



X-ray transients in the Chandra archive

Introducing the cumulative distribution discriminator (CuDiDi)

I. Saathoff^{1,2}  and J. Larsson^{1,2,*} 

¹ KTH Royal Institute of Technology, Department of Physics, 10691 Stockholm, Sweden

² The Oskar Klein Centre for Cosmoparticle Physics, AlbaNova University Centre, 10691 Stockholm, Sweden

Received 16 September 2025 / Accepted 13 March 2026

ABSTRACT

X-ray transients on sub-observation timescales represent a diverse and underexplored class of astrophysical phenomena, from stellar flares and magnetar bursts to extragalactic fast transients and supernova shock breakouts. We present a systematic search for such events across 20 212 *Chandra* ACIS observations using a new detection pipeline that combines source identification, light-curve analysis, catalogue cross-matching, and a novel statistical classifier, the cumulative distribution discriminator (CuDiDi). From 1420 initial candidates, we identified a high-confidence golden sample of 765 transients spanning a broad range of timescales, fluxes, and spectral shapes. The candidates are distributed across the whole sky and show a wide range of durations with a median of 10 ks. A subset of fast events lasting $\lesssim 30$ s displays very soft spectra and is likely due to flaring dwarf stars, although extragalactic phenomena cannot be ruled out for all of them. The comparison with previously published samples showed that CuDiDi identifies most known transients while imposing somewhat stricter variability criteria, and it also extends the total sample of *Chandra* transients to include shorter events. We deliver a comprehensive catalogue of sub-observation *Chandra* X-ray transients and establish a general method for exploiting archival datasets to uncover rare short-lived high-energy phenomena.

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

1. Introduction

The X-ray sky is home to a rich variety of transient phenomena, ranging from magnetic reconnection in stellar atmospheres to tidal disruption events (TDEs) and various cataclysmic explosions. The importance of X-ray observations of these phenomena is highlighted by the results from the recently launched *Einstein Probe* (EP; Yuan et al. 2025), which has opened up a new discovery space for X-ray transients. Observations with the EP have revealed a large diversity in the X-ray emission that accompanies the explosive deaths of massive stars as supernovae (SNe) and gamma-ray bursts (GRBs), providing new insights into the nature of the central engines and the properties of relativistic jets. This diversity is illustrated by detailed studies of individual EP transients, including cases associated with GRBs (EP240315a, Liu et al. 2025b; EP240801a, Jiang et al. 2025; EP241107a, Eappachen et al. 2026), collapsars (EP240414a, Sun et al. 2025; van Dalen et al. 2025; EP250108a, Rastinejad et al. 2025; EP250827b, Srinivasaragavan et al. 2025), and a likely binary neutron star merger (EP250207b, Becerra et al. 2026; Jonker et al. 2026).

At the same time, archival data from X-ray telescopes operating since the early 2000s provide a rich resource for identifying transients serendipitously caught within the field of view (FOV) of other targets. Of particular interest are data from *XMM-Newton* (Jansen et al. 2001) and *Chandra* (Weisskopf et al. 2000), both launched in 1999, as well as *Swift* (Burrows et al. 2005), launched in 2004, which all cover the soft X-ray band (~ 0.5 – 10 keV). A number of searches for transients in the archives of these telescopes have been presented, including Alp & Larsson (2020),

Pastor-Marazuela et al. (2020), De Luca et al. (2021), Ruiz et al. (2024), and Khan et al. (2025) for *XMM-Newton*, Yang et al. (2019), Lin et al. (2022), Quirola-Vásquez et al. (2022), Quirola-Vásquez et al. (2023), Zhang & Feng (2023) and Dillmann et al. (2025) for *Chandra*, and Eyles-Ferris et al. (2022) for *Swift*. Since the transients detected in these searches are discovered years to decades after they occurred, it has not been possible to perform multi-wavelength follow-up observations, although we note that the field of the *Chandra* fast X-ray transient (FXT) 141001 was serendipitously observed at optical wavelengths with the Very Large Telescope Visible MultiObject Spectrograph only 80 min after it occurred. The observation did not reveal any counterpart (Bauer et al. 2017). There are now efforts to identify X-ray transients from *Swift* and *XMM-Newton* in real time to enable follow-up observations (Evans et al. 2023; Quintin et al. 2024).

The transients identified in the archival searches have uncovered diverse phenomena, ranging from flares from Galactic objects (e.g. Khan et al. 2025) to a population of energetic extragalactic transients referred to as FXTs, some of which may be powered by millisecond magnetars created in binary neutron star mergers (Bauer et al. 2017; Xue et al. 2019; Lin et al. 2022; Quirola-Vásquez et al. 2022, 2023, 2024, 2025). Adding to the diversity of X-ray transients, Alp & Larsson (2020) identified events with properties consistent with SN shock breakouts (SBOs). This included the case of FXT 110621 (also identified by Novara et al. 2020), whose properties are similar to the X-ray SBO from SN2008D, for which the connection between the X-ray signal caught by *Swift* and the optical SN could be firmly established (Soderberg et al. 2008).

Although multi-wavelength follow-up observations of transients discovered in archival data are not possible, the identification and studies of their host galaxies can provide insights into

* Corresponding author: josla@kth.se

their origin by constraining their luminosities and environments. A number of such studies have been carried out (Novara et al. 2020; Lin et al. 2022; Eappachen et al. 2022, 2023; Inkenhaag et al. 2024; Eappachen et al. 2024; Quirola-Vásquez et al. 2025), but in many cases, there are either multiple host candidates or no detected host galaxy (Lin et al. 2022; Eappachen et al. 2022, 2023; Inkenhaag et al. 2024). A handful of the FXTs identified in *Chandra* observations are associated with very nearby galaxies ($\sim 4\text{--}30$ Mpc), while the majority are more distant, with redshifts $z \gtrsim 0.1$ (Quirola-Vásquez et al. 2022, 2023). The transients uncovered in *XMM* observations by Alp & Larsson (2020) whose host galaxies are identified also fall into the latter category (Novara et al. 2020; Eappachen et al. 2024). The FXTs with clear host galaxy associations reveal a wide range of peak luminosities, ranging from $\sim 10^{38}\text{--}10^{40}$ erg s $^{-1}$ cm $^{-2}$ for the nearby subset (Quirola-Vásquez et al. 2022) to between $\sim 2 \times 10^{43}$ erg s $^{-1}$ cm $^{-2}$ (FXT 110621, Novara et al. 2020; Eappachen et al. 2024) and $\sim 3 \times 10^{47}$ erg s $^{-1}$ cm $^{-2}$ (FXT 141001, Quirola-Vásquez et al. 2025) for the distant cases. This indicates multiple progenitor channels.

The transients are also diverse in terms of timescales, spectra, and fluxes, which presents a challenge for identifying them in archival searches. The situation is further complicated by the inhomogeneous nature of the archival data, including large variations in the duration of individual observations. The different transient searches above have therefore uncovered transients with partly different properties, depending on the datasets and detection algorithms used.

We present a new transient detection algorithm that we applied to the full public archive of the *Chandra* advanced CCD imaging spectrometer (ACIS; Garmire et al. 2003). This transient search was designed to further exploit the potential of the archival data by also capturing faint and short transients, as well as transients occurring in observations with short durations. The durations of many EP transients are $\sim 100\text{--}1000$ s (e.g. Sun et al. 2025; Liu et al. 2025b; Zhou et al. 2026), down to ~ 12 s for EP240408a (Zhang et al. 2025), while the durations of the transients so far uncovered in the *Chandra* archive are $\sim 3\text{--}50$ ks (Quirola-Vásquez et al. 2022, 2023). This suggests that a population of short *Chandra* transients might not have been identified yet.

To identify robust candidates, we developed a multi-stage pipeline involving automated source detection, light-curve analysis, and catalogue cross-matching, followed by a novel statistical classification based on the cumulative distribution of the total light-curve counts. A final round of visual screening was used to flag artefacts and define a golden sample of high-confidence transients. The full pipeline and data selection process are described in Sect. 2. In Sect. 3, we present the X-ray properties of the golden sample. This is followed by a discussion in Sect. 4, covering the properties of a small population of very short transients and an assessment of the algorithm performance. The paper concludes with a summary in Sect. 5.

2. Identifying transient candidates

An X-ray source detected as a transient in a given observation may be either a true transient in the sense that it happens only once, or be a flare from a source that often also emits faint persistent emission together with occasional bright flares. The former category includes GRBs originating from collapsars and binary neutron star mergers (Kumar & Zhang 2015; Sun et al. 2025; Liu et al. 2025b; Jonker et al. 2026), SBOs (Waxman & Katz 2017;

Levinson & Nakar 2020; Alp & Larsson 2020), and TDEs (Rossi et al. 2021), where the type of TDE most relevant for producing short transients is the disruption of a white dwarf by an intermediate mass black hole (Jonker et al. 2013; Maguire et al. 2020; Li et al. 2025; Levan et al. 2025).

The category of sources with flaring activity in the X-ray band include active galactic nuclei (AGNs; Grupe et al. 2001; Stathopoulos et al. 2022; Liu et al. 2025a), X-ray binaries (XRBs, Galloway et al. 2020; Sguera et al. 2020; Romano et al. 2023; Haberl et al. 2025), magnetars (Woods & Thompson 2006; Rea & Esposito 2011; Kaspi & Beloborodov 2017), cataclysmic variables/novae (Starrfield et al. 2016; König et al. 2022), and stars (Pye et al. 2015; De Luca et al. 2020; Kowalski 2024; Zhao et al. 2026). We focused on identifying true transients and new flaring objects, meaning that we excluded transients associated with known X-ray sources. We also excluded transients associated with stars.

2.1. Sample selection

We used data from the *Chandra*/ACIS archive, which is well suited for identifying faint transients due to the combination of high angular resolution and a low background. ACIS provides imaging spectroscopy in the 0.3–8 keV energy range and comprises ten CCDs, where the four ACIS-I chips are arranged in a square covering $17' \times 17'$ and the six ACIS-S chips are arranged in a row covering $50' \times 8.3'^1$. Up to six (I and S) chips can be active in a given observation, which sets the maximal FOV.

All publicly available ACIS observations obtained between 1999-08-14 and 2024-12-31 were retrieved from the high energy astrophysics science archive research center (HEASARC) using the *chanmaster* mission interface². There are a total of 21 914 observations, out of which 20 212 had data available for download. The total exposure time of these observations is 467 Ms, where 63% (37%) has ACIS-S (ACIS-I) as the aimpoint. The sample includes all ACIS operating modes, including various subarrays and observations with gratings.

The necessary event and auxiliary files (evt2, asol, bpix, msk) were downloaded for each observation to enable the extraction of images, light curves, and spectra. The data were not reprocessed because the benefits are small for most observations, while the computational cost is high. We verified this by reprocessing ~ 100 randomly selected observations and comparing the results, which showed minimal effect on transient selection. We did note small differences in the spectra in some cases, but this will not significantly impact our objectives to identify transients and characterize their basic properties. The flow chart in Figure 1 shows an overview of the analysis procedure applied to each observation, and the full details are provided below.

2.2. Source detection and extraction of science products

Images were generated using the *Chandra* interactive analysis of observations (CIAO; Fruscione et al. 2006, version 4.15.2 with CALDB 4.10.7) tool *fluximage*, restricted to the 0.3–2 keV energy range to optimise the detection of soft X-ray sources. This energy range was chosen to optimise the search for SBOs in future analysis of the selected transients, as these are expected to have very soft spectra (Waxman & Katz 2017). The encircled

¹ <https://cxc.harvard.edu/cal/ACIS/>

² The data were queried via the *astroquery.heasarc* module in Python, using the *chanmaster* mission catalogue.

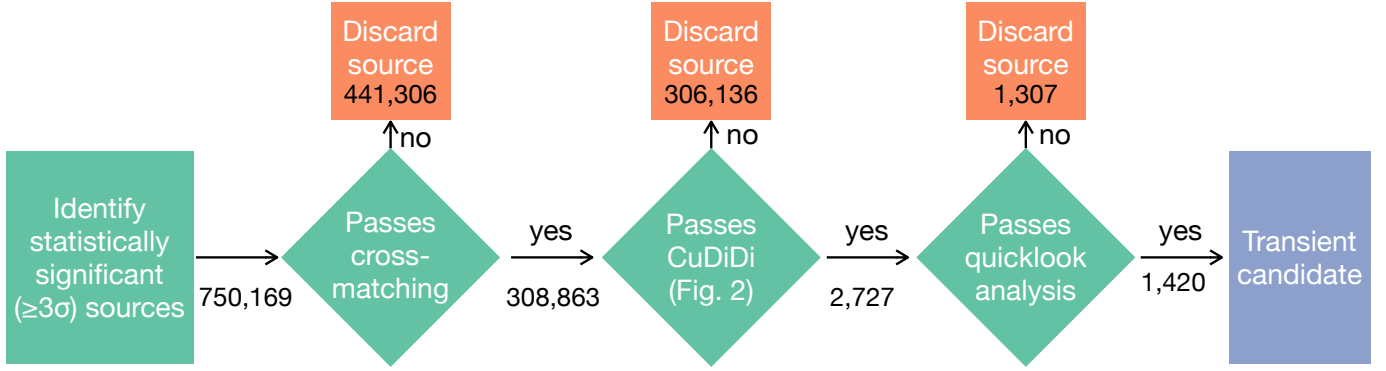


Fig. 1. Flow chart of the transient candidate search method. The numbers shown at each step indicate how many sources were retained and excluded at that stage of the analysis.

counts fraction was set to 0.393, corresponding to the 1σ integrated volume of a 2D Gaussian. Source detection was performed using `wavdetect`, with wavelet scales of [1 2 4 6 8 12 16 24 32] pixels to cover a broad range of spatial scales, and a detection threshold of 10^{-6} . A total of 1 019 308 sources were detected in all the *Chandra* observations, out of which 750 169 had a significance $> 3\sigma$ (as reported by `wavdetect`) and were retained for further analysis. The sources detected at lower significance were used in a later stage of the analysis to flag possible recurring sources among the final transient candidates (Sect. 2.4).

The light curves were extracted using `dmextract` in the 0.5–7 keV energy range. The source regions were defined as the 3σ elliptical output regions of `wavdetect`, excluding any areas that overlapped with the 3σ regions of nearby sources. The background counts were extracted from annular regions with inner and outer radii of $2r_{\max}$ and $3r_{\max}$, respectively, where r_{\max} is the semi-major axis of the source ellipse. As for the source regions, any areas of background regions that overlapped with nearby sources were excluded. Point sources that overlapped with extended X-ray sources were flagged at a later stage of the analysis (Sect. 2.4).

Although the nominal *Chandra* ACIS frame time is 3.2 s, a 1 s binning was adopted for the cumulative distribution analysis to provide consistent temporal sampling across all observations, including observations that have a better time resolution due to the use of subarrays. For reference, 21% of the observations in our sample have a time resolution < 3 s and 13% have a time resolution < 1 s. The choice of 1 s bins has a negligible impact on the CuDiDi selection (see Sect. 2.3). Good-time intervals (GTIs), which exclude periods of high background due to flaring or instrumental artefacts, were applied to all light curves, and the same GTIs were used for the spectral extraction. These GTIs were generated following standard CIAO tools and filtering criteria³.

For spectral analysis, source spectra were extracted directly from the event files using `specextract`, using the same source and background regions as for the light curves. CIAO generates background-subtracted spectra along with the corresponding response matrix files and ancillary response files for each source. The spectra were grouped to 1 count per spectral bin, needed for fitting with C-stat in XSPEC, and the point spread function (PSF) correction was applied. In cases where the standard background region resulted in extraction errors, larger annuli (up to $12r_{\max}$) were used to ensure a valid background spectrum. These spectral

products are used for quick-look inspection and subsequent spectral fitting. A full list of tools and parameters used throughout the pipeline is provided in Table 1.

2.3. Selection criteria

Each detected source was then required to pass a series of selection criteria to be classified as a transient candidate. The criteria are as follows: (a) cross-matching with external catalogues to remove known sources; (b) application of the CuDiDi criterion; and (c) visual inspection via quick look analysis. Each step is described in detail below.

2.3.1. (a) Cross-matching

Many sources identified by the `wavdetect` algorithm are expected to be known persistent or recurring sources, or associated with stars (see e.g. Yang et al. 2019). To identify and exclude such sources, cross-matching was performed using a $3''$ matching radius against several astronomical catalogues. This radius was chosen to account for the combined astrometric uncertainties of the *Chandra* source positions and the comparison catalogues. While *Chandra* offers sub-arcsecond astrometric precision for on-axis sources (typically $\leq 0.6''$), the PSF broadens significantly at larger off-axis angles, leading to increased positional uncertainties that can exceed $1''$ (Evans et al. 2024)⁴. Additionally, small systematic offsets between different datasets may be present. The $3''$ radius thus provides a conservative but effective compromise.

For reference, out of the 750 169 sources detected at $>3\sigma$ significance with `wavdetect`, $\sim 38\,000$ ($\sim 5\%$) have 1σ positional uncertainties $> 1''$, while only 27 sources have 1σ positional uncertainties $> 3''$. Assuming that the positional probability distributions are described by circular 2D Gaussians, we estimate the number of sources that will have their true positions outside the $3''$ matching radius as $N_{\text{outside}} = \sum_{i=1}^N e^{-3^2/(2\sigma_i^2)} \approx 3300$, where σ_i is the positional uncertainty of each source and the sum is over all N detected sources. This corresponds to 0.43% of the sample and shows that the probability of missing real counterparts is very low, although we note that the positional errors in comparison catalogues will add to the uncertainties. The probability of excluding genuine new sources due to mismatches is higher, as discussed further below.

First, matches were sought in the SIMBAD database (Wenger et al. 2000), where we filter out sources categorised under the ‘‘Taxonomy of Stars’’ and ‘‘Sets of Stars’’ classifications. Sources

³ See <https://cxc.cfa.harvard.edu/ciao/ahelp/dmextract.html> for details.

⁴ See also <https://cxc.harvard.edu/cal/ASPECT/celmon/>

Table 1. Overview of software tools, input parameters, and selection criteria.

Level	Tool/criterion	Parameters
Observation	download_chandra_obsid fluximage wavdetect	evt2, asol, bpix, msk energy range 0.3–2 keV encircling counts fraction = 0.393 [1 2 4 6 8 12 16 24 32] (default) sigthresh = 10^{-6} (default) $\geq 3\sigma$ significance threshold
Sources	dmextract specextract	energy range 0.5–7 keV 3σ region wavdetect, bkg annulus $(2-3)r_{\max}$ 3σ region wavdetect, bkg annulus $(2-3)r_{\max}$
Criteria	NED cross-matching SIMBAD cross-matching Gaia cross-matching STONKS cross-matching eRASS1 cross-matching Cumulative distribution	$3''$ radius; excluding star-like types $3''$ radius; excluding star-like object types $3''$ radius $3''$ radius $3''$ radius $\mu_{\text{upper}} > \mu_{\text{lower}} + 0.5 + 0.2$

without a counterpart in SIMBAD were then cross-matched against the NASA/IPAC extragalactic database (NED), excluding objects classified as stellar, planetary nebulae, or SNe⁵. We verified that there were no X-ray sources detected shortly before the optical detections of SNe, implying that this cut did not exclude any SBO candidates. Finally, unmatched sources were cross-matched with Gaia (Gaia Collaboration 2016) data release 3 (Gaia DR3; Gaia Collaboration 2023) to identify stellar counterparts, ensuring that stars not already flagged in SIMBAD or NED were excluded. The Gaia cross-match helps to exclude nearby stars that may show low-level variability or flares. The relatively large matching radius of $3''$ also ensures that stars with significant proper motion are excluded, although we note that rare cases of stars with very high proper motion may have been missed.

Additional cross-matches were performed with the STONKS catalogue (Quintin et al. 2024), which compiles known X-ray sources across a broad range of timescales (including data from *XMM-Newton*, *Swift*, *Chandra*, *eROSITA*, and *ROSAT*; see Quintin et al. 2024 for details), and the *eROSITA* all-sky survey first data release (eRASS1) (Merloni et al. 2024), which lists X-ray sources detected in the 0.2–2.3 keV band with detection likelihoods greater than 6. In cases where the only match in STONKS was with a source in the *Chandra* source catalog (CSC), we searched the latest version of the catalogue (CSC 2.1, Evans et al. 2024) and our own catalogue to exclude sources that were detected at $>3\sigma$ in more than one observation. As the version of the CSC included in STONKS only contains observations up to the year 2014 (CSC 2.0), we used our own catalogue of detections to identify and exclude recurring sources in later observations. These cross-matches serve to remove known and recurring X-ray emitters from our sample, such as AGN and XRBs. The sample could potentially be further cleaned from flaring stars by using the various observed correlations for stellar flares in X-rays (e.g. Pye et al. 2015), although this is likely to provide a clear diagnostic only for the brightest cases.

In addition to the catalogue cross-matching, we manually compared our candidates with the transient samples reported by Quirola-Vásquez et al. (2022, 2023) (QV transients). Out of the 22 published QV transients, we recovered 13 in our analysis

⁵ See https://ned.ipac.caltech.edu/help/ui/nearposn-list_objecttypes?popup=1 for a full list.

and removed them from our catalogue. The remaining 9 sources were not retained, either because they were associated with known objects in SIMBAD or NED, or did not pass our CuDiDi threshold. A detailed overview is provided in Table A.1 (Appendix A). This is also discussed further in Sect. 4.2.

In total, the cross-matching described above removed 441 306 out of 750 169 sources from the sample. The majority of the excluded sources (70%), were removed in the cross-matching with SIMBAD, NED, and Gaia, while the rest were removed in the X-ray cross-matching.

The conservative matching radius that was used to minimise contamination from known and recurring sources comes at the expense of a higher risk of excluding genuine transients due to matches with nearby unrelated sources. This is exemplified by the case of CDSF-S XT2 (FXT 150322, Xue et al. 2019), which was initially misidentified with a nearby AGN. The probability of a mismatch will be different for each source, depending on the source density in the field, as well as its positional uncertainty and uncertainties in the matching catalogues. To obtain a rough estimate of the impact of mismatches, we reran the matching step for 500 randomly selected observations (containing $\sim 19\,000$ sources) using a more conservative matching radius of $1.5''$. The 3σ positional uncertainties on the *Chandra* sources are smaller than this for 80% of the sample, so matches with more distant sources have a higher probability to be due to unrelated nearby sources.

We found that the more conservative radius reduces the matches in the SIMBAD/NED/Gaia step by 27%, while the X-ray matches are reduced by 14%. The smaller difference in the latter case is partly due to matches with recurring *Chandra* sources, and when these matches were excluded, the reduction in the X-ray cross-matching was 22%. While this investigation indicates that the probability of mismatches is relatively high, we note that this conservative approach also implies that the analysis of the final selected transients is unlikely to be affected by contamination of nearby sources.

2.3.2. (b) Cumulative distribution discriminator

The CuDiDi is a novel statistical tool we developed to identify transient sources based on the shape of their light-curve distributions (see Figure 2). While the method is in principle

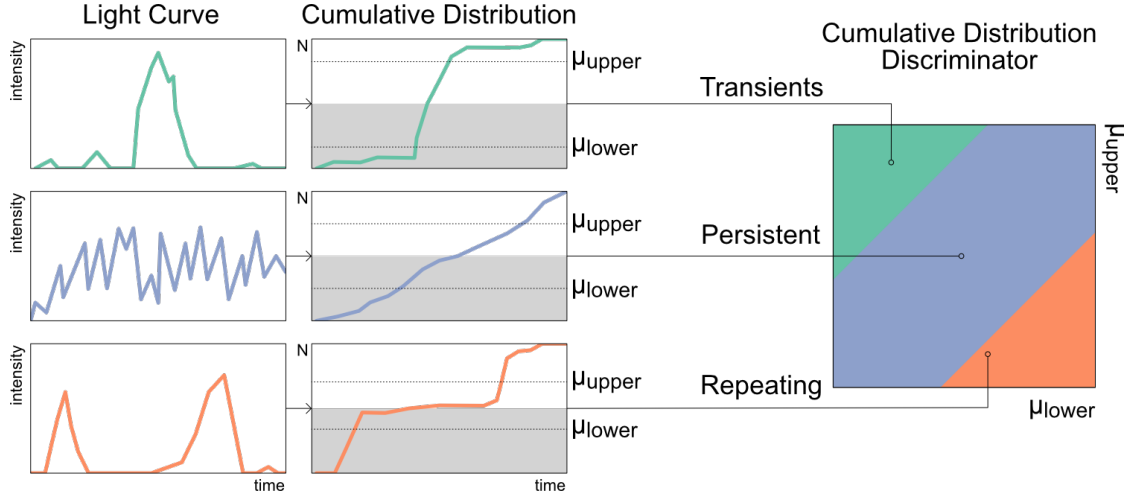


Fig. 2. Source classification method using CuDiDi. Left: sketches of representative light-curve behaviour, which illustrate typical variability patterns for persistent, repeating, and transient sources. Centre: corresponding CDFs. The y -axis shows the cumulative number of counts N , normalised by the total so that each distribution reaches unity. Each CDF is split at the median (0.5) into lower (0–0.5, shaded light grey) and upper (0.5–1.0) halves. The mean values of each half, μ_{lower} and μ_{upper} , are indicated. Right: CuDiDi plot showing μ_{lower} vs. μ_{upper} for each source. The diagonal band (± 0.2 around $\mu_{\text{upper}} = \mu_{\text{lower}} + 0.5$) defines the persistent region. Sources above the band are classified as transients, while sources below the band are considered repeating sources. This visualisation shows how CuDiDi uses statistical asymmetries in cumulative flux distributions to separate source classes.

sensitive to repeating and persistent behaviours as well, its design and application here are focused specifically on isolating transients.

For each source, the cumulative distribution function (CDF) of the photon arrival times is calculated from the 0.5–7 keV energy range light curves, normalised and divided at the median ($y=0.5$) value of the CDF into two halves: the lower half (0–0.5) and the upper half (0.5–1.0). The mean values of each half are computed and denoted as μ_{lower} and μ_{upper} , respectively. Background counts, which are usually evenly spread over time, can slightly flatten the shape of the CDF, reducing the contrast between the two halves and affecting the calculated mean arrival times (μ_{lower} and μ_{upper}). Subtracting the background is not practical in our case because it can produce negative counts in bins with low signal, complicating the analysis. However, since *Chandra* generally has low background levels, this effect is expected to be minimal in most cases (see Appendix B for further details.)

The CuDiDi diagram (Figure 2) plots μ_{lower} versus μ_{upper} for each source, with classification boundaries defined by a diagonal band centred on $\mu_{\text{upper}} = \mu_{\text{lower}} + 0.5$, with a width of ± 0.2 . Sources above this band are classified as transients, those within it as persistent, and those below as repeating. The choice of the band width was guided empirically by visual inspection of representative examples, balancing sensitivity to genuine transient behaviour against robustness to statistical fluctuations and noise. We also find that the selected threshold is somewhat stricter in terms of transient selection than previous searchers, but sufficiently low to recover the majority of the previously identified FXTs (see Sect. 4).

To illustrate the effectiveness of the CuDiDi, we present in Figure 3 a set of representative sources. The top row shows three transient sources, while the bottom row displays three non-transient sources chosen from the same observations for comparison. The light curves (left column) highlight the contrasting temporal behaviours: transients exhibit sharp, localised bursts of activity, whereas non-transient sources show more uniform or recurrent variability. The corresponding cumulative distributions (middle column) emphasize this difference, with transients

producing strongly asymmetric distributions compared to the smoother rise of persistent sources. The right column summarises these properties in the CuDiDi plane, where each source is represented by its mean values μ_{lower} and μ_{upper} from the lower and upper halves of the cumulative distribution. The green shaded area indicates the region within which most transient sources fall. This example illustrates how the chosen band width provides a practical threshold for distinguishing source classes: narrower bands risk excluding borderline transients, whereas wider bands would allow contamination from persistent sources.

The choice of time binning has a negligible impact on the values of μ_{lower} and μ_{upper} as long as the bins are small. For the six examples shown in Figure 3, the values change by $< 10^{-4}$ if the time bins are increased from 1 s to 3.2 s (the frame time of the observations), which implies a negligible impact on the transient classification. By contrast, increasing the time bins to 1 ks for the same sources changes the values of μ_{lower} and μ_{upper} by up to 0.03, which could change the classification of sources near the boundaries (see also Appendix B for examples with 5 ks bins). The impact of increasing the time bins is expected to be larger for shorter observations and for sources with low count rates, which justifies our approach of using a uniform 1 s binning for all observations. Only sources that satisfied the CuDiDi criterion were retained for further analysis, leading to the exclusion of 306 136 additional sources, leaving 2727 transient candidates.

2.3.3. (c) Quick-look analysis

Each candidate passing the automated criteria was subjected to a visual quick-look inspection of the *Chandra* images. In this step, we checked for obvious instrumental artefacts, in particular the pronounced readout streaks that can mimic transient sources. Candidates affected by such artefacts were excluded. This visual inspection led to the exclusion of 1307 sources. Only sources that passed all three stages – cross-matching, CuDiDi filtering, and visual inspection – were classified as transient candidates and retained for detailed follow-up analysis.

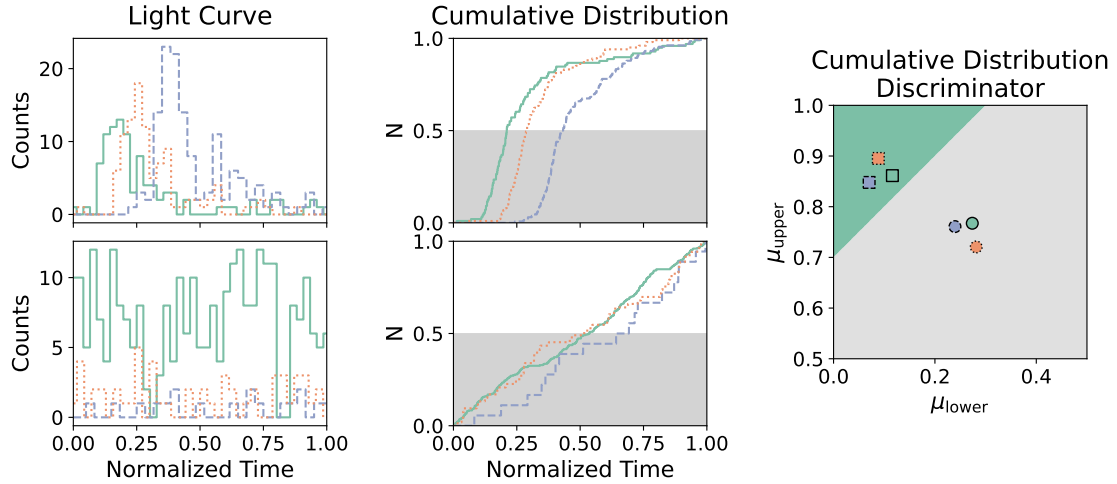


Fig. 3. Example X-ray light curves and their classification with CuDiDi. Left column: total counts binned at 1 ks for visual clarity and normalised in time for three representative sources per row. Transient sources are shown in the top row, and non-transient sources in the bottom row. Middle column: corresponding CDFs, obtained from the light curves with 1 s bins, normalised to unity. The shaded grey area indicates the lower half (0–0.5) of the cumulative fraction. Right column: CuDiDi diagram showing μ_{lower} vs. μ_{upper} , where μ_{lower} and μ_{upper} denote the mean values of the lower and upper halves of the cumulative distributions, respectively. Transient sources are marked as squares, and non-transient sources are shown as circles. The shaded green band indicates the transient region, defined as $+0.2$ above the diagonal $\mu_{\text{upper}} = \mu_{\text{lower}} + 0.5$. Sources in this region are classified as transients, while those below fall into the non-transient category. The colours and line styles in all panels indicate observation ID: solid green line = 7048, dashed blue line = 16174, dotted orange line = 20120. This demonstrates how CuDiDi captures the distinct variability patterns of transient and non-transient sources.

Out of the 750 169 unique sources detected across all ACIS observations, 1420 passed the full sequence of selection criteria and were classified as transient candidates. The number of sources excluded at each stage of the analysis is summarised in Figure 1.

2.4. Diagnostic flags

To facilitate further investigation and vetting of individual candidates, a series of diagnostic flags were added to the catalogue. These are summarised in Table 2 and discussed below. This approach allows us to define a golden sample of high-confidence transients, while preserving the broader dataset for users interested in exploring more ambiguous or potentially interesting cases.

Uniqueness flag (non_unique). While sources detected at high significance in more than one observation were excluded in the cross-matching described above, we performed an additional check by matching the transient candidates against all other sources detected by `wavdetect` using a radius of $3''$. This finds matches with sources detected at a significance below 3σ , which may be spurious or due to weak nearby sources, but could also indicate that the transient is in fact a recurring source. Candidates with at least one positional match in a different observation from our database were flagged as non-unique.

Galactic exclusion flags (snr_match and extended_fov). Since our primary goal is to identify a high-purity sample of transient candidates, we excluded sources associated with prominent extended X-ray sources, which add uncertainties to the source detection and analysis. Specifically, we flagged all candidates found in observations targeting well-known supernova remnants (SNRs) such as Cas A and Tycho, identified through observation ID matching, under the flag `snr_match`. Additionally, extended Galactic sources were flagged with `extended_fov` via manual inspection of FOV images.

These associations do not definitively rule out a background transient event located behind the foreground structure, so we

retained the sources in the full candidate catalogue but flagged them for caution. We stress, however, that the vast majority of these transients are expected to be spurious. Based on the surface area of the main SNRs and the number of observations, we expect less than one genuine transient in our sample to be superposed on them by chance.

Artefact flags (readout_streak and edge_affected). Two main observational artefacts were identified that could mimic transient behaviour. First, sources lying along readout streaks⁶ (McCullough & Rots 2005; Evans et al. 2010), particularly in observations of bright sources such as Cyg X–1, were flagged following visual inspection of the images, using the flag `readout_streak`. Although the quick-look analysis excluded the most obvious cases, these flags capture more ambiguous sources that were not removed in that step.

Secondly, sources near the detector edge may show apparent variability due to dithering motion⁷ moving the source region partially on and off the detector (Evans et al. 2010; Primini et al. 2011). To mitigate this, we used the CIAO tool `dither_region`⁸ to estimate the fractional exposure of the source region. Any source with a less than full coverage (i.e. a coverage fraction < 1.0) was flagged as `edge_affected`.

Background flag (bkg_correlated). Since the CuDiDi analysis is based on total (non-background-subtracted) counts, background flaring can produce artificial variability in the light curve. To identify such cases, we computed the Pearson correlation coefficient between the source and background light curves, restricted to the interval between the 5th and 95th percentiles of the source photon arrival times (i.e. the T_{90} interval). Sources with a correlation coefficient exceeding 0.15 were flagged as `(bkg_correlated)`

⁶ https://cxc.cfa.harvard.edu/ciao/dictionary/readout_streak.html

⁷ <https://cxc.cfa.harvard.edu/ciao/why/dither.html>

⁸ https://cxc.cfa.harvard.edu/ciao/ahelp/dither_region.html

Table 2. Overview of diagnostic flags included in the transient candidate catalogue.

Flag Name	Type	Description	Count
non_unique	Boolean	True if the transient is matched with a source detected at low significance in other <i>Chandra</i> observations (sources detected at $>3\sigma$ in more than one observation are already excluded).	58
snr_match	Boolean	True if the source is associated with a known Galactic SNR (e.g. Cas A, Tycho), based on observation ID cross-matching.	5
extended_fov	Boolean	True if the source lies in or near an extended Galactic structure identified by manual inspection of the FOV.	13
readout_streak	Boolean	True if the source overlaps with a known instrumental readout streak, commonly from bright sources like Cyg X-1.	29
edge_affected	Boolean	True if the source region (as determined by <code>wavdetect</code>) is not fully covered due to detector dithering; computed via <code>dither_region</code> .	392
bkg_correlated	Boolean	True if the Pearson correlation between the source and background light curves during the T_{90} interval exceeds 0.15, indicating likely background flaring.	49
start_or_end	Boolean	True if the T_5 or T_{95} time falls within the first or last 5% of the observation, indicating the event may lie at the edge of the observing window.	266

Notes. These flags helped us to identify likely artefacts or ambiguous cases that may not represent genuine transient events. Out of a total of 1420 candidates, 765 form a golden sample with none of the listed flags set to true.

and are likely dominated by background flaring. This threshold was empirically determined by testing several candidate values and visually inspecting sources around the corresponding decision boundary to assess how the flagging behaviour changed.

Although standard GTI filtering excludes intervals with high global background, which typically removes strong, long-duration flares, it does not always capture low-level or localised background variability. In our sample, 49 sources showed significant correlation between source and background light curves, indicating that some background fluctuations remained even after GTI filtering. This additional check helps to identify and exclude false positives that would otherwise contaminate the transient candidate list.

Start- or end-of-observation flag (`start_or_end`). Transient classification requires sufficient temporal baseline to establish quiescence before and after the event. Sources whose variability occurred at the very beginning or end of the observation are ambiguous in this regard. We therefore flagged any source for which the start (T_5) or end (T_{95}) of the light curve overlapped with the first or last 5% of the observation duration.

Based on the full set of diagnostic flags described above, we define a high-confidence golden sample of transient candidates as those for which all cautionary flags are set to `False`. This conservative subset is intended to minimise contamination from instrumental artefacts, background variability, and foreground Galactic sources. In Sect. 3, we explore the spatial distribution, temporal characteristics, and X-ray spectral properties of the 765 sources in this golden sample.

3. X-ray properties

3.1. General properties

Figure 4 shows the distribution of exposure times for all 20 212 *Chandra* ACIS observations included in our search, along with the subset containing the 765 transient candidates. The exposure times span a wide range, from brief observations of less than

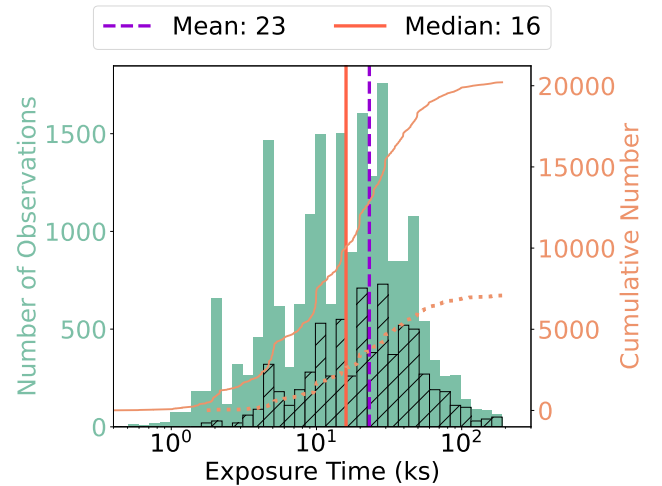


Fig. 4. Histogram (green; left y -axis) and cumulative distribution (orange lines; right y -axis) of the exposure times for the 20 212 *Chandra* observations. The vertical dashed purple and solid dark orange lines indicate the mean and median exposure times of the full sample, respectively. The distribution for 765 transient candidates is scaled by a factor of 10 for visual clarity (hatched bars and dotted orange line).

1 ks to deep observations exceeding 100 ks. The overall distribution peaks around 30 ks, with a mean of 23 ks and a median of 16 ks, consistent with the typical duration of *Chandra* imaging observations (Primini et al. 2011).

Importantly, transient candidates (hatched bars) are found across the full range of exposure times. Their distribution closely tracks that of the full dataset, indicating that the detection method does not significantly favour either short or long observations. In particular, the presence of candidates in short ($\lesssim 5$ ks) and long ($\gtrsim 50$ ks) exposures demonstrates the method's robustness across varying observational depths. This also suggests that the transient events we identify are not systematically tied to any

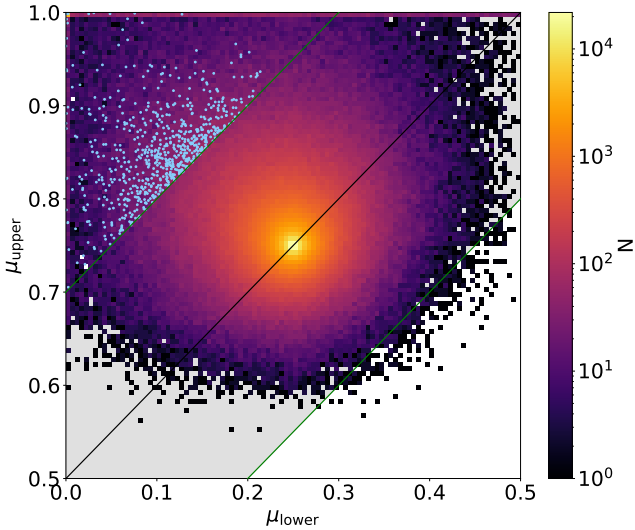


Fig. 5. CuDiDi distribution of the photon-arrival-time asymmetry, showing μ_{lower} vs. μ_{upper} for all sources we analysed (background-density map). The diagonal black line indicates the locus of temporally symmetric light curves, $\mu_{\text{upper}} = \mu_{\text{lower}} + 0.5$. The green lines mark the empirical persistent band, defined as $|\mu_{\text{upper}} - (\mu_{\text{lower}} + 0.5)| \leq 0.2$. Sources within and below this band are considered non-transient, while sources above are classified as transients. The 765 transient candidates in our golden sample are overplotted as light blue points. By construction they lie above the band, but their spread within this region reflects a wide range of asymmetry strengths. The bulk of the overall population clusters around (0.25, 0.75), as expected for approximately constant emission.

particular exposure regime and that detection efficiency is not strongly biased by exposure time.

To illustrate the identification of transients with the CuDiDi method, we applied it to all 750 169 light curves in our sample, with the resulting distribution shown in Figure 5. The majority of sources form a tight cluster centred near $(\mu_{\text{lower}}, \mu_{\text{upper}}) \approx (0.25, 0.75)$, consistent with constant or weakly variable emission. The figure also shows the boundaries used for defining transient and persistent sources.

The 765 sources that constitute our golden sample are overplotted in Figure 5 as light-blue points. Their concentration just above the upper edge of the band reflects the applied threshold rather than an intrinsic segregation of the population. Within the transient region, these sources span a wide range of asymmetries in their cumulative distributions, indicating diverse variability behaviours rather than a single characteristic profile. This spread suggests that CuDiDi captures a continuum of transient morphologies, from brief, sharply peaked events to more extended episodes of enhanced emission. The empirical cut therefore provides a practical, though somewhat arbitrary, means of separating transients from the dominant non-transient population. In contrast, relatively few sources are found below the persistent band. While this region could in principle contain sources with repeating or periodic flaring behaviour within a single observation, our classification boundaries were not tuned or validated for such cases, and we do not further explore this population in this work.

We next examined the brightness distribution of the transient candidates, based on the net background subtracted counts. Figure 6 shows a histogram spanning from 1 to over 200 counts, with a strong peak at the low-count end. The median count is 9, and the mean is 12. An inset zooms into the 0–50 count range, where the bulk of the population lies. The large number of low-count candidates reflects the method’s sensitivity to

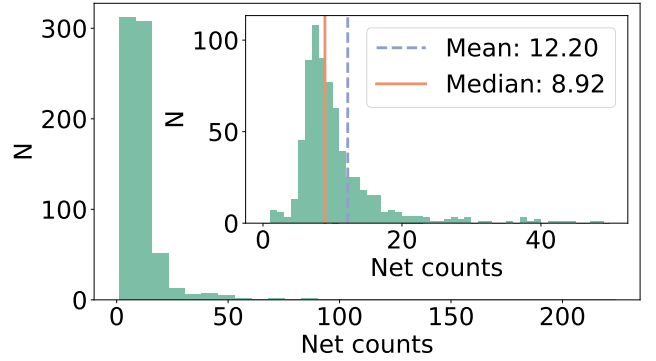


Fig. 6. Distribution of net counts for all 765 transient candidates. The histogram spans a wide range, with a mean of 12 and a median of 9 counts per candidate. The inset shows a zoomed-in view of the 0–50 count range to highlight the peak and shape of the low-count population. The prevalence of low-count events reflects the sensitivity of the method to faint or short-duration transients.

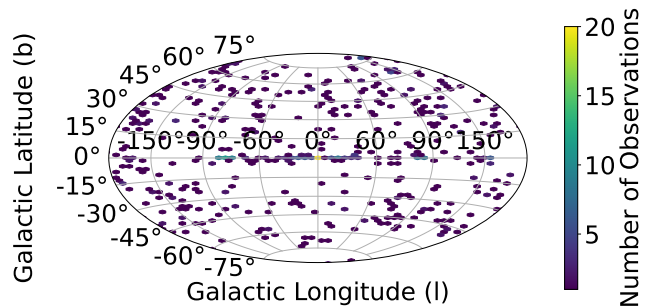


Fig. 7. Sky distribution of the 765 transient candidates in Galactic coordinates using a Hammer–Aitoff projection. The hexbin density map shows the number of candidates per bin across the sky.

faint or short-duration transients, which may not stand out in time-averaged images but exhibit clear temporal structure. At the same time it is clear that sources with a small number of counts are more likely to belong to the category of transients that exhibit occasional flares in addition to persistent faint emission, where the persistent emission would be below the *Chandra* sensitivity limit. Our check for recurring sources would, however, have excluded any such sources that have been observed to flare more than once.

3.2. Spatial distribution

If the transient candidates are predominantly extragalactic in origin, their sky distribution should be approximately isotropic, aside from modulations due to spatial variations in the total exposure time of all the *Chandra* observations. Figure 7 shows the distribution of all 765 candidates in Galactic coordinates. This shows that the transients are distributed across the whole sky, without any strong large-scale anisotropy. In particular, there is no strong concentration of candidates toward the Galactic plane, which would be expected if a substantial fraction were tied to Galactic stellar populations or compact objects. This suggests that many of the candidates may be extragalactic in nature, although definitive classification requires further multi-wavelength analysis.

3.3. Temporal properties

To quantify the characteristic duration of each transient, we measured the T_{90} interval, defined as the time over which 90% of

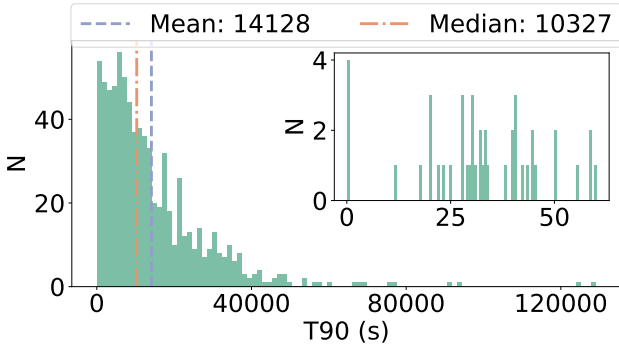


Fig. 8. Distribution of T_{90} durations (the time interval containing 90% of the total counts) for all 765 transient candidates. The zoomed inset highlights the short-duration regime ($T_{90} \leq 60$ s). The durations span from effectively zero to over 120 000 s, with a mean of approximately 14 000 s and a median of 10 000 s. Four candidates with $T_{90} = 0$ s are discussed in Sect. 4.1.

the total source counts are accumulated, between the epochs at which 5% (T_5) and 95% (T_{95}) of the counts are reached. A direct determination from the raw cumulative distributions was found to be unreliable in some cases due to statistical fluctuations and irregular light-curve shapes. We therefore fitted the cumulative total counts of each candidate with a sigmoid function of the form

$$f(t) = \frac{L}{1 + \exp[-k(t - t_0)]},$$

where L is the asymptotic total counts, t_0 is the midpoint, and k is the steepness parameter. The fit provides a smooth monotonic representation of the cumulative distribution, from which T_5 and T_{95} were determined as the times at which the fitted curve reaches 5% and 95% of L . The duration is then given by $T_{90} = T_{95} - T_5$. Figure 8 shows the resulting distribution, which spans over five orders of magnitude, from second durations to a substantial fraction of the longest *Chandra* observations. The distribution has a median of approximately 10,000 s and a mean of 14,000 s, indicating that most events are relatively long-lived but with substantial spread.

The inset zooms into the short-duration regime, revealing a substantial tail of fast events with $T_{90} < 60$ s. Notably, four candidates exhibit a measured $T_{90} = 0$ s. These extreme cases are discussed in more detail in Sect. 4.1, where we find that they are most likely genuine transients, while the $T_{90} = 0$ is an artefact of the sigmoid fit. The broad distribution of T_{90} values reflects the sensitivity of the method to rapid and slowly evolving events.

To assess the duration of each transient in the context of its observation, we compute the fractional timescale T_{90}/T_{obs} , where T_{obs} is the total exposure time. Figure 9 shows the distribution of this quantity across all 765 candidates. A prominent peak near zero reflects the population of short-duration events. The rest of the distribution is relatively broad, extending to $T_{90}/T_{\text{obs}} \sim 0.8$, indicating a significant subset of events that last throughout most of the exposure. This diversity in relative timescales suggests that the candidate sample includes multiple transient classes with differing physical origins.

3.4. Spectral properties

To characterise the spectral properties of our transient candidates, we calculated the hardness ratio (HR), a commonly used metric

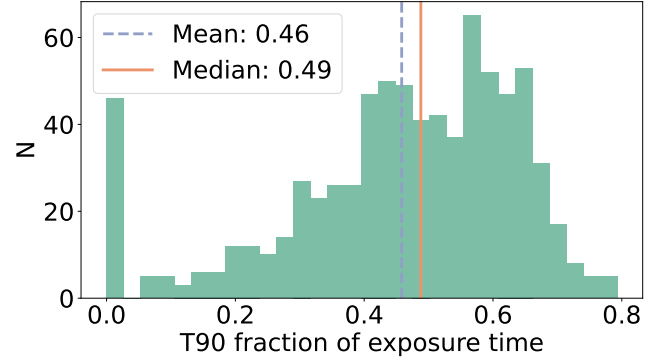


Fig. 9. Distribution of T_{90}/T_{obs} , the fractional duration of each transient relative to its parent observation. The sharp peak near zero reflects a population of very short transients, while the broader tail extending up to ~ 0.8 indicates events that persist for a substantial fraction of the exposure.

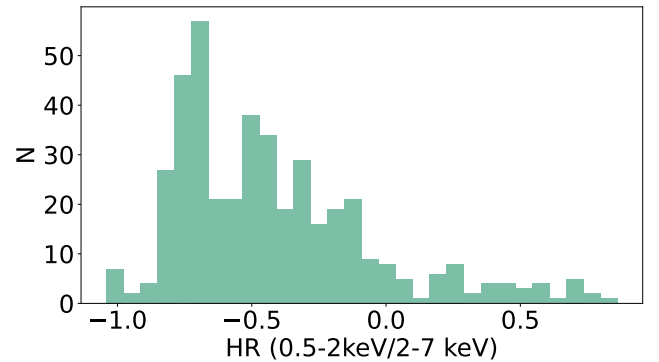


Fig. 10. Distribution of HRs for the transient candidates. The HR is defined as $\text{HR} = (H - S)/(H + S)$, where H and S are counts in the hard (2–7 keV) and soft (0.5–2 keV) bands, respectively. The distribution peaks toward negative values, indicating that the majority of candidates exhibit relatively soft X-ray spectra.

for low-count X-ray sources. The HR is defined as

$$\text{HR} = \frac{H - S}{H + S},$$

where H and S correspond to the number of photons detected in the hard (2–7 keV) and soft (0.5–2 keV) bands, respectively. We obtained valid HR measurements for 424 candidates. The remaining sources lack HR values due to empty background regions or errors encountered during spectral extraction, which prevented a reliable separation of counts into soft and hard bands. Figure 10 shows the distribution of hardness ratios for the 424 candidates. The clear peak toward negative HR values indicates that the majority of sources are dominated by soft X-ray emission, as expected since the source detection was performed on images in the soft band. We also investigated potential trends between HR and event duration (T_{90}), but found no significant correlations or evidence of clustering in the distribution. This is shown in Figure 11.

4. Discussion

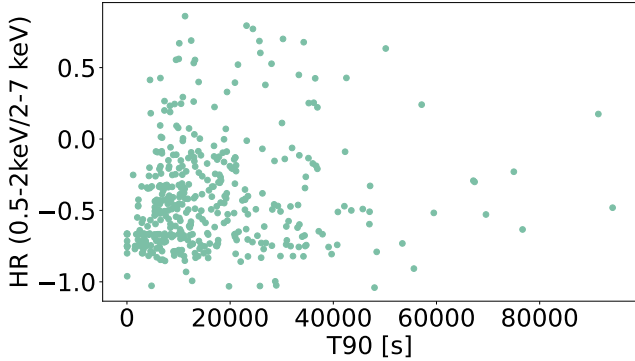
4.1. The $T_{90} = 0$ s population

Four candidates stand out in the duration distribution (Figure 8, inset), where they appear formally consistent with $T_{90} = 0$ s. This

Table 3. Selected candidates with zero-duration T_{90} .

Source ID	Obs. ID	RA (deg)	Dec (deg)	Obs. Time YYYY Mmm dd hh:mm:ss	Exp. Time (ks)	Notes
XT 040916	5024	161.161145	58.733960	2004 Sep. 16 06:54:50	67.5	Close to end of observation
XT 111205	13610	100.216376	9.538021	2011 Dec. 05 01:53:23	93.9	NGC 2264
XT 150429	17632	206.620123	26.423577	2015 Apr. 29 09:07:16	49.8	
XT 160625	18131	56.775122	-36.777104	2016 Jun. 25 01:10:04	4.9	

Notes. Each row includes the source ID, observation ID, sky position, observation start time, approximate exposure time and relevant notes.

**Fig. 11.** Distribution of the HR as a function of T_{90} for the transient candidates.

behaviour arises from the way durations are measured rather than from physical properties. Specifically, T_{90} is derived from a sigmoid fit to the cumulative photon arrival distribution (Sect. 3.3), with T_5 and T_{95} defined as the epochs when the fitted curve reaches 5% and 95% of the total counts. For short transients embedded within long observations, the fit becomes extremely steep, and the resulting T_{90} values collapse toward zero, even though the events span multiple seconds in real time. The light curves of these candidates (Figure 12) illustrate this point. In the zoomed-in insets, the photon arrival times span intervals from ~ 1 s to ~ 32 s. In the most extreme case, X-ray transient (XT) 160625, several photons are recorded within what appears as a 1 s interval, followed by additional isolated photons before and after the peak, suggesting that the true activity may extend over a longer period. We note, however, that the instrumental frame time in these four observations is 3.14 s. To enable consistent comparison across the dataset, all light curves were extracted with a uniform 1 s binning, which oversamples the intrinsic resolution for many observations but preserves the finest accessible temporal structure. This choice does not bias the transient selection (see Sect. 2.3) and allows flexible rebinning to coarser resolutions if desired.

The basic observational properties of the four candidates are listed in Table 3. Two of the events (XT 040916 and XT 160625) were detected under conditions that limit duration constraints – one near the end of a 67.5 ks observation, and one during a short 4.9 ks pointing. The other two (XT 111205 and XT 150429) were observed in the middle of long exposures (93.9 ks and 49.8 ks, respectively), making it unlikely that extended flaring activity was missed. In summary, the apparent $T_{90} = 0$ s values should not be interpreted as single-frame transients. Instead, they reflect limitations of the sigmoid-based estimator when applied to very short events in long exposures. Improved methods, for example by incorporating direct photon arrival distributions,

may yield more accurate duration measurements in the future. Nevertheless, these four transients still probe the extreme end of the timescale distribution for the transients identified in our search, and we therefore examine their properties in more detail below.

Before considering physical interpretations, we checked for potential instrumental origins. The four candidates are not concentrated near chip edges, nor are they located in regions with known extended sources (e.g. the vicinity of Cas A). Events of that kind would have been flagged by our diagnostic criteria and would not be included in the high-confidence (golden) sample. None of the candidates were flagged by the standard *Chandra* artefact diagnostics (e.g. streaks, hot pixels, cosmic-ray afterglows), which are designed to identify spurious detections of instrumental origin. We therefore consider it unlikely that these events arise from known detector effects. With instrumental origins unlikely, we can turn to the spectral properties of these events.

The spectra of the four transients were fitted with XSPEC version 12.14.1 (Arnaud 1996), using the C-stat fit statistic (Cash 1979). The spectra were grouped to one count per bin. Parameter uncertainties were obtained with the `error` command and correspond to 90% confidence intervals for one parameter of interest. The spectra are reasonably well fitted by blackbody models or soft power laws, as summarised in Table 4. The only exception is XT 160625, for which the parameters of the power-law model could not be constrained. The Galactic column density of hydrogen (N_{H}) in the direction of these transients is $7.61 \times 10^{19} \text{ cm}^{-2}$ (XT 040916), $4.43 \times 10^{21} \text{ cm}^{-2}$ (XT 111205), $1.00 \times 10^{20} \text{ cm}^{-2}$ (XT 150429), and $1.75 \times 10^{20} \text{ cm}^{-2}$ (XT 160625), determined from the `NHtot` tool⁹ (Willingale et al. 2013). These levels of absorption have a negligible impact on the fit results in all cases except XT 111205, for which we also report results that include absorption fixed at the Galactic value in Table 4, modelled with the XSPEC `tbabs` model (Wilms et al. 2000). The spectra of the four transients fitted with the blackbody model without absorption are shown in Figure 13.

The results in Table 4 highlight the very soft nature of the spectra, with kT in the range ~ 0.06 – 0.4 keV, or photon indices in the range $\Gamma \sim 2.5$ – 7 . This is also seen in Figure 13, which shows that XT 040916 is the only source with detected counts above 2 keV. For reference, we report fluxes from the best-fit models in the full *Chandra* energy range 0.5–7 keV in Table 4. XT 040916, XT 111205, and XT 150429 have similar unabsorbed fluxes in the range ~ 2 – $5 \times 10^{-15} \text{ erg s}^{-1} \text{ cm}^{-2}$, while XT 160625 is brighter at $\sim 7 \times 10^{-15} \text{ erg s}^{-1} \text{ cm}^{-2}$.

We carried out a search for host galaxies and possible Galactic counterparts in order to gain further insight into the nature of the sources, and ideally be able to determine their luminosities. The aim was to assess whether these events could be plausibly

⁹ <https://www.swift.ac.uk/analysis/nhtot/>

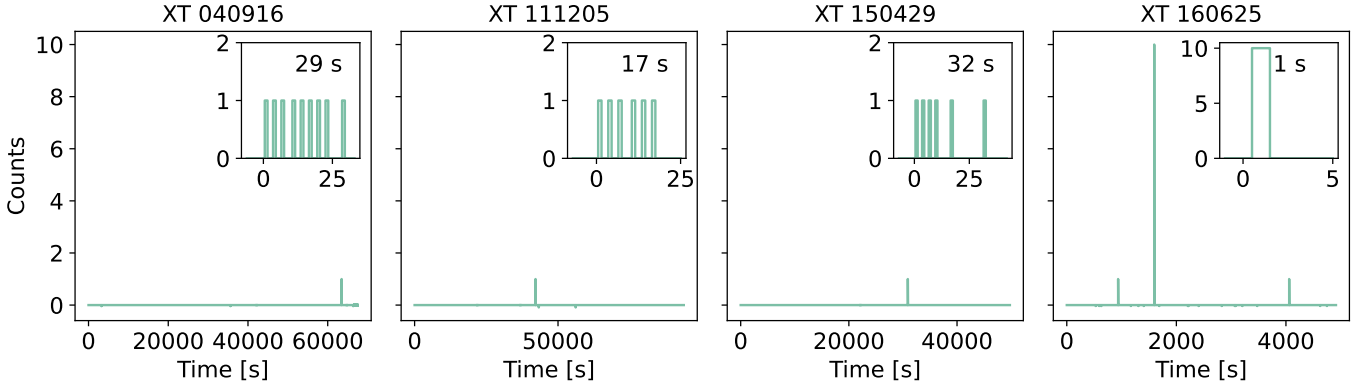


Fig. 12. Background-subtracted light curves of the transients with zero-duration T_{90} . The main panels show the full observation light curves, with the x -axes set such that 0 s corresponds to the start of the observation. The insets zoom in on the transient events, with x -axes shifted so that the first bin with signal corresponds to 0 s, and the durations of the events are indicated. Counts are binned to 1 s. These events last between ~ 1 and ~ 32 s.

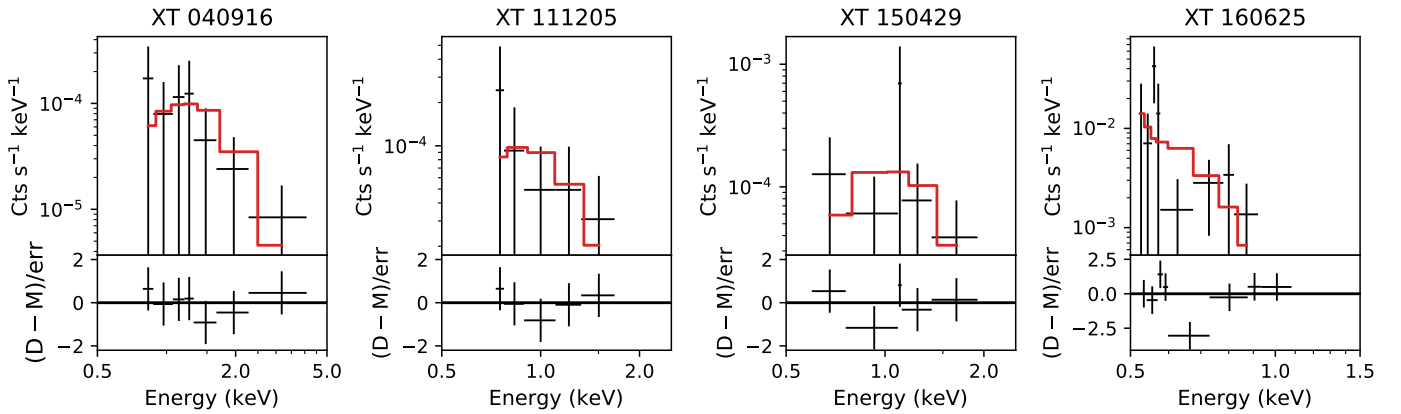


Fig. 13. Spectra of the transients with zero-duration T_{90} . The upper panels show the spectra (black crosses) and the best-fit blackbody models (red lines), while the lower panels show the residuals between the data (D) and models (M) normalised by the uncertainty. Note the different upper limits on the x -axes, set to cover the energy ranges where there are counts in the spectra.

associated with galaxies in their vicinity. To this end, visual inspections were performed using multi-wavelength imaging data. Sky views centred on each candidate were generated with the Aladin Lite (Baumann et al. 2022) interface across several optical, infrared, and ultraviolet surveys, including the Sloan digital sky survey (SDSS; Ahn et al. 2012), Pan-STARRS (Chambers et al. 2016), the two micron all sky survey (2MASS; Skrutskie et al. 2006), the wide-field infrared survey explorer (WISE; Wright et al. 2010), the digitized sky survey (DSS; Morrison 1995), and the galaxy evolution explorer (GALEX; Martin et al. 2005). Each region was examined at wide (3°) and narrow ($1'$) fields of view to capture the surrounding environment and the immediate vicinity of the candidate. These visual inspections were complemented by catalogue queries within a $1'$ radius around each candidate's position. Sources from major astronomical databases (Vizier, SIMBAD, and NED) were retrieved to identify nearby objects and assess potential host associations.

Among the four candidates, only XT 111205 shows a clear association, being located within NGC 2264, a well-studied star-forming region in the Milky Way at a distance of 722 pc (Flaccomio et al. 2023). XT 111205 is also the only candidate projected onto the Galactic plane. At the distance of NGC 2264, its 0.5–7 keV luminosity is $L_X = 1.2^{+1.9}_{-0.7} \times 10^{29}$ erg s $^{-1}$, where we have used the flux from the blackbody model without Galactic absorption in Table 4, considering the small inferred distance. For

the remaining three transients, no convincing host associations were identified.

The luminosity and spectrum of XT 111205 are consistent with the properties of X-ray flares from dwarf stars, which typically have soft spectra and peak luminosities below $\sim 10^{30}$ erg s $^{-1}$, but reaching $\sim 10^{32}$ erg s $^{-1}$ in extreme cases (Güdel & Nazé 2009; Robrade et al. 2010; Williams et al. 2014; Pye et al. 2015; De Luca et al. 2020). Such flares last $\sim 10^3$ – 10^4 s (Pye et al. 2015), longer than the 17 s duration observed for XT 111205, which may be attributed to the fact that the source is very faint and that we likely only detect the peak of the flare. The location and X-ray properties thus both favour the interpretation of XT 111205 as a stellar flare. Given the similar properties of the other three transients it is highly likely that they are also due to dwarf stars, although the lack of multi-wavelength counterparts and the locations away from the Galactic plane leave more ambiguity in the interpretation.

If extragalactic in nature, these transients may alternatively be interpreted as SBOs. In particular, a SBO from a blue supergiant is expected to produce a ~ 100 s transient with a soft spectrum with $kT \sim 0.3$ keV, assuming a non-relativistic spherically symmetric scenario (Waxman & Katz 2017), which is broadly consistent with the properties of XT 040916 and XT 150429. Given the observed fluxes, a typical luminosity of a few times 10^{44} erg s $^{-1}$ for the SBOs would then imply a redshift $z \sim 3$ for these sources. Deeper

Table 4. Results of spectral fits for candidates with zero-duration T_{90} .

Source ID	Blackbody model			Power law model		
	T (keV)	$F_{0.5-7 \text{ keV}}$ ($\text{erg s}^{-1} \text{cm}^{-2}$)	C-stat/d.o.f	Γ	$F_{0.5-7 \text{ keV}}$ ($\text{erg s}^{-1} \text{cm}^{-2}$)	C-stat/d.o.f
XT 040916	$0.40^{+0.38}_{-0.16}$	$2.5^{+2.5}_{-1.3} \times 10^{-15}$	1.84/7	$2.5^{+1.9}_{-1.6}$	$3.7^{+4.3}_{-1.9} \times 10^{-15}$	0.69/7
XT 111205	$0.17^{+0.27}_{-0.07}$	$1.9^{+3.0}_{-1.2} \times 10^{-15}$	1.39/4	$4.7^{+3.9}_{-3.6}$	$3.2^{+13.9}_{-2.1} \times 10^{-15}$	0.87/4
	$0.12^{+0.11}_{-0.04}$	$1.0^{+1.4}_{-0.8} \times 10^{-14}$	2.07/4	$7.0^{+3.0}_{-3.8}$	$2.9^{+14.4}_{-2.5} \times 10^{-14}$	1.25/4
XT 150429	$0.19^{+0.24}_{-0.07}$	$4.7^{+7.4}_{-2.9} \times 10^{-15}$	3.06/3	$4.0^{+2.8}_{-2.9}$	$6.9^{+9.8}_{-4.2} \times 10^{-15}$	2.8/3
XT 160625	$0.062^{+0.021}_{-0.015}$	$7.3^{+5.5}_{-3.5} \times 10^{-13}$	9.84/8	–	–	–

Notes. The photon index (Γ) of the power-law model could not be constrained for XT 160625. The second row for XT 111205 includes Galactic absorption fixed at $N_{\text{H}} = 4.43 \times 10^{21} \text{cm}^{-2}$, and the corresponding fluxes are unabsorbed. All other results are for models without Galactic absorption. The N_{H} of the other sources are in the range $0.76\text{--}1.75 \times 10^{20}$ and have a negligible effect on the fits.

follow-up observations are clearly needed in order to identify the stars and/or host galaxies associated with these transients.

4.2. Algorithm performance

Different approaches to transient searches with *Chandra* data make different assumptions about source selection, exposure properties, and detection thresholds, which strongly affect completeness and contamination. Many previous studies rely on the CSC (Evans et al. 2010, 2024) as their starting point (e.g. Quirola-Vásquez et al. 2022; Dillmann et al. 2025), which ensures well-characterised sources but will miss faint transients and recent events (the latest catalogue is complete until the end of 2021, Evans et al. 2024). Other works deliberately avoid catalogued sources, such as the single-frame search of Zhang & Feng (2023), which targets only events in regions without known point sources. The method we introduced is independent of pre-compiled source lists and applies only minimal pre-filtering (a source significance of 3σ). This makes it sensitive to catalogued and uncatalogued sources, across a wide range of sky regions, including the Galactic plane, and in fields that might otherwise be excluded.

Restrictions on exposure properties can also play an important role. For example, Yang et al. (2019) limited their analysis to sources within $8'$ off-axis angle and divided exposures longer than 50 ks into shorter segments to mitigate background accumulation. The method we introduced imposes no such cuts, allowing transients to be found in virtually all available observations. As a result, the algorithm is able to identify events on timescales ranging from single-frame spikes to longer-duration variability, although it is naturally restricted to sub-observation transients and primarily detects events lasting less than $\sim 70\%$ of the exposure time (Figure 9).

The detection thresholds for the sources and variability also vary widely between methods (Yang et al. 2019; Lin et al. 2022; Quirola-Vásquez et al. 2022, 2023; Zhang & Feng 2023; Dillmann et al. 2025). Algorithms such as `glvary` in CIAO assign a variability index to each source based on deviations from a steady light curve (Lin et al. 2022), which is powerful for detecting persistent variability but less sensitive to isolated, short-duration events. Our method instead relies on `wavdetect` for initial source identification. Contamination in our search arises mainly from statistical fluctuations in faint sources, instrumental artefacts, or background variations; the conservative choices in `sigthresh` and the CuDiDi threshold are designed to minimise these false positives.

Our subsequent cross-matching with external catalogues is also deliberately cautious, using a $3''$ radius, which further reduces contamination. We note, however, that this comes at the expense of a higher risks of excluding transients due to false associations (see Sect. 2.3). This step could be improved in the future by using the positional uncertainties of all the individual sources instead of a fixed matching radius. Finally, while the algorithm can detect events as short as a single frame, the T_{90} duration is derived from a sigmoid fit that is not always optimal for all light curves, especially for short transients embedded in very long exposures as discussed above.

The comparison with the transient samples of Quirola-Vásquez et al. (2022, 2023) demonstrates that our method is sensitive to previously reported transients, while applying stricter selection criteria. Out of the 22 published QV transients, our algorithm independently recovered 13, which were subsequently excluded from the catalogue (see Table A.1 in Appendix A for a detailed overview). Out of the remaining nine sources, six were excluded by the CuDiDi threshold, while three were rejected due to associations with known objects in SIMBAD or NED (one of which was also below the CuDiDi threshold).

We plot all 22 FXTs on the CuDiDi diagram in Figure 14, which shows that the majority of sources that fall below the threshold do so by a very small margin, illustrating that the CuDiDi cut is slightly more conservative than the method used in Quirola-Vásquez et al. (2022, 2023). The source furthest away from the transient region in Figure 14 is FXT 5, followed by FXT 13 and FXT 18. The light curve of FXT 5 (shown in Quirola-Vásquez et al. 2022, Fig. 6), exhibits two bursts of emission spanning $T_{90} = 32.8$ ks of the total 51.7 ks exposure, that is, $\sim 63\%$ of the observation, which is at the upper end of our sample (Figure 9). FXT 13 and 18 are both among the faintest of the FXT sample, with FXT 18 also having a $T_{90} =$ of 2.9 ks, which spans nearly half of the short 6.1 ks exposure time.

The recovery of more than half of the QV transients provides an independent validation of our approach. The somewhat stricter criteria help guard against contamination from spurious detections, but can of course also miss real transients that fall just below the threshold. A future development of the CuDiDi classifier could therefore be to use a lower threshold, combined with additional tests of variability for sources that fall just above the boundary. Further analysis of the golden sample, including host galaxy identifications, will be necessary to determine whether the transients uncovered in this search differ from those reported in previous studies.

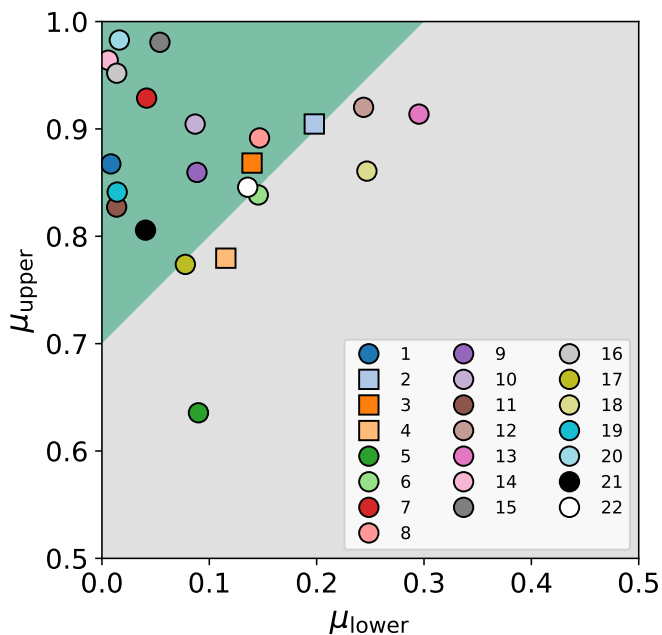


Fig. 14. CuDiDi diagram of the FXTs reported by Quirola-Vásquez et al. (2022) and Quirola-Vásquez et al. (2023), labelled by numbers. The squares indicate the three FXTs that were excluded by our cross-matching criterion. Sources in the shaded green region are classified as transients according to the CuDiDi threshold. Further details about the FXTs are provided in Table A.1.

5. Summary and conclusions

We have presented a search for sub-observation transients in the full public *Chandra* ACIS archive, covering 20 212 observations obtained before 2024-12-31. Our transient detection pipeline integrates automated source detection, light-curve extraction, diagnostic flagging, and the novel CuDiDi classifier, followed by visual screening to minimise contamination. Previously known sources were removed by cross-matching with other catalogues, ensuring that the resulting sample comprises single-occurrence transients, such as cataclysmic events, and new detections of flaring objects. The main outcomes of this work are summarised below:

1. From 1420 initial detections, we constructed a golden sample of 765 high-confidence transients, suitable for follow-up studies;
2. The sample spans a wide range of temporal properties, with T_{90} values covering five orders of magnitude. While most events lasted thousands of seconds, we also identified a tail of very fast events, with durations $\lesssim 30$ s. The transients are also spectrally diverse, with a wide range of hardness ratios;
3. A more detailed investigation of four fast candidates with durations $\lesssim 30$ s showed that their spectra are very soft. The candidates can be characterised as blackbodies with kT in the range $\sim 0.06\text{--}0.4$ keV. One of them (XT 111205) is associated with a star-forming region in the Galactic plane, implying a low luminosity of $\sim 10^{29}$ erg s $^{-1}$ cm $^{-2}$ and making it consistent with a flare from a dwarf star. The other three candidates lack multi-wavelength/host galaxy associations, but their properties are similar to those of XT 111205, which makes it likely that they are also flaring stars;
4. A comparison with the transients reported in Quirola-Vásquez et al. (2022, 2023) showed that CuDiDi performed well in

terms of identifying transients, while the threshold we used here imposed a somewhat stricter variability criterion.

Future work will include further analysis of the transient catalogue, including host galaxy identification and more detailed studies of the X-ray properties. The CuDiDi classifier will also be explored further, including its potential use as a tool for identifying repeating sources.

Data availability

The following data products are available on Zenodo: the final candidate catalogue: <https://doi.org/10.5281/zenodo.16356970>, the complete list of observation IDs analysed: <https://doi.org/10.5281/zenodo.17099991>, the types of objects excluded in the SIMBAD and NED searches: <https://doi.org/10.5281/zenodo.16759113>.

Acknowledgements. The authors thank Dennis Alp for his consultations and fruitful discussions. This research was supported by the Knut & Alice Wallenberg foundation. This research has used the SIMBAD database, operated at CDS, Strasbourg, France. This research has used the NASA/IPAC Extragalactic Database (NED), which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology. This research has made use of Aladin sky atlas developed at CDS, Strasbourg Observatory, France. This work has used the following software: Astropy 5.3.4 (Astropy Collaboration 2013, 2018, 2022), ds9 8.4.1, CIAO 4.15.2 with CALDB 4.10.7 (Fruscione et al. 2006), Heasoft ftools 22Aug2023_V6.32.1 (Blackburn et al. 1999) Matplotlib 3.8.1 (Hunter 2007), Numpy 1.26.0 (Harris et al. 2020), Pandas 2.2.0 (Wes McKinney 2010).

References

- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, *ApJS*, 203, 21
- Alp, D., & Larsson, J. 2020, *ApJ*, 896, 39
- Arnaud, K. A. 1996, in *Astronomical Society of the Pacific Conference Series*, 101, *Astronomical Data Analysis Software and Systems V*, eds. G. H. Jacoby, & J. Barnes, 17
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, 558, A33
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, 156, 123
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, *ApJ*, 935, 167
- Bauer, F. E., Treister, E., Schawinski, K., et al. 2017, *MNRAS*, 467, 4841
- Baumann, M., Boch, T., Pineau, F.-X., et al. 2022, in *Astronomical Society of the Pacific Conference Series*, 532, *Astronomical Data Analysis Software and Systems XXX*, eds. J. E. Ruiz, F. Pierfederici, & P. Teuben, 7
- Becerra, R. L., Yang, Y.-H., Troja, E., et al. 2026, *A&A*, 705, A233
- Blackburn, J. K., Shaw, R. A., Payne, H. E., Hayes, J. J. E., & Heasarc. 1999, FTOOLS: A general package of software to manipulate FITS files, Astrophysics Source Code Library [[record asc1:9912.002](https://arxiv.org/abs/record/asc1:9912.002)]
- Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, *Space Sci. Rev.*, 120, 165
- Cash, W. 1979, *ApJ*, 228, 939
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints [[arXiv:1612.05560](https://arxiv.org/abs/1612.05560)]
- De Luca, A., Stelzer, B., Burgasser, A. J., et al. 2020, *A&A*, 634, L13
- De Luca, A., Salvaterra, R., Belfiore, A., et al. 2021, *A&A*, 650, A167
- Dillmann, S., Martínez-Galarza, J. R., Soria, R., Stefano, R. D., & Kashyap, V. L. 2025, *MNRAS*, 537, 931
- Eappachen, D., Jonker, P. G., Fraser, M., et al. 2022, *MNRAS*, 514, 302
- Eappachen, D., Jonker, P. G., Levan, A. J., et al. 2023, *ApJ*, 948, 91
- Eappachen, D., Jonker, P. G., Quirola-Vásquez, J., et al. 2024, *MNRAS*, 527, 11823
- Eappachen, D., Balasubramanian, A., Swain, V., et al. 2026, *MNRAS*, 545, staf2062
- Evans, I. N., Primini, F. A., Glotfelty, K. J., et al. 2010, *ApJS*, 189, 37
- Evans, P. A., Page, K. L., Beardmore, A. P., et al. 2023, *MNRAS*, 518, 174
- Evans, I. N., Evans, J. D., Martínez-Galarza, J. R., et al. 2024, *ApJS*, 274, 22
- Eyles-Ferris, R. A. J., Starling, R. L. C., O’Brien, P. T., & Evans, P. A. 2022, *MNRAS*, 515, 4402
- Flaccomio, E., Micela, G., Peres, G., et al. 2023, *A&A*, 670, A37
- Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, *SPIE Conf. Ser.*, 6270, 62701V
- Gaia Collaboration (Prusti, T., et al.) 2016, *A&A*, 595, A1

- Gaia Collaboration (Vallenari, A., et al.) 2023, *A&A*, **674**, A1
- Galloway, D. K., in't Zand, J., Chenevez, J., et al. 2020, *ApJS*, **249**, 32
- Garmire, G. P., Bautz, M. W., Ford, P. G., Nousek, J. A., & Ricker, Jr., G. R. 2003, *SPIE Conf. Ser.*, **4851**, 28
- Grupe, D., Thomas, H. C., & Beuermann, K. 2001, *A&A*, **367**, 470
- Güdel, M., & Nazé, Y. 2009, *A&A Rev.*, **17**, 309
- Haberl, F., Xu, Y., Maitra, C., et al. 2025, *MNRAS*, **542**, 583
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Inkenhaag, A., Jonker, P. G., Leván, A. J., et al. 2024, *A&A*, **689**, A343
- Jansen, F., Lumb, D., Altieri, B., et al. 2001, *A&A*, **365**, L1
- Jiang, S.-Q., Xu, D., van Hoof, A. P. C., et al. 2025, *ApJ*, **988**, L34
- Jonker, P. G., Glennie, A., Heida, M., et al. 2013, *ApJ*, **779**, 14
- Jonker, P. G., Leván, A. J., Liu, X., et al. 2026, *MNRAS*, **545**, staf2021
- Kaspi, V. M., & Beloborodov, A. M. 2017, *ARA&A*, **55**, 261
- Khan, N., Quintin, E., Webb, N. A., et al. 2025, *A&A*, **697**, A50
- König, O., Wilms, J., Arcodia, R., et al. 2022, *Nature*, **605**, 248
- Kowalski, A. F. 2024, *Living Rev. Sol. Phys.*, **21**, 1
- Kumar, P., & Zhang, B. 2015, *Phys. Rep.*, **561**, 1
- Levan, A. J., Martín-Carrillo, A., Laskar, T., et al. 2025, *ApJ*, **990**, L28
- Levinson, A., & Nakar, E. 2020, *Phys. Rep.*, **866**, 1
- Li, D., Zhang, W., Yang, J., et al. 2025, arXiv e-prints [arXiv:2509.25877]
- Lin, D., Irwin, J. A., Berger, E., & Nguyen, R. 2022, *ApJ*, **927**, 211
- Liu, M., Zhang, Y., Wang, Y., et al. 2025a, *ApJ*, **984**, 5
- Liu, Y., Sun, H., Xu, D., et al. 2025b, *Nat. Astron.*, **9**, 564
- Maguire, K., Eracleous, M., Jonker, P. G., MacLeod, M., & Rosswog, S. 2020, *Space Sci. Rev.*, **216**, 39
- Martin, D. C., Fanson, J., Schiminovich, D., et al. 2005, *ApJ*, **619**, L1
- McCullough, M. L., & Rots, A. H. 2005, in *Astronomical Society of the Pacific Conference Series*, **347**, Astronomical Data Analysis Software and Systems XIV, eds. P. Shopbell, M. Britton, & R. Ebert, 478
- Merloni, A., Lamer, G., Liu, T., et al. 2024, *A&A*, **682**, A34
- Morrison, J. E. 1995, in *Astronomical Society of the Pacific Conference Series*, **77**, Astronomical Data Analysis Software and Systems IV, eds. R. A. Shaw, H. E. Payne, & J. J. E. Hayes, 179
- Novara, G., Esposito, P., Tiengo, A., et al. 2020, *ApJ*, **898**, 37
- Pastor-Marazuela, I., Webb, N. A., Wojtowicz, D. T., & van Leeuwen, J. 2020, *A&A*, **640**, A124
- Primini, F. A., Houck, J. C., Davis, J. E., et al. 2011, *ApJS*, **194**, 37
- Pye, J. P., Rosen, S., Fyfe, D., & Schröder, A. C. 2015, *A&A*, **581**, A28
- Quintin, E., Webb, N. A., Georgantopoulos, I., et al. 2024, *A&A*, **687**, A250
- Quirola-Vásquez, J., Bauer, F. E., Jonker, P. G., et al. 2024, *A&A*, **683**, A243
- Quirola-Vásquez, J., Bauer, F. E., Jonker, P. G., et al. 2022, *A&A*, **663**, A168
- Quirola-Vásquez, J., Bauer, F. E., Jonker, P. G., et al. 2023, *A&A*, **675**, A44
- Quirola-Vásquez, J., Bauer, F. E., Jonker, P. G., et al. 2025, *A&A*, **695**, A279
- Rastinejad, J. C., Leván, A. J., Jonker, P. G., et al. 2025, *ApJ*, **988**, L13
- Rea, N., Esposito, P. 2011, in *Astrophysics and Space Science Proceedings*, **21**, High-Energy Emission from Pulsars and their Systems, eds. D. F. Torres, & N. Rea, 247
- Robrade, J., Poppenhaeger, K., & Schmitt, J. H. M. M. 2010, *A&A*, **513**, A12
- Romano, P., Evans, P. A., Bozzo, E., et al. 2023, *A&A*, **670**, A127
- Rossi, E. M., Stone, N. C., Law-Smith, J. A. P., et al. 2021, *Space Sci. Rev.*, **217**, 40
- Ruiz, A., Georgakakis, A., Georgantopoulos, I., et al. 2024, *MNRAS*, **527**, 3674
- Sguera, V., Sidoli, L., Bird, A. J., Paizis, A., & Bazzano, A. 2020, *MNRAS*, **491**, 4543
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163
- Soderberg, A. M., Berger, E., Page, K. L., et al. 2008, *Nature*, **453**, 469
- Srinivasaragavan, G. P., Li, D., Hall, X. J., et al. 2025, arXiv e-prints [arXiv:2512.10239]
- Starrfield, S., Iliadis, C., & Hix, W. R. 2016, *PASP*, **128**, 051001
- Stathopoulos, S. I., Petropoulou, M., Giommi, P., et al. 2022, *MNRAS*, **510**, 4063
- Sun, H., Li, W. X., Liu, L. D., et al. 2025, *Nat. Astron.*, **9**, 1073
- van Dalen, J. N. D., Leván, A. J., Jonker, P. G., et al. 2025, *ApJ*, **982**, L47
- Waxman, E., & Katz, B. 2017, in *Handbook of Supernovae*, eds. A. W. Alsabti & P. Murdin (Springer International Publishing), 967
- Weisskopf, M. C., Tananbaum, H. D., Van Speybroeck, L. P., & O'Dell, S. L. 2000, *SPIE Conf. Ser.*, **4012**, 2
- Wenger, M., Ochsenein, F., Egret, D., et al. 2000, *A&AS*, **143**, 9
- Wes McKinney 2010, in *Proceedings of the 9th Python in Science Conference*, eds. S. van der Walt, & J. Millman, 56
- Williams, P. K. G., Cook, B. A., & Berger, E. 2014, *ApJ*, **785**, 9
- 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
- Woods, P. M., & Thompson, C. 2006, in *Compact stellar X-ray sources*, 39, eds. W. H. G. Lewin, & M. van der Klis (Cambridge University Press), 547
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868
- Xue, Y. Q., Zheng, X. C., Li, Y., et al. 2019, *Nature*, **568**, 198
- Yang, G., Brandt, W. N., Zhu, S. F., et al. 2019, *MNRAS*, **487**, 4721
- Yuan, W., Dai, L., Feng, H., et al. 2025, *Sci. China Phys. Mech. Astron.*, **68**, 239501
- Zhang, Y., & Feng, H. 2023, *MNRAS*, **523**, 2513
- Zhang, W., Yuan, W., Ling, Z., et al. 2025, *Sci. China Phys. Mech. Astron.*, **68**, 219511
- Zhao, G., Zheng, W., Shen, R.-F., et al. 2026, *J. High Energy Astrophys.*, **51**, 100537
- Zhou, H., Ren, J., Wang, C.-W., et al. 2026, *ApJ*, **998**, 163

Appendix A: Comparison with Quirola-Vásquez et al. (2022, 2023)

Table A.1 lists the transient candidates reported by Quirola-Vásquez et al. (2022, 2023), their *Chandra* observation IDs, positions, and whether the sources were recovered in our analysis. For non-recovered sources, the reason for exclusion is indicated.

Appendix B: Impact of background

The CuDiDi method used in this work is based on total cumulative counts without background subtraction. This is motivated by the generally very low background in *Chandra*, combined with the use of short 1 s bins to improve sensitivity to short timescale variability. Background subtraction in finely time resolved bins leads to bins with negative counts, which in turn results in a cumulative distribution that does not increase monotonously, for which the CuDiDi discriminator was not designed.

To explore the impact of background we take 10 transient candidates that were identified in observations with exposure times >100 ks, considering that the background is expected to be more important in long observations. We bin the background-subtracted light curves to 5 ks and set the count rate in any intervals that have negative counts to 0 when constructing the cumulative distributions. The results are shown in Figure B.1, where we also compare with the results of using 5 ks bins without background subtraction, and the original 1 s bins. As can be seen, the background subtraction has a small effect on the results and do not shift the points in any systematic way. In fact, Figure B.1 shows that changing the bins from 1 s to 5 ks generally has a larger impact on the positions in the CuDiDi diagram than the background subtraction.

These results indicate that the transient selection based on total cumulative counts is not strongly impacted by background. However, extending this method to other telescopes with higher background would most likely require some adjustments of the method. This could involve using wider background-subtracted bins as shown here, or simply adjusting the selection criteria in the CuDiDi plane to account for different background levels.

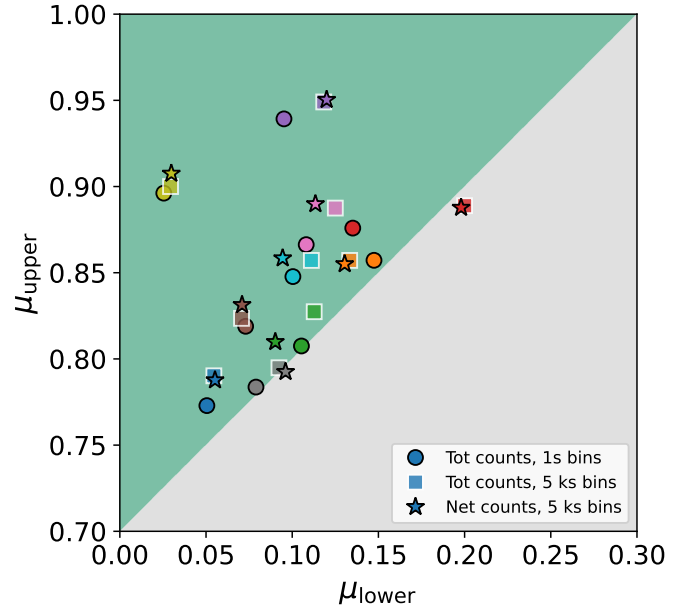


Fig. B.1. CuDiDi diagram of 10 transit candidates that were identified in observations with exposure times >100 ks. Each transient is shown with a different colour, where the different symbols correspond to different choices of binning and background subtraction as indicated in the legend. Only the upper left corner of the CuDiDi diagram is shown in order to highlight the differences between the different points.

Table A.1. Comparison with transient candidates reported by Quirola-Vásquez et al. (2022, 2023).

FXT	Obs. ID	RA (deg)	Dec (deg)	Status / Reason (this work)
Quirola-Vásquez et al. (2022)				
1	803	186.38125	13.06607	Recovered (removed from catalogue)
2	2025	167.86792	55.67253	Known HMXB (SIMBAD), (above CuDiDi threshold)
3	8490	201.24329	-43.04060	Planetary Nebula (SIMBAD/NED), (above CuDiDi threshold)
4	9546	211.25113	53.65706	Known HMXB (SIMBAD), below CuDiDi threshold
5	9548	170.07296	12.97189	Below CuDiDi threshold
6	14904	345.49250	15.94871	Below CuDiDi threshold
7	4062	76.77817	-31.86980	Recovered (removed from catalogue)
8	5885	318.12646	-63.49914	Recovered (removed from catalogue)
9	9841	175.00504	-31.91743	Recovered (removed from catalogue)
10	12264	90.00450	-52.71501	Recovered (removed from catalogue)
11	12884	212.12063	-27.05784	Recovered (removed from catalogue)
12	13454	15.93558	-21.81272	Below CuDiDi threshold
13	15113	45.26725	-77.88095	Below CuDiDi threshold
14	16454	53.16158	-27.85940	Recovered (removed from catalogue)
Quirola-Vásquez et al. (2023)				
15	16093	233.73496	23.46849	Recovered (removed from catalogue)
16	16453	53.07672	-27.87345	Recovered (removed from catalogue)
17	18715	40.82972	32.32390	Below CuDiDi threshold
18	19310	36.71489	-1.08317	Below CuDiDi threshold
19	20635	356.26437	-42.64494	Recovered (removed from catalogue)
20	21831	207.34711	26.58421	Recovered (removed from catalogue)
21	23103	50.47516	41.24704	Recovered (removed from catalogue)
22	24604	207.23523	26.66230	Recovered (removed from catalogue)

Notes. Listed are the QV designation (FXT, followed by a number), *Chandra* observation ID, position, and whether the source was recovered in our analysis. For non-recovered sources, the reason is indicated.