consistent with the points in panel (d) of Fig. 10 – data taken after the end of the plateau and a subsequent break – which are generally inconsistent with post-jet-break models with no energy injection.

The data set presented in this paper represents the best diagnostic tool for afterglow models currently available, and can be used to place specific constraints on any given model for the X-ray emission. For example, considering the external forward shock model, Fig. 10 tells us following points.

- (1) During the plateau phase, energy must be injected into the shock.
- (2) The so-called ‘post-jet-break’ phase in the ‘canonical’ light curve is in fact better explained as occurring before the jet-break but after the cessation of energy injection than by the standard interpretation of arising after the jet-break and cessation of energy injection.
- (3) Some mechanism must cause a steepening of the light curve, independent of energy injection. It must not invoke any spectral change.

The latter point arises because the break seen between the plateau and normal phases cannot always be caused by the cessation of energy injection: too many points in panel (c) lie above the grey band and hence must be undergoing energy injection.

<sup>16</sup> Via <http://heasarc.gsfc.nasa.gov/docs/swift/results/>

There is a reasonable number of bursts whose normal decay phase is consistent with a standard forward shock with no energy injection [i.e. points in Fig. 10 panel (c) which lie within the grey band] as well as many which do require energy injection during this phase. Thus, in the above description, it must be possible for the unknown-origin break (in point 3 above) to occur before, or after, the cessation of energy injection. Before this break while energy injection is ongoing, a GRB lies on panel (b) of Fig. 10, after the break and once energy injection has ceased it lies on panel (d). Whether the cessation of energy injection or the unknown-origin break occurs first would then determine whether the GRB lies in or above the grey band during its time on panel (c).

There is a significant number of bursts lying above the grey band permitted by the closure relationships in Panel (c) of Fig. 10, which do not show a subsequent break. Similarly, the black points in panel (f) represent the last observed state for many GRBs. This implies that for external, forward shock model of the afterglow, significant energy injection must last for days, if not weeks, after the trigger. Ghisellini et al. (2009) suggest that this is possible, however it is not clear that their mechanism can produce sufficient levels of energy injection to sustain the shallow decay. None the less, if energy injection from lasting central engine activity (or slow-moving shells ejected at the time of the burst) is responsible for the shallower-than-expected decay, we may expect bursts whose prompt emission is relatively faint compared to the afterglow emission to show little evidence of energy injection (unless the central engine gets brighter with time!). The red points in panels (e)–(f) of Fig. 10, and the points in panel (g) are such bursts: their afterglows show no steep-decay phase, which (see Section 5) may mean that from an early time, the afterglow dominated any prompt emission. As can be seen, the majority of these are consistent with having no energy injection, supporting this model.

The above discussion does not tell us that the forward shock model for X-ray afterglows is the correct model for XRT afterglow emission, however it demonstrated that, with a little reorganising in light of the constraints placed by our data set, it is still consistent with observations. None the less, two difficulties remain: some mechanism must be found to produce a spectrally invariant temporal break with a wide range of  $\Delta\alpha$ ; and it must be possible to inject a significant amount of energy into the external shock for days to weeks after the explosion.

## 7 CONCLUSIONS

We have developed software to automatically produce light curves and hardness ratios, spectra and high-precision enhanced XRT positions of GRBs. Preliminary versions of these are available within minutes of a trigger, and the full versions are available within a few hours. Users can interact with and customize these products as desired. We also provide general-purpose versions of these tools to run for any object observed by XRT, available via a web interface.

These products are available online:

- (i) Index: [http://www.swift.ac.uk/xrt\\_products](http://www.swift.ac.uk/xrt_products)
- (ii) Positions: [http://www.swift.ac.uk/xrt\\_positions](http://www.swift.ac.uk/xrt_positions)
- (iii) Light curves: [http://www.swift.ac.uk/xrt\\_curves](http://www.swift.ac.uk/xrt_curves)
- (iv) Spectra: [http://www.swift.ac.uk/xrt\\_spectra](http://www.swift.ac.uk/xrt_spectra)

Using this software, we have performed a homogeneous analysis of all GRBs observed by the XRT to date, and presented positions and temporal and spectral indices, in various formats. An analysis of these data shows that a variety of light-curve morphologies exist, and the so-called ‘canonical’ curve, while the most common case,

accounts for less than half of the light curves seen by *Swift*. Defining a subsample of 162 GRBs with sufficient coverage to detect the canonical shape, if it existed, we found

- (i) 7 (4 per cent) have no breaks,
- (ii) 49 (30 per cent) have one break (25 shallow, 24 steepen),
- (iii) 67 (41 per cent) are canonical,
- (iv) 38 (24 per cent) are oddballs.

We have, however, demonstrated that this range of morphologies can be explained by a single underlying behaviour; the two-component model suggested by Willingale et al. (2007), which involves a ‘prompt’ component and an ‘afterglow’ component. To achieve this, we require a range of prompt-to-afterglow emission ratios, and a range of energy injection rates, both of which are easy to accept given the variations seen from burst to burst.

If the afterglow emission is due to the external forward shock model, then in many cases this scenario can only explain the data if energy injection continues beyond the plateau phase, and lasts for days to weeks after the GRB. The data also require a mechanism which can cause a light-curve break (i.e. the end of the plateau) without terminating energy injection, and without causing a change in the X-ray spectrum.

### 7.1 Usage policy

Anybody is welcome to use the products and tool details in this paper for their work. Although we have verified these tools as far as possible, we still strongly advise users to ‘sanity check’ their results, particularly with regard to light-curve binning (Section 3.3).

If these products or tools are used in any publication, we ask that this paper be cited, and that users include the following statement in the acknowledgements:

‘This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester’.

## ACKNOWLEDGMENTS

PAE, APB, KLP, JPO, RLCS, OG, KW and CJM acknowledge STFC support. DNB, LV, JR, JAK and CP are supported by NASA contract NAS5-00136. MC, GC, RM, MP and PR are funded by grants SWIFT I/011/07/0, PRIN MIUR 2007TNYZXL, MAE.

## REFERENCES

- Abbey A. F. et al., 2005, in Wilson A., ed., ESA-SP 684, The X-Ray Universe. ESA, Noordwijk, p. 94
- Anders E., Grevesse N., 1989, *Geochim. Cosmochim. Acta*, 53, 197
- Barthelmy S. D. et al., 2005a, *Space Sci. Rev.*, 120, 143
- Barthelmy S. D. et al., 2005b, *ApJ*, 635, L133
- Burrows D. N. et al., 2005, *Space Sci. Rev.*, 120, 165
- Butler N. R., 2007, *AJ*, 133, 1027
- Butler N. R., Kocevski D., 2007a, *ApJ*, 663, 407
- Butler N. R., Kocevski D., 2007b, *ApJ*, 668, 400
- Campana S. et al., 2006, *Nat*, 442, 100
- Chincarini G. et al., 2007, *ApJ*, 671, 1903
- Curran P. A., van der Horst A. J., Wijers R. A. M. J., 2008a, *MNRAS*, 386, 859
- Curran P. A., Starling R. L. C., O’Brien P. T., Godet O., van der Horst A. J., Wijers R. A. M. J., 2008b, *A&A*, 487, 533
- Cusumano G. et al., 2007, *A&A*, 462, 73
- de Pasquale et al., 2009, *MNRAS*, 392, 153
- Evans P. A. et al., 2007, *A&A*, 469, 379
- Evans P. A., Beardmore A. P., Osborne J. P., Goad M. R., Kennea J., Burrows D. N., Gehrels N., 2008a, *GCN Circ* 7328

- Evans P. A., Osborne J. P., Burrows D. N., Barthelmy S. D., 2008b, GCN Circ 7955
- Falcone A. D. et al., 2006, ApJ, 641, 1010
- Falcone A. D. et al., 2007, ApJ, 671, 1912
- Gehrels N. et al., 2004, ApJ, 611, 1005
- Gehrels N. et al., 2008, ApJ, 689, 1161
- Ghisellini G., Ghirlanda G., Nava L., Firmani C., 2007, ApJ, 658, L75
- Ghisellini G., Nardini M., Ghirlanda G., Celotti A., 2009, MNRAS, 393, 253
- Goad M. R. et al., 2006, A&A, 449, 89
- Goad M. R. et al., 2007, A&A, 476, 1401 (G07)
- Godet O. et al., 2009, A&A, 494, 775
- Hill J. E. et al., 2005, Proc. SPIE, 5898, 313
- Kalberla P. M. W., Burton W. B., Hartmaan D., Arnal E. M., Bajaja E., Morras R., Pöppel W. G. L., 2005, A&A, 440, 775
- Kennea J. A. et al., 2005, Proc. SPIE, 5898, 329
- Kocevski D., Butler N., Bloom J. S., 2007, ApJ, 667, 1024
- Kraft R. P., Burrows D. N., Nousek J. A. et al., 1991, ApJ, 374, 344
- Liang E. et al., 2006, ApJ, 646, 351
- Moretti A. et al., 2005, Proc. SPIE, 5898, 360
- Nousek J. A. et al., 2006, ApJ, 642, 389
- O'Brien P. T. et al., 2004, ApJ, 647, 1213
- Panaitescu A., Mészáros P., Gehrels N., Burrows D., Nousek J., 2006, MNRAS, 366, 1357
- Piran T., 2005, Rev. Mod. Phys., 76, 1143
- Racusin J. L. et al., 2009, ApJ, 698, 43
- Rees M. J., Mészáros P., 1994, ApJ, 430, L93
- Romano P. et al., 2006, A&A, 456, 917
- Roming P. W. A. et al., 2005, Space Sci. Rev., 120, 95
- Roming P. W. A. et al., 2009, ApJ, 690, 163
- Sakamoto T. et al., 2008, ApJS, 175, 179
- Sari R., Piran T., Narayan R., 1998, ApJ, 497, L17
- Shao L., Dai Z. G., 2007, ApJ, 660, 1319
- Shen R.-F., Willingale R., Kumar P., O'Brien P. T., Evans P. A., 2009, MNRAS, 393, 598
- Starling R. L. C., van der Horst A. J., Rol E., Wijers R. A. M. J., Kouveliotou C., Wiersema K., Curran P. A., Weltevrede P., 2008, ApJ, 672, 433
- Tagliaferri G. et al., 2005, Nat, 436, 985
- Uhm Z., Beloborodoc A. M., 2007, ApJ, 665, L93
- Vaughan S. et al., 2006, ApJ, 638, 920
- Willingale R. et al., 2007, ApJ, 662, 1093
- Zhang B., 2007, Chin. J. Astron. Astrophys. 7, 1
- Zhang B., Mészáros P., 2004, Int. J. Mod. Phys. A, 19, 2385
- Zhang B., Fan Y. Z., Dyks J., Kobayashi S., Mészáros P., Burrows D. N., Nousek J. A., Gehrels N., 2006, ApJ, 642, 354

## SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

- Table 5.** Enhanced XRT positions of GRBs observed by *Swift*.
- Table 6.** Power-law decay indices from light-curve fits. *Swift*.
- Table 7.** Times of the breaks in the light-curve fits.
- Table 8.** Best-fitting results for time-averaged spectra.
- Table 9.** Spectral indices for the time-resolved spectra.
- Table 10.** Column densities (in excess of the Galactic value) for the time-resolved spectra, in units of  $10^{20} \text{ cm}^{-2}$ .

Please note: Wiley-Blackwell are not responsible for the content or functionality of any supporting material supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.