Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **518,** 174–184 (2023) https://doi.org/10.1093/mnras/stac2937
Advance Access publication 2022 October 20

A real-time transient detector and the living *Swift*-XRT point source catalogue

P. A. Evans ,[1]★ K. L. Page ,[1] A. P. Beardmore ,[1] R. A. J. Eyles-Ferris ,[1] J. P. Osborne ,[1] S. Campana,[2] J. A. Kennea[3] and S. B. Cenko[4,5]

[1]*School of Physics and Astronomy, University of Leicester, Leicester LE1 7RH, UK*
[2]*INAF, Osservatorio Astronomico di Brera, via E. Bianchi 46, I-23807 Merate, Italy*
[3]*Department of Astronomy and Astrophysics, Pennsylvania State University, 525 Davey Lab, University Park, PA 16802, USA*
[4]*NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771, USA*
[5]*Joint Space-Science Institute, University of Maryland, College Park, MD 20742, USA*

Accepted 2022 October 10. Received 2022 October 10; in original form 2022 August 30

ABSTRACT

We present the Living *Swift*-XRT Point Source (LSXPS) catalogue and real-time transient detector. This system allows us for the first time to carry out low-latency searches for new transient X-ray events fainter than those available to the current generation of wide-field imagers, and report their detection in near real time. Previously, such events could only be found in delayed searches, e.g. of archival data; our low-latency analysis now enables rapid and ongoing follow-up of these events, enabling the probing of time-scales previously inaccessible. The LSXPS is, uniquely among X-ray catalogues, updated in near real time, making this the first up-to-date record of the point sources detected by a sensitive X-ray telescope: the *Swift* X-ray Telescope. The associated upper limit calculator likewise makes use of all available data allowing contemporary upper limits to be rapidly produced on demand. These facilities, which enable the low-latency transient system, are also fully available to the community, providing a powerful resource for time-domain and multimessenger astrophysics.

Key words: methods: data analysis – catalogues – X-rays: general.

1 INTRODUCTION

Serendipitous source catalogues have an important function in X-ray astronomy, helping us to constrain the population of X-ray emitting objects, identify new phenomena not intentionally targeted by observations, and characterize the time dependence of the X-ray sky. Such is their value that a source catalogue is a standard product of most imaging X-ray telescopes, e.g. *ROSAT* (White, Giommi & Angelini 1994; Voges et al. 1999; Boller et al. 2016), *XMM–Newton* (Saxton et al. 2008; Watson et al. 2009; Rosen et al. 2016; Traulsen et al. 2019), *Chandra* (Evans et al. 2010), and the *Neil Gehrels Swift Observatory* (Puccetti et al. 2011; D'Elia et al. 2013; Evans et al. 2014, 2020b). Some of these (e.g. *Chandra* and the *XMM* pointed catalogues) are narrow-field but deep; others [the *ROSAT* all-sky survey (RASS), the *XMM* slew survey, and the much-anticipated *eROSITA* (Predehl et al. 2021) catalogues] are all-sky but with less sensitivity.

The compilation of *Swift* X-ray telescope (*Swift*-XRT) observations lies between these two extremes, covering over 12 per cent of the sky, with typically more sensitivity than the all-sky surveys: in the 2SXPS catalogue (Evans et al. 2020b), we reported a median 0.3–10 keV source flux of 4.7×10^{-14} erg cm^{-2} s^{-1}, not much above the value from 3*XMM*-DR8 (2.2×10^{-14} erg cm^{-2} s^{-1}, 0.2–12 keV). Additionally, *Swift* typically observes a given object multiple times, giving a powerful insight into the variability of the X-ray sky. This combination of sensitivity, sky area, and variability information is likely to remain unique until the multiyear *eROSITA* catalogues are published.

However, there is a drawback to typical X-ray catalogues: they are always out of date.[1] Due to the time it takes to process the data, compile the results, create a data base and/or interface for catalogue access, and document everything, the newest data are often at least a year old by the time the catalogue is released and the observatory has collected more data in the meantime. For many purposes, this does not matter, but with the recent growth of time-domain and multimessenger astronomy, having an up-to-date catalogue is increasingly important. For example, when observing a new gravitational wave event, *Swift* regularly finds new X-ray sources, but to reliably determine which (if any) is related to the triggering event requires prior knowledge of that sky location either to identify the source as a known object or to determine whether it is above historical upper limits (Klingler et al. 2019).

Additionally, transients can be found serendipitously in X-rays, e.g. the fast X-ray transients found in *Chandra* data (e.g. Jonker et al. 2013; Bauer et al. 2017; Quirola-Vásquez et al. 2022). These were found by mining the data, but were already years old by

★ E-mail: pae9@leicester.ac.uk

[1]Except those created after a mission has ended.

the time of their discovery, preventing rapid follow-up and thus hampering the ability to expose the phenomenology of the transient, and the underlying physics. Further, the variability time-scales on which such transients can be studied are limited either to the (hours to days) time-scale of the discovery observation (e.g. the fast X-ray transients or nuclear quasi-periodic eruptions; Miniutti et al. 2019; Giustini, Miniutti & Saxton 2020) or potentially to much longer time-scales on which follow-up can be carried out (e.g. Starling et al. 2011; Strotjohann et al. 2016). In contrast, low-latency announcement of transients allows rapid, multiwavelength follow-up to probe the nature of the transient object, and enables ongoing observations that can probe the source's behaviour on time-scales of days to weeks/months. The scientific value of such low-latency announcement and response is well demonstrated by the case of SN2008D. This supernova occurred in the same galaxy as another supernova (SN2007uy) that was being regularly observed by *Swift*, and its sudden appearance in the X-rays was spotted by Soderberg et al. (2008). This rapid discovery enabled targeted observations of what proved to be the first X-ray detection of the shock breakout from a supernova.

Swift-XRT is an excellent tool for searching for new, serendipitous X-ray transients. Its combination of good effective area and low background, combined with *Swift*'s unique observing strategy and the rapid availability of its data, allows us to look for transients in near real time, with sensitivity to much fainter transient events than from typical wide-field detectors such as the RXTE All Sky Monitor (Levine et al. 1996) or Monitor of All-sky X-ray Image (Matsuoka et al. 2009).

In this work, we present the 'Living' *Swift*-XRT Point Source (LSXPS) catalogue and real-time transient detector. Software developed and running at the UK Swift Science Data Centre (UKSSDC) at the University of Leicester is executed each time new *Swift*-XRT data are received. Those data are searched for new transients, and added to the LSXPS catalogue, meaning it is kept up to date in nearly real time. The provision of the real-time catalogue is essential to the transient search: a reference image is needed to determine whether a source just detected is already known in X-rays, and if not, what upper limit can be placed on its historical flux. Because *Swift* often reobserves locations on the sky (and a single XRT snapshot is usually more sensitive than the RASS or *XMM* Slew Survey observations), it is commonly the case that the best reference image to use is a previous *Swift* data set, which is now always available in LSXPS.

2 DATA FLOW, FROM *Swift* TO LSXPS

In order to understand the 'living' nature of LSXPS and the latencies associated with transient detection, it is necessary to briefly introduce the structure and delivery of *Swift* data. A key and novel feature of *Swift*'s observing strategy is that it is always observing[2]; when a target goes into Earth eclipse, *Swift* immediately slews to the next target in its observing plan. Due to its low-Earth orbit, the longest continuous observing window is ~ 2700 s and for operational reasons these are usually no longer than 1800 s. This means that *Swift* both observes many targets during a day and that the same target may be observed multiple times on different orbits, with observations of other targets taking place in between. These individual pointings of a given target are referred to as *snapshots*, and (ordinarily)[3] all snapshots of a given target within a single UT day are collected together on the ground into a single *observation*: the basic data unit used for archiving and data access. Each observation is identified by its unique 11-digit obsID.

Swift typically has 8–12 ground-station contacts per day, and (again, unlike most missions) the downlinked data are made available rapidly[4] via the quick-look sites in the USA, UK, and Italy. These data are not necessarily 'complete': not all snapshots in a given observation may have occurred by the time of the downlink or not all observed data may have been downlinked during the pass. With each subsequent ground-station pass, observations may be updated and the data on the quick-look site replaced with a revised (more complete) version; to track this, metadata indicates how many times the observation has been processed. Eventually, each observation will be marked as 'complete' by the Swift Data Center (SDC), normally 3–7 d after the data were collected. The quick-look sites will label the data as 'FINAL FOR ARCHIVE', and shortly thereafter these data are moved from the quick-look area to the archive.

A simplified pictoral representation of this is given in Fig. 1. The coloured boxes on the top row represent consecutive pointings, ground station passes are shown with the state of the quick-look site after each pass.

2.1 Adding data to LSXPS

Quick-look XRT data are by definition incomplete, and can also occasionally be 'incorrect', in the sense that the incomplete housekeeping data can result in the pipeline processing inadequately filtering the data. For this reason, observations are not added to the actual LSXPS catalogue until they are marked as 'complete' by the SDC. Analysis of these incomplete observations is, however, extremely important for transient detection – not just for the obvious reason of finding transients as early as possible, but also to provide the most complete possible set of reference images for future searches.[5]

For this reason, observations are analysed each time data are received by the UKSSDC quick-look site, as shown in Fig. 1. The receipt of quick-look data triggers the LSXPS source detection, and the results are stored in a cache area that is not made publicly available. Transient checks (Section 4) are immediately carried out for each detected object and the XRT team are notified of any transient candidates, but none of the sources found in the observation, nor the observation itself, are added to the public catalogue (although the observation is available to the upper-limit server, Section 6). The transient checks for subsequent observations can make of use these cached data. When a revised version of an observation is received on the UK quick-look site, it is analysed and the cached results for the observation are replaced with those from this new analysis. In order to track this evolving situation, for LSXPS we index data not by the ObsID, but by a new property, the DatasetID. This is simply an incremental integer number starting from 1, and each new version of an observation is given a new DatasetID, so 'Obs 1 v1' and 'Obs 1

[2]More accurately, nearly always: it does not observe while passing through the Southern Atlantic Anomaly or slewing; due to its rapid slew speed, the latter is not a significant factor.

[3]For technical reasons, sometimes an observation may span multiple days, or there may be multiple observations in 1 d.

[4]Typically within about 2 h of the downlink.

[5]For example, imagine *Swift* observes the field of a new GRB and finds no new serendipitous transient. The GRB is re-observed the following day – before the first observation is marked as complete. If X-ray sources are found in this second observation, the analysis of the first observation is vital to determine whether any of these sources are new transients.

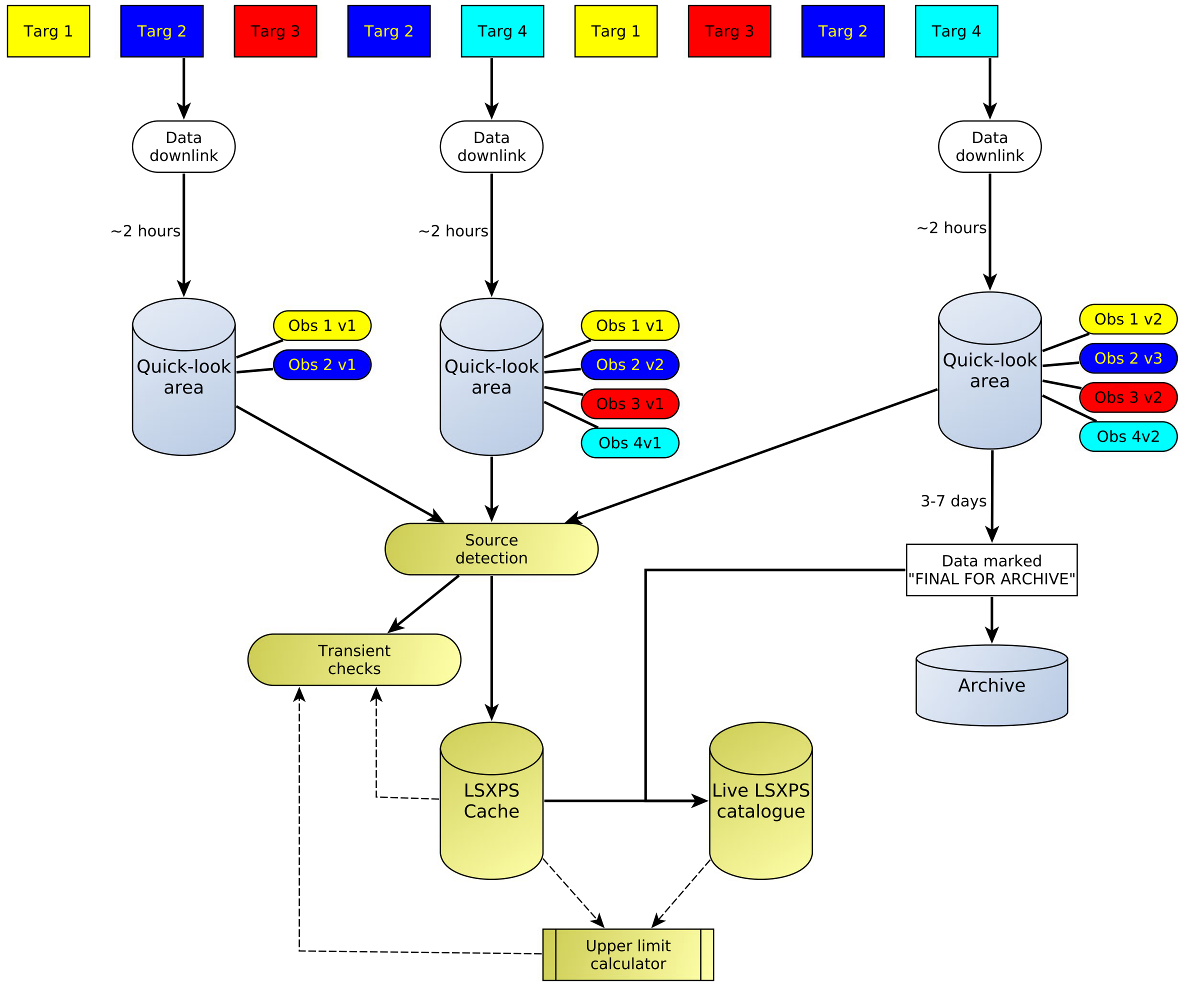

Figure 1. A simplified schematic of the *Swift* data flow. The top row shows a timeline of *Swift* pointings in a given day. Each box represents a pointing with *Swift*; multiple boxes with the same colour and label are multiple pointings at the same object ('snapshots'). Solid lines show data flow; dashed lines show data access.

v2' in Fig. 1 will have different DatasetIDs.[6] Once the observation is marked as complete, the cached analysis is merged into the catalogue (see Section 3.2).

2.2 Targets, observations, stacked image definition and evolution

Swift observations are planned in terms of *targets*, where a target, with a unique targetID, corresponds to a pointing location on the sky. Thus multiple observations of the same target will share a targetID; indeed, the 11-digit obsID is just the 8-digit targetID followed by a 3-digit number tracking how many times the target has been observed. Occasionally, either due to a human error or an inaccurate slew, two observations with the same targetID have significantly different pointings; if observations with the same targetID are more than 6 arcmin offset from each other then within LSXPS we define new targetIDs for each of these pointings, with observations being assigned to the targets as appropriate.

As with the previous SXPS catalogues, as well as analysing individual observations, we also create *stacked images*, where all observations of the same portion of the sky are combined to maximize the exposure.[7] A stacked image is defined as a specific set of overlapping targets (or a single target, with multiple observations), and the population of stacked images is created such as to produce the minimum number of stacks while ensuring that every observation and every overlap between observations is included in at least one stack. As in the previous catalogues, stacked images are assigned fake obsIDs, incremental numbers beginning with 10^{10}, and the term *data set* is a generic term covering individual observations and stacked images.

In the previous catalogues, stacked images were statically defined, but this is clearly not possible for a living catalogue: not only are new stacks being constantly added, but also the existing stacks can be modified. This can happen in two ways. First, a new observation may be taken of a target in a stack. In this case, the LSXPS processing does not change the 'ObsID' of the stack, but does assign it a new DatasetID (how many times this stack has been analysed is also

[6] DatasetIDs are set simply by the order in which data are processed; thus, two different versions of the same observation will not normally have consecutive DatasetIDs.

[7] We slightly increased the maximum size of a single stacked image to 2700 pixels (106 arcmin) to a side for LSXPS.

recorded); this is deliberately analogous to what happens to single observations when they are updated with more data. Secondly, a new target may be observed which overlaps an existing stack. In this case, the definitions of the stacks are updated. It is impossible to say a priori what the outcome of this will be. In the simplest case, a new stack is defined which is the same as an old stack but with the new target added; however, the effect may be that the addition of the new target causes a stack to grow beyond the 2700-pixel limit in which case multiple new stacks may be defined, or the new target may bridge a gap between two existing targets and so be able to merge all of their data into a single (new) stack. In either case, stacks that are superseded are marked as such, and are not analysed when new data are received. The handling of superseded stacks is discussed further in Section 3.3.

2.2.1 *Analysing stacks*

Stacks can take much longer to analyse than single observations – the runtime scales approximately with the number of snapshots – and in the most extreme cases stacked image analysis can take many weeks to run and its temporary disc usage can near a terabyte. It is thus impractical to simply run stacked analyses whenever new data are received, as is done for observations. Further, there is less need for this: as a general rule, each new observation of a stack adds only a (ever-decreasing) fractional increase in exposure time, and we do not search stacks for transients (Section 4) – their primary function for transients being the provision of historical upper limits. Therefore, stacked image processing is instead managed by a CRON job. This checks how many stacked image analyses are in progress, and if it is above some threshold[8] it will terminate without submitting any new jobs. Otherwise, it identifies n observations (where n is again a tunable parameter) which were marked as complete at least 24 h ago and have not yet been subjected to stacked analysis, finds all stacks that these observations contribute to, and triggers the analysis of those stacks. This analysis will include all observations in the stack which are marked as complete.

Long-running, high disc-usage processes are still an issue for this approach, since they can fill up all of the available stack-processing 'slots', causing the CRON job to exit without scheduling any new fields for days or even weeks. To avoid this, every time a stacked image is analysed the runtime is recorded in a data base. Stacks which took more than 24 h to run (or which supersede a stack which took more than 24 h) are not handled by the routine CRON job just described, but are instead flagged in the data base, and analysed at a lower cadence (loosely connected to the runtime) in a separate queue, as resources allow. By definition, these slow-running fields already contain a lot of data, therefore the incremental changes are small and the impact of this low-cadence update is minimal.

3 CATALOGUE CONSTRUCTION

The basic processing of a data set for LSXPS remains almost the same as in 2SXPS (Evans et al. 2020b) and is not described in detail here: a brief overview is given in Appendix A and the interested reader is encouraged to read the 2SXPS paper for full details; all we will do here is to restate the four energy bands used in LSXPS:

[8]Both this threshold and the frequency with which the CRON job runs are tunable parameters, allowing us to change things as our compute capactiy varies.

Table 1. The warning flag bits in the detection flag; taken from Evans et al. (2020b), table 4.

Bit	Value	Meaning
2	4	Source is within the extent of a known extended source.
3	8	Source likely a badly fitted piled-up source.
4	16	Position matched area flagged by manual screening.

Table 2. The definition of the data set warning flag, taken from Evans et al. (2020b), table 5.

Bit	Value	Meaning
0	1	Stray light was present, and fitted.
1	2	Diffuse emission identified.
2	4	Stray light badly/not fitted.
3	8	Bright source fitting issues[a]

Note. [a]The field contained a source that was heavily piled up in one band, but not fitted as such in another band. See section 3.7 of Evans et al. (2020b) for details.

(i) ***Total***: 0.3–10 keV
(ii) ***Soft***: 0.3–1 keV
(iii) ***Medium***: 1–2 keV
(iv) ***Hard***: 2–10 keV

and reproduce details of the source classification system. Sources are allocated a detection flag based on their statistical properties, these can have values of *Good*, *Reasonable*, and *Poor* (in the data base tables, these are numerical values of 0, 1, 2, respectively). The thresholds for these parameters were set such that the fraction of spurious sources in a given sample of objects is 0.003, 0.01, and 0.1, if that sample includes *Good*, *Good*, and *Reasonable* or all sources, respectively; individually the classifications correspond roughly to 3σ, 2σ, and 1σ significance. The spurious source rate will be considerably higher in the presence of artefacts such as stray light, or an extended/diffuse source. For this reason, objects that are affected by such things (identified either automatically or by the manual screening, Section 3.1) have extra warning flags set; on the website this is shown as descriptive text; in the data base these flags are bitwise values in the detection flag, as shown in Table 1. Data sets likewise have flags assigned to them to warn if there is a potential issue affecting the data set as it is possible (but unlikely) that these issues could alter the spurious detection fraction even among objects in that field with none of their warning flags set. This flag is also a binary flag, detailed in Table 2.

There is one small but important extension to the 2SXPS algorithm for LSXPS for individual observations: the search for new transients, which is described in Section 4.

LSXPS consists of all XRT data with at least 100 s of Photon Counting (PC) mode exposure remaining after the standard XRT-PIPELINE processing at the UKSSDC, and after the removal of times affected by bright Earth or unreliable pointing (Appendix A). These data are split into snapshots, only snapshots with at least 50 s of exposure are accepted.[9] All XRT data meeting these criteria were

[9]This can result in LSXPS data sets with only 50 s of exposure, for example if the original 100 s of usable data was spread over multiple snapshots, only one of which exceeded 50 s in duration. This was true of the previous SXPS catalogues but has not been explicitly stated before.

reprocessed with HEASOFT v6.29, and the XRT CALDB v20210915[10] and then the catalogue creation software was run on all historical data (we did not copy the results from the earlier catalogues). This processing 'caught up' with the archive data towards the end of 2022 March, and after a series of short 'catch-up' runs, the 'live' processing software to handle new quick-look data as described above was enabled on 2022 April 1.[11] The analysis of the stacked images built from historical data was also enabled at this point, with the queue of images to process automatically updated by the LSXPS processing as new data were received and so new stacks defined. Once the stacked image processing was up to date,[12] the CRON job described in Section 2.2.1 was activated.

While the source detection algorithm is almost unchanged from 2SXPS, the aggregation of the results to form the final catalogue required significant modification for LSXPS. The method by which the various detections are rationalized to a unique source list is unchanged (see section 3 of Evans et al. 2020b), but for 1/2SXPS this step was done once, after all data sets had been processed and the list of detected objects was final. For the dynamic LSXPS catalogue this is a continuous process, and takes place when the data set is added to the already-existing catalogue. For observations this is done when the data are marked as complete; for stacks this step is carried out immediately after the data set processing is completed.

There are four steps involved in adding a data set to LSXPS, which are:

(i) Manual screening of the field†.
(ii) Creating a new entry in the data sets table.
(iii) Incorporating the objects detected in the data set into LSXPS.
(iv) Handling older versions of a stacked image†.

Items marked with † may not take place for all fields. With the exception of item (ii) (which is just a data base operation), these operations are described in the following sections.

3.1 Manual screening of fields

In order to make the quality of LSXPS as high as possible, the analysis software identifies data sets in which the results may be unreliable and flags these for human inspection, which must be completed before the data set can enter the catalogue. There are two criteria that trigger this flagging.

The first of these relates to stray light and is not applicable to stacked images. X-rays from a source ∼30–80 arcmin from the XRT boresight (and so out of the field of view), can be diverted on to the XRT detector via a single reflection off the hyperbolic mirror surface; if the source is bright enought this can cause a pattern of concentric rings to appear on the XRT detector, which can give rise to spurious detections. We developed a technique to identify and fit this emission for 2SXPS (see Evans et al. 2020b, Appendix A), but this is not infallible, especially if there are real X-ray sources in or near the stray light. Such errors are immediately obvious to the human eye, so any field for which the automated system deemed stray light to be present is marked for human checking. Once this data set is marked as 'final', the XRT team are notified and directed to a web interface which allows them to compare the image with the fitted background model (which includes the stray light), and to either accept the model, mark the stray light as not being present (i.e. a spurious fit), or mark it as badly fitted and provide a better input position for the source of stray light. In the first case, the data set is accepted into LSXPS. In the latter two cases, the data set is reanalysed (reusing its existing DatasetID, as the underlying data have not changed), i.e. with stray light modelling disabled or with the new input parameters as appropriate. As soon as this reanalysis is complete, the field is again marked for screening. In the vast majority of cases, it is then accepted into the catalogue. In a small number of cases, the catalogue software is not able to produce an adequate model for the stray light. In this case, the XRT team will define an artefact (see below) covering the parts of the image affected by stray light, causing all detections therein to be flagged. As described in Appendix A of Evans et al. (2020b), stray light is not fully fitted to stacked images; instead the results of the fits to the observations making up the stacked image are taken and applied to the stacks, with only the normalization allowed to vary. For this reason, if the stray light fitting for a stack is poor, this must be handled by the creation of artefacts, as below.

The second reason for human verification of a data set is if median intersource distance in the data set is less than 80 arcsec. This can simply indicate a full or crowded field, but it can also indicate problems such as badly or unfitted stray light, instrumental artefacts such as residual contamination by bright Earth or hot pixels/columns/rows in the detector, or an area of diffuse emission which is not handled by our point-source detection system and so causes spurious detections. Such fields are flagged and the XRT team again notified and directed to a website where they can investigate the image. If there is evidence of a problem, they will define a circular or elliptical region, or a set of regions, and all of the objects inside these and in this specific data set have a warning flag applied to them. If the 'problem' is astrophysical in origin (as opposed to being of instrumental origin), e.g. an extended source, then the team member will mark these artefacts for propagation, which will mean that they are automatically applied to all future observations of these location on the sky. This helps to reduce the number of fields needing screening, since fields are only marked for inspection if the median distance between sources *that have not already had their warning flag set* is below 80 arcsec.

For data sets corresponding to observations in the 2SXPS catalogue, the results of the human screening of those data in that catalogue were automatically applied where possible. That is, if the automatic stray light model closely corresponded to that accepted for 2SXPS, it was accepted for LSXPS. Artefacts identified in 2SXPS were automatically applied to their corresponding LSXPS fields, and those fields were therefore only flagged for LSXPS screening if the unflagged sources has a low median separation.

3.2 Adding sources to LSXPS

When a data set is added to LSXPS, the objects detected in that data set are compared with the sources already in LSXPS. An object is deemed to match a source in LSXPS if their positions – including

[10]At the time of writing, all LSXPS data use these software and CALDB versions; however, they will doubtless be updated again in the future at which point only new data in LSXPS will use the newer software and calibration: we have no plans to periodically rebuild LSXPS from scratch following such updates.

[11]I did consider delaying by 1 d...

[12]Or *almost* up to date. Various technical challenges were encountered relating to certain stacks with very long (>100 h) runtimes, almost all of which cover the Galactic centre. Analysis of these images was deferred until the backlog was cleared, and is now – slowly – catching up.

astrometric uncertainty – agree at the 5σ level.[13] There are three potential results for each new object:

(i) The object has no counterpart in LSXPS.
(ii) The object is an additional detection of a source in LSXPS.
(iii) The object matches an LSXPS source, but improves the position of the LSXPS source (i.e. its position uncertainty is smaller than that currently in LSXPS).

The first two are straightforward: in case (i) the object is a new source which is added to LSXPS; in case (ii) the products and average properties of the existing source are updated to include the results from the new data set.

Case (iii) is more complex as there are several possible permutations. In the simplest scenario, the new position is identical to the LSXPS position,[14] but the uncertainty is smaller. Here, the catalogue entry is updated to use this smaller error and the processing continues as per case (ii), above. If, however, the new position is different from the LSXPS position, the situation is more complex as the association between individual (per-data set) detections and LSXPS sources may change. For example, if there are several X-ray sources close together on the sky, then moving one of these sources 1 arcsec to the East could result in some detections previously assumed to be this source now being better associated with a different object; or indeed the inverse. On the other hand, if a source position is revised it may now be the case that some of the detections, which had been attributed to that source, are not consistent with the new position, nor are they consistent with that of any existing source; i.e. the original 'source' is now found to be two distinct sources, which of course each need a record in LSXPS.

In order to identify and handle all of these cases, when a data set is added to the catalogue we select every LSXPS source that lies within the field of view of the new data set, and every individual detection of those sources; then we find any other sources within 5σ of those detections, and all detection of these; and so on recursively until no extra objects are found. This is analogous to picking up a paperclip with a magnet and so also retrieving all connected paperclips until a gap (in this case, a 5σ gap) is found, and essentially ensures that we collect every single detection in any data set and band that could potentially be affected by the new data set. We then apply the source matching algorithm from 2SXPS (Evans et al. 2020b, section 3.7) to this subset of detections. This yields the list of unique sources and which detections are associated with each, and the catalogue is revised accordingly. If a source moves such that its name changes, a new entry is created in the catalogue and the old entry deleted, but a record of this is kept. Records are also kept of which sources are deconvolved into multiple sources, and which individual detections are reassigned from one source to another. This book-keeping is extremely important for the usability of a dynamic catalogue such as LSXPS: if someone studies a particular source, perhaps announcing it to collaborators or the wider community, and then the source disappears without a trace this would be very unhelpful! The user interface (Section 6) ensures that if a user requests a 'superseded' source they will either be redirected to the new entry which has replaced this source, or if necessary presented with a list of the sources into which it has been deconvolved.

[13] More correctly, at the probability associated with a Gaussian 5σ confidence. Since the radial position errors should follow a Rayleigh distribution, this level was determined based on Rayleigh, not Gaussian, statistics.

[14] We define 'identical' as meaning 'does not change the IAU-format name of the object', where the IAU format is 'LSXPS JHHMMSS.S±ddmmss'.

3.2.1 *Source products and updates*

For each source, a series of products and measurements are created, and these are the same as in 2SXPS (see Evans et al. 2020b, section 4). In short, these comprise mean count-rates in each energy band, two time-averaged hardness ratios and light curves of all energy bands and hardness ratios, with one bin per observation, and one bin per snapshot. Unlike some catalogues, these light curves and count-rates are fully corrected for effects such as pile-up, vignetting, dead columns on the detector etc. and so can be used with no further corections needed. A measurement of variability using the Pearson's χ^2 is calculated for each time series. The two hardness ratios are defined as

$$\mathrm{HR1} = (M - S)/(M + S) \quad (1)$$

$$\mathrm{HR2} = (H - M)/(H + M), \quad (2)$$

where S, M, and H refer to the background-subtracted count-rates in the soft, medium, and hard bands, respectively.

Spectral and flux information is also provided for each source, with up to three methods employed per source, depending on the characteristics. Two spectral models are used: an absorbed power law, and an absorbed optically thin plasma model (APEC; Smith et al. 2001); absorption was calculated using the TBABS model (Wilms, Allen & McCray 2000). Flux conversions and (where appropriate) spectral properties were derived using XSPEC. For every source, fluxes were deduced using fixed spectral components: a power law with photon index 1.7 and an APEC with a plasma temperature of 1 keV; the absorption column was set to the Galactic value along the line of sight to the source, from Willingale et al. (2013). These values are available for every source but are the least accurate since they simply assume the spectral shape, regardless of the data. Where possible, the spectral parameters and flux are also inferred from the hardness ratios, using look-up tables linking (HR1, HR2) to the model parameters – for full details see Evans et al. (2020b) section 4.1 and Evans et al. (2014) section 4.2. For sources with at least 50 net counts in the total band, and a detection flag of *Good* or *Reasonable* (Section 5) with no stray light or diffuse source warning set, we also extract and fit a spectrum. This is done using the tools of Evans et al. (2009), and we fit both absorbed power-law and APEC models.

As LSXPS is a living catalogue, these source properties are not static. Every time a new data set is added to the catalogue, all sources that lie within the field of view of that data set (regardless of whether or not they were detected in it) have their products updated.

3.3 Handling older versions of a stacked image

When a new version of a stacked image is added to the catalogue, it is preferable to remove the old version, rather than keeping large numbers of obsolete data sets. If every source detected in the old version is also detected either by the new version, or by another data set, then there is no reason to keep the old version of the stacked image. Its data are removed from disc and from most of the data base tables.[15] It can happen, however, that a source is detected in a stacked image, but not in the updated version of that stack – especially for faint and variable or transient sources. If the old version of a stacked image contains the only detection of a source in LSXPS, that stacked image will not be removed from the catalogue. Instead, the source is flagged internally as an 'orphan', and the data set is marked as being

[15] A record of all such deleted data sets is still maintained but this is not part of the public catalogue.

an ‘obsolete stack’. Should an orphan source ever be detected in a new data set (single observation or stacked image), its orphan status is removed; should this happen to all orphans in an obsolete stack, that stacked image is then removed. The same process is followed when a new stacked image is analysed which supersedes a stack currently in the catalogue; that is, the new stack has a different definition (and obsID), but contains all of the targets in an old stack (see Section 2.2).

3.4 Static catalogue releases and data reprocessing

The introduction of a living catalogue, with near real-time updates, does not remove the need for ‘static’ catalogues, for projects which require a stable and unchanging reference such as where the reproducibility of a project is crucial. It is therefore our intention to issue frozen releases of the catalogue (3SXPS, 4SXPS, etc.) at regular intervals, likely every 1 or 2 yr. These will essentially be ‘snapshots’ of LSXPS, but will be full decoupled from LSXPS, so completely insulated from updates. Because of the time-scales on which stacked image analysis runs, it may not be possible to fully ‘synchronize’ the contents of these releases; that is, to ensure that the stacks in a static release contain all of the observations in the release, and only those observations. We do not see this as a major issue, as the key point of these releases is that their contents can be quantified and are unchanging, and this will be ensured. The first static release, 3SXPS, is envisioned for 2023.

Before building LSXPS, we reprocessed all historical XRT data with what was, at the time, the most recent HEASOFT and XRT calibration releases (Section 3). However, these are regularly updated[16] and calibration in particular is, by its nature, always retrospective. While the UKSSDC does periodically reprocess the entire XRT archive following such updates, we are not planning to rebuild the LSXPS catalogue following such processing. The CALDB and HEASOFT versions used in the live LSXPS analysis will be updated, but the data already in the catalogue will not be revised. The reason for this is that such a revision would essentially involve completely rerunning the LSXPS analysis: a process that takes around 6 months and monopolises compute time, and so clearly cannot be done routinely! It is not anticipated that any software or calibration revisions will have a major impact on the catalogue contents.[17] Changes in the XRT gain could potentially have a (very) small impact on the hardness ratios and spectral properties of sources; however, the catalogue web pages and the API (Section 6) are fully integrated with the UKSSDC software to provide high-quality analysis of point sources (Evans et al. 2009), making it easy to generate custom spectral products for sources of interest, making use of the latest software and calibration, and using the archive data locally reprocessed with these at the UKSSDC.

4 TRANSIENT DETECTION AND CLASSIFICATION

We define a serendipitous transient as an object which is uncatalogued in X-rays, has a flux above historical upper limits and was not the target of the *Swift* observation. In order to minimize the delay in detecting and announcing these transients, we do not wait for an observation to be marked as complete before carrying out transient checks, but rather carry out these checks immediately after each analysis of an observation.

The transient checks applied to each detected object can be broken into the following steps:

(i) determine whether the object is a known X-ray source
(ii) determine whether it is above historical upper limits
(iii) carry out basic analysis of the object
(iv) classification of the object

and these steps are described in the following sections. At each step the process may terminate; e.g. if the source is a known X-ray emitter, then steps (ii) onwards are not carried out.

During the bulk processing of historical data, the transient checks were carried out, but notifications of new transients (see below) were not produced and the candidate transients were not added to the main LSXPS transients data base but stored separately. This is because the events are, by definition, old and so not pressing, and may have been announced or published already. The primary purpose of enabling transient checking and analysis during the historical process was to test and debug the software and tune the automated classification process (Section 4.4); however, we may release the data base of these historical transients or some analysis thereof at a future date.

4.1 Is the object an known X-ray source?

We deem an object to be uncatalogued in X-rays, if it does not correspond to a source already in LSXPS (or in the cache)[18], and is inconsistent at the 5σ level with any object in the X-ray Master catalogue provided by HEASARC,[19] CSC 2.0 (Evans et al. 2020a) or the live XRT GRB catalogue (Evans et al. 2009).[20] This set of reference catalogues may be updated in future (e.g. with the addition of *eROSITA* catalogues), an up-to-date list will be maintained on the LSXPS website.

4.2 Comparison with historical upper limits

We compare the flux at discovery with upper limits obtained from *Swift* itself and also *XMM–Newton* and *ROSAT*.[21] *XMM* has a similar band-pass to *Swift*-XRT and we can use the total (0.3–10 keV) LSXPS band for this analysis. *ROSAT*, however, has a much softer band-pass, 0.1–2.4 keV, thus for comparison with *ROSAT* we use the combination of the soft and medium (0.3–1, 1–2 keV) LSXPS energy bands, hereafter the ‘SM’ band.

We create light curves with one bin per snapshot (i.e. *Swift* orbit) in the total, soft and medium LSXPS energy bands, summing the latter two to give the SM-band light curve. The peak per-snapshot count-rate is then identified from these light curves, defined as the bin with the highest 1σ lower limit.[22] We also calculate the mean count-rate of the object across the observation in the total and SM bands; if this has a higher 1σ lower limit this is instead taken as the peak rate. To enable us to compare upper limits in the next step, we also calculate the ratio, *R*, of the peak count rates in the total and SM bands.

[16]Indeed, HEASOFT 6.30 was released during the processing.
[17]In the unlikely event that something changes that is impactful, we will investigate how best to respond, and in such cases a complete rebuild of the catalogue cannot be ruled out.

[18]Excluding sources found in an earlier version of the observation being analysed.
[19]https://heasarc.gsfc.nasa.gov/W3Browse/all/xray.html
[20]https://www.swift.ac.uk/xrt_live_cat
[21]This list may also be revised in future, especially to include *eROSITA*; the LSXPS website will maintain an up-to-date list.
[22]This definition is robust against data points with very large errors; a count rate of 0.5 ± 0.05 is deemed higher than a rate of 2 ± 1.9.

3σ upper limits are then obtained from LSXPS (including cached data sets and stacked images, see Section 6), *XMM–Newton* and *ROSAT*; the latter two instruments' limits are obtained using the HILIGT service[23] (König et al. 2022; Saxton et al. 2022). Of these, only the RASS is an all-sky survey thus not all upper limits are always available. Upper limits are created in count-rate units using the native energy band of the instrument (0.2–12 keV for *XMM*, 0.2–2 keV for *ROSAT*, and 0.3–10 keV for LSXPS); for *XMM* limits from both pointed and slew data are requested. These limits are then converted to XRT count rates, using fixed conversion factors determined using PIMMS and assuming a power-law spectrum with spectral index 1.7 and an absorbing column of 10^{21} cm^{-2}; for *XMM* the conversions were calculated for each instrument and filter and so the appropriate one is selected. *XMM* upper limits are converted to total-band rates, and *ROSAT* limits to SM band rates. The deepest of these upper limits is then taken as the reference value to use. Since the *ROSAT* limit is in a different energy band from the others, for the comparison we multiply it by the ratio *R* defined above which helps to account for the spectral shape of the object, which is unknown at this point. Having determined the deepest upper limit, the peak count-rate from the appropriate band is compared with this limit, and if it is at least 1σ above it, the object is flagged as a candidate transient.[24]

During the historical processing, we found that all transient candidates that had a detection flag of *Poor* or one of the warning flags set (see Appendix A) were spurious detections, and indeed a significant fraction of alerts came from such objects (this is expected: a real source with a *Poor* detection is unlikely to be bright enough to pass the above tests, and while transients occurring within a stray light ring or known extended object are not impossible, they will be rare, whereas spurious detections with high apparent count rates are much more common). Accordingly, such candidates are automatically classified as *Spurious* (Section 4.4) and the analysis (Section 4.3) is not carried out and the XRT team is not alerted to such events. They are still posted to the team-only website discussed below, and they can be 'upgraded' if a team member studying the image believes them to be real, but they are not routinely analysed.

4.3 Basic transient analysis

For each candidate transient identified by the previous steps, custom data products are built using the tools that power the UKSSDC on-demand analysis services[25] (Evans et al. 2009). A light curve is built using all available XRT data, including those taken before the transient discovery (if available); this can result in evidence of historical emission from the source even though it has not been previously detected, since forced photometry at the location of a known source is more sensitive than a blind search.[26] This is binned in three ways: with fixed numbers of counts per bin, with one bin per snapshot, and with one bin per observation. A spectrum is also created using only the discovery observation; this is partly to maximize the S/N, and also to determine the spectral state specifically at the moment of outburst. This spectrum is automatically fitted with an absorbed power-law model, and if the automatic fit is successful and the reference upper limit (Section 4.2) is not from LSXPS, then the upper limit is revised, the conversion from *XMM* or *ROSAT* to XRT being recalculated with this fitted model, using PIMMS and the best-fitting model parameters.

These data products and the revised upper limit are posted to a website accessible only to the XRT-team, who are then notified by email of the potential transient and asked to confirm its status and classify it. This happens even if the redetermined upper limit is no longer below the peak count-rate, in case the spectral fit was poor; however, the notification email does indicate if this was the case.

4.4 Transient classification

The final step of transient analysis requires human checks to filter out objects which are definitely not serendipitous transients, and to further classify those which are. This step is currently semi-automated: the analysis system tries to suggest a classification, but a human is still required to verify or change this before the transient is made public.

Objects may be deemed not to be serendipitous transients for the following reasons:

(i) **Spurious detection**: the 'transient' may be an obviously spurious detection. Many of these cases are identified automatically as noted above but some may not be – for example if the X-ray 'source' is actually an artefact of optical loading, the automated system will (usually) warn that this is possible, but will not automatically flag the source as spurious.

(ii) **Targeted transient**: the object is a transient, but is not serendipitous, it is the object *Swift* was observing (e.g. as a target-of-opportunity observation of a transient discovered elsewhere). These cases can usually be flagged automatically by the analysis code, which uses the OBSQUERY class in the SWIFTTOOLS PYTHON package[27] to determine what object was the target of the observation. In cases where the target observation had a large position uncertainty, the automatic flagging is less reliable.

(iii) **It is not a transient**: analysis of the historical light curve may show a series of previous bins in which the source flux can be constrained, and is consistent with that at the time of discovery. In this case the source is clearly not transient, but as a result of the specifics of the current observation happens to have been blindly detected for the first time. Alternatively, the spectral fit to the transient data may cause the historical upper limit to be above the peak flux; if the XRT team member confirms that this fit is reliable, the transient nature of the source is retracted.

There is an additional category that is not listed above, a *proprietary transient*. This is an object that passes the above tests and is a real transient, but the parent observation was part of a programme designed to monitor a specific region to look for new transients. In this case whether the transient should be deemed 'serendipitous' is open to interpretation, but it is clearly inappropriate for the *Swift* team to routinely and rapidly announce the primary science of someone's approved observing programme, and before they have even had time to examine the data! In these cases, the *Swift* team will endevaour to contact the observation PI and encourage them to announce the transient, and the object will not be posted to the public transients catalogue or webpage without their consent.

[23] xmmuls.esac.esa.int/hiligt/

[24] Note that the *ROSAT* limit extrapolated to the 0.3–10 keV band is *only* used for comparing with the upper limits to determine which is deepest. If *ROSAT* is selected by this test, it is the upper limit in the SM band that is used, and compared with the peak SM-band count rate.

[25] https://www.swift.ac.uk/user_objects

[26] In each bin, if the 3σ lower limit on the count rate is non-zero, a 'detection' comprising a count rate and 1σ errors is produced; otherwise, a 3σ upper limit is created.

[27] https://www.swift.ac.uk/API

Items which are not caught by the above are still not necessarily transients, and are manually classified into one of the following groups, collectively referred to as 'interesting classes'.

(i) **Outburst**: If the light curve shows historical detections of the source (from the forced photometry) but at a flux level below that of the new observation, the source is likely an outburst rather than a transient.

(ii) **Low significance**: The Eddington bias (Eddington 1940) causes objects detected close to the sensitivity limit to be preferentially recorded with count-rates significantly above their true value; this can result in sources which are in reality below the historical upper limit (i.e. not transients) to be misclassified as transients. We are not able to say a priori whether a given source is affected by this bias, although Evans et al. (2014, section 6.2 and fig. 10) showed that SXPS sources with fewer than ∼ 30 counts are affected by this bias. Sources whose peak count-rate is $<3\sigma$ above the $3-\sigma$ upper limit and which have fewer than 30 counts in the detection are usually classed as *low significance*.

(iii) **Needs follow-up**: This is somewhat of a catch-all category, where the source does not obviously fall into either category above, but there is some doubt as to its nature and further observations are felt necessary before it can be definitively identified as a transient.

(iv) **Confirmed transient**: The object is definitely a new transient.

The reader will immediately realize that these classifications are to a considerable degree subjective, and they should be interpreted as such. All events falling into these categories are made publicly available (Section 4.5) and the classifications are intended as a guide and to enable filtering, rather than a definitive statement; this also helps the clear candidates to be easily identifiable from among the (larger number of) low significance and outburst events.

The *outburst* category is particularly subjective, in that in reality the only phenomenological difference between an outburst and a transient is the sensitivity of previous observations. Two identical sources, with quiescent fluxes of (say) 10^{-14} erg s^{-1} and outburst fluxes 100 times higher, could be classed one as a transient and one as an outburst, if the latter had a historical observation long enough for forced photometry to detect it in quiescence.

Swift J051001.8−055931 presents a good case study of this subjectivity. It was discovered in observation 00014892016 (version 4, LSXPS DatasetID: 225587) which began at 23:59 UT on 2022 July 20. The light curve (Fig. 2) showed a detection in the previous observation that had been accumulated over the previous 16 h,[28] and several non-detections (grey arrows) more than 240 d earlier. At the time of the discovery, the LSXPS stacked image contained only those non-detections, from which an upper limit of 0.0084 ct s^{-1} was deduced (blue arrow); the count-rate at discovery is clearly well above the upper limit, but the automated system recommended this be classified as an outburst due to the detection in the previous observation. As the day-old 'historical' detection could indicate the rise of the transient, this was instead manually classified as *needs follow-up*. A *Swift* ToO observation was requested to investigate further; this revealed that the source, while still detectable, had dropped below the historical upper limit and thus the 'transient' discovery was probably either a short-lived outburst or just a low-significance fluctuation. In this case, we classified it as an outburst, but *low significance* would be equally justifiable.

[28]Not accumulated continuously; that observation contained 3.2 ks of exposure gathered over the 16-h interval.

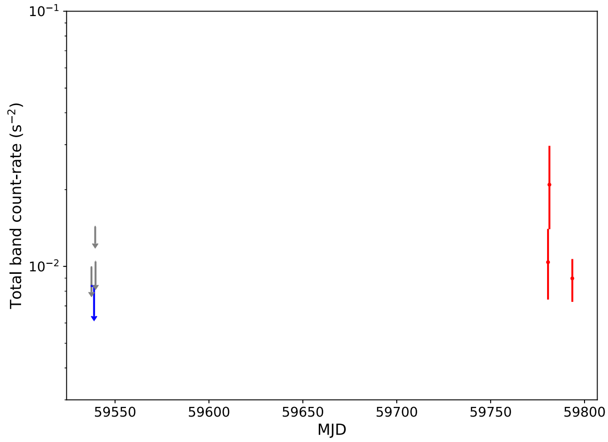

Figure 2. 0.3–10 keV light curve of *Swift* J051001.8–055931, one bin per *Swift* observation. The transient was discovered in the observation corresponding to the penultimate data point; the final data point comes from the follow-up ToO observations, i.e. this data point was not available when the transient was first reported to the XRT team. The grey upper limits are from the individual observations; the blue upper limit is the limit from the stacked image which combined these observations.

4.5 Transient announcement and data access

As soon as a transient is classified into one of the interesting classes (Section 4.4), it is published on the public transient webpage: https://www.swift.ac.uk/LSXPS/transients, and added to the public transients data base (Section 5). At this time, no push notifications are produced, although we are discussing the recording of these transients (or some subset thereof) in the Transient Name Server.

If a transient is reobserved after discovery, whether intentionally or serendipitously, its light curve is updated and also a second spectrum constructed and fitted, using all data from the discovery observation onwards. Normally, these products are only updated until the discovery data set enters the live LSXPS catalogue (Section 3), at which point the transient will correspond to an LSXPS source (and the transient page and data base will contain this link), and the source products are kept up to date in the main catalogue (Section 3.2.1). However, because source products are only updated to include observations that have been added to LSXPS, they lag a few days behind real time. On occasion there may be a transient for which follow-up observations have been requested and for which it is desirable to maintain up-to-date products in real time. In such cases, the XRT team will manually flag the transient and its products in the transient area will continue to be updated each time new data are received. An example of such an event is *Swift* J023017.0+283603, a source that demonstrates the power of our new transient detector. It is consistent with the nucleus of a known galaxy and has a very soft X-ray spectrum, and so was initially announced as a probable tidal disruption event (Evans, Campana & Page 2022). Had this source only been found in later analysis, as was hitherto typical, this classification would likely have been unchallenged. However, the prompt transient alert from our new system allowed rapid and long-term multiwavelength monitoring, exposing variability on timescales of weeks, challenging the initial classification of the object and exposing its true, and much more enigmatic, nature (Evans et al. 2022, in preparation).

The transient data and data base tables can be accessed through the LSXPS website and SWIFTTOOLS PYTHON module as described in Section 6.

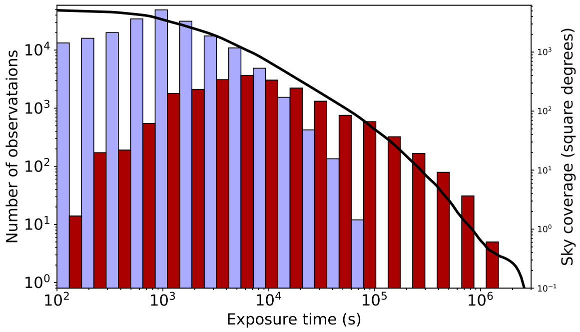

Figure 3. The sky coverage and exposure details of the LSXPS catalogue as of 2022 August 30. The solid line shows the sky coverage (corrected for overlaps) as a cumulative function of exposure time (i.e. area with at least the exposure indicated). The histogram shows the distribution of exposure time per data set, with the individual observations shown in light blue and the stacked images in red; the different colours are each half the width of the actual bins.

5 CATALOGUE CHARACTERISTICS AND CONTENTS

Being a dynamic catalogue, any statement of characteristics is obviously tied to the time at which they were generated: the 'live' status of the numbers given here is available on the LSXPS website. The numbers given below were obtained at 15:00 UT on 2022 August 30.

The LSXPS catalogue contains 279 021 unique sources, composed from 1.5 million blind detections across all four energy bands, or 756 769 'obs sources'.[29] The median 0.3–10 keV flux (assuming a power-law spectrum) is 4.1×10^{-14} erg cm^{-2} s^{-1}. The sky coverage of LSXPS is shown in Fig. 3. The total unique sky coverage is 5186 deg^2; 3503 deg^2 have at least 1 ks of exposure and $\sim$700 deg^2 have been observed with at least 10 ks. The 2SXPS catalogue contained data up to 2018 August 1; thus, on average LSXPS has grown by 49 sources per day and the unique sky coverage is increasing at 0.94 deg^2 per day.

The source detection mechanism employed for LSXPS was the same as for 2SXPS, and so only summarized here. For full details, see Evans et al. (2020b), especially sections 3.5 and 7, figs 6 and 7, and tables 3 and 6.

As for 2SXPS, we provide 'clean' and 'ultraclean' subsets of the sources and data set. Clean sources are those with a best detection flag of 0 or 1 (i.e. *Good* or *Reasonable* with no other warning bit set); OpticalLoadingWarning, StrayLightWarning, and NearBrightSourceWarning all unset; and a field flag of 0 or 1 (see Table 2). Ultraclean sources are a subset of the clean sources, with detection and field flags of 0. There are 190 902 clean sources and 170 372 ultraclean sources in LSXPS, a growth rate of 30 (26) per day for the clean (ultraclean) classes.

The actual catalogue contents are organized into six publicly available tables, which are described in Appendix B.

6 ACCESS TO LSXPS: THE WEBSITE AND API

The LSXPS catalogue is available from the UKSSDC website via https://www.swift.ac.uk/LSXPS. This website provides interactive tools to explore the catalogue (i.e. cone searches, filters, etc.) and visualizations of all data sets and sources and relevant properties. The catalogue tables (Appendix B) can also be downloaded from here as FITS or CSV files. Due to the dynamic nature of the catalogue, these files are updated every hour. These hourly snapshots are available via the website for one week after creation; daily snapshots are created at mid-night UT each day and available indefinitely. It is thus recommended that for any serious project based on a frozen snapshot of LSXPS, one of these daily snapshots is used to enable the work to be reproducible in future.

The website also includes an upper limit calculator, as provided for the previous catalogues; the operation of this system was described in Evans et al. (2014), section 4.4. For LSXPS, this server queries all data, including those still in the cache (i.e. observations that are not marked as complete; see Section 2.1). The ESA HILIGT upper limit server (König et al. 2022; Saxton et al. 2022) currently queries 2SXPS for upper limits from *Swift*, but this will soon be updated to use LSXPS.

The catalogue can also be accessed through the SWIFTTOOLS PYTHON module. This module, which can be installed via PIP, contains two components, both of which are extensively documented. SWIFTTOOLS.SWIFT_TOO is maintained by Jamie Kennea (at Penn State University) and provides functionality including access to *Swift* data, the ability to query the observability and observing history of a given object, and the ability to submit ToO requests; it is documented at https://www.swift.psu.edu/too_api/. SWIFTTOOLS.UKSSDC is maintained by Phil Evans, and was added in SWIFTTOOLS v3.0. This version includes tools to access data and query different *Swift* catalogues; this includes GRB data and catalogues, and LSXPS. All functionality of the LSXPS website, including the upper limit server, is available through the API, including integration with the tools that allow on-demand analysis of sources (themselves provided through the SWIFTTOOLS.UKSSDC.XRT_PRODS module),[30] and much of the functionality has also been backported to 2SXPS, so that this can also be queried by the module. This module is documented at https://www.swift.ac.uk/API/.

We ask that users of the catalogue, via the website, data tables, or API, acknowledge this in their subsequent publications: by citing this paper and also including in the acknowledgements the sentence: *This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester.*

7 SUMMARY

We have presented the first-ever low-latency, sensitive X-ray point source catalogue and transient detector. This enables rapid follow-up of X-ray transients too faint to trigger all-sky monitors and so opens up a new region of transient phase space for exploration. This is enabled by the unique provision of a 'living' source catalogue, LSXPS, which is kept constantly up to date every time a new *Swift* observation is marked as completed. This catalogue and associated upper limit server provide not only a reference for our transient detector, but also a valuable resource for the rapidly growing field of time-domain and multimessenger astronomy generally, in which it is often useful to have quick access to the most up-to-date analysis of a given sky location.

[29] An 'obs source' is produced by merging all detections of the same object across the energy bands within a single data set; this is analogous to what the *XMM* catalogues call a detection.

[30] Prior to SWIFTTOOLS v3.0, this module was SWIFTTOOLS.XRT_PRODS, and while that path is still active and supported, it now simply silently wraps SWIFTTOOLS.UKSSDC.XRT_PRODS. There are excellent reasons for this, but they are very boring.

ACKNOWLEDGEMENTS

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester. PAE, KLP, ABP, and RAJE-F acknowledge UK Space Agency support. SC acknowledges the support from the Italian Space Agency, contract ASI/INAF no. I/004/11/5. This work is partially supported by a grant from the Italian Ministry of Foreign Affairs and International Cooperation Nr. MAE0065741.

DATA AVAILABILITY

All of the original *Swift* data are available via the three *Swift* data centres: (https://www.swift.ac.uk/swift_live/, https://swift.gsfc.nasa.gov/archive, https://www.ssdc.asi.it/mmia/index.php?mission=swiftmastr). The data derived for this paper are available through the LSXPS catalogue web pages (https://www.swift.ac.uk/LSXPS/) and can also be accessed using the SWIFTTOOLS PYTHON module, available via PIP.

REFERENCES

Bauer F. E. et al., 2017, MNRAS, 467, 4841
Boller T., Freyberg M. J., Trümper J., Haberl F., Voges W., Nandra K., 2016, A&A, 588, A103
D'Elia V. et al., 2013, A&A, 551, A142
Eddington Sir A. S., 1940, MNRAS, 100, 354
Evans P. A. et al., 2009, MNRAS, 397, 1177
Evans I. N. et al., 2010, ApJS, 189, 37
Evans P. A. et al., 2014, ApJS, 210, 8
Evans I. N. et al., 2020a, Bull. Am. Astron. Soc., 52, 154.05
Evans P. A. et al., 2020b, ApJS, 247, 54
Evans P. A., Campana S., Page K. L., 2022, Astron. Telegram, 15454, 1
Giustini M., Miniutti G., Saxton R. D., 2020, A&A, 636, L2
Jonker P. G. et al., 2013, ApJ, 779, 14
Klingler N. J. et al., 2019, ApJS, 245, 15
König O. et al., 2022, Astron. Comput., 38, 100529
Levine A. M., Bradt H., Cui W., Jernigan J. G., Morgan E. H., Remillard R., Shirey R. E., Smith D. A., 1996, ApJ, 469, L33
Matsuoka M. et al., 2009, PASJ, 61, 999
Miniutti G. et al., 2019, Nature, 573, 381
Predehl P. et al., 2021, A&A, 647, A1
Puccetti S. et al., 2011, A&A, 528, A122
Quirola-Vásquez J. et al., 2022, A&A, 663, A168
Rosen S. R. et al., 2016, A&A, 590, A1
Saxton R. D., Read A. M., Esquej P., Freyberg M. J., Altieri B., Bermejo D., 2008, A&A, 480, 611
Saxton R. D. et al., 2022, Astron. Comput., 38, 100531
Smith R. K., Brickhouse N. S., Liedahl D. A., Raymond J. C., 2001, ApJ, 556, L91
Soderberg A. M. et al., 2008, Nature, 453, 469
Starling R. L. C. et al., 2011, MNRAS, 412, 1853
Strotjohann N. L., Saxton R. D., Starling R. L. C., Esquej P., Read A. M., Evans P. A., Miniutti G., 2016, A&A, 592, A74
Traulsen I. et al., 2019, A&A, 624, A77
Voges W. et al., 1999, A&A, 349, 389
Watson M. G. et al., 2009, A&A, 493, 339
White N. E., Giommi P., Angelini L., 1994, IAU Circ., No. 6100, #1
Willingale R., Starling R. L. C., Beardmore A. P., Tanvir N. R., O'Brien P. T., 2013, MNRAS, 431, 394
Wilms J., Allen A., McCray R., 2000, ApJ, 542, 914

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Figure S1. Diagrammatic outline of the source detection mechanism, reproduced from fig. 2 of Evans et al. (2020b).
Table S1. Contents of the 'Sources' data base table.
Table S2. Contents of the 'Data sets' data base table.
Table S3. Contents of the 'Detections' data base table.
Table S4. Contents of the 'ObsSource' data base table.
Table S5. Contents of the 'xCorr' data base table.
Table S6. Contents of the 'oldStacks' data base table.
Table S7. Contents of the 'Transients' data base table.