

The XMM-Newton serendipitous survey[★]

IX. The fourth XMM-Newton serendipitous source catalogue

N. A. Webb¹, M. Coriat¹, I. Traulsen², J. Ballet³, C. Motch⁴, F. J. Carrera⁵, F. Koliopanos¹, J. Authier¹, I. de la Calle¹¹, M. T. Ceballos⁵, E. Colomo¹¹, D. Chuard³, M. Freyberg⁶, T. Garcia¹, M. Kolehmainen⁴, G. Lamer², D. Lin⁷, P. Maggi⁴, L. Michel⁴, C. G. Page⁹, M. J. Page⁸, J. V. Perea-Calderon¹⁰, F.-X. Pineau⁴, P. Rodriguez¹¹, S.R. Rosen¹¹, M. Santos Lleo¹¹, R. D. Saxton¹¹, A. Schwoppe², L. Tomás¹¹, M. G. Watson⁹, and A. Zakardjian¹

¹ IRAP, Université de Toulouse, CNRS, CNES, Toulouse, France
e-mail: Natalie.Webb@irap.omp.eu

² Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

³ AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Diderot, Sorbonne Paris Cité, F-91191 Gif-sur-Yvette, France

⁴ Université de Strasbourg, CNRS, Observatoire astronomique de Strasbourg, UMR 7550, 67000 Strasbourg, France

⁵ Instituto de Física de Cantabria (CSIC-UC), Avenida de los Castros, 39005 Santander, Spain

⁶ Max-Planck-Institut für extraterrestrische Physik, Giessenbachstraße 1, 85748 Garching, Germany

⁷ Space Science Center, University of New Hampshire, Durham, NH, USA

⁸ Mullard Space Science Laboratory, University College London, Holbury St Mary, Dorking, Surrey RH5 6NT, UK

⁹ Department of Physics & Astronomy, University of Leicester, Leicester, LE1 7RH, UK

¹⁰ RHEA for ESA/ESAC. European Space Astronomy Center (ESAC-ESA). Madrid. 28691. Spain

¹¹ ESAC, European Space Astronomy Center (ESAC-ESA), Madrid 28691, Spain

Received , accepted

ABSTRACT

Context. Sky surveys produce enormous quantities of data on extensive regions of the sky. The easiest way to access this information is through catalogues of standardised data products. *XMM-Newton* has been surveying the sky in the X-ray, ultra-violet and optical bands for 20 years.

Aims. The *XMM-Newton Survey Science Centre* has been producing standardised data products and catalogues to facilitate access to the serendipitous X-ray sky.

Methods. Using improved calibration and enhanced software we re-reduced all of the 14041 *XMM-Newton* X-ray observations. 11204 observations contained useful data and with these we created a new, high quality version of the *XMM-Newton* serendipitous source catalogue, 4XMM-DR9.

Results. 4XMM-DR9 contains 810795 detections of 550124 unique sources, covering 1152 degrees² (2.85%) of the sky. 336 columns of information on each detection are provided, along with images. The quality of the source detection is shown to have improved significantly with respect to previous versions of the catalogues. Spectra and lightcurves are also made available for more than 288000 of the brightest sources (36% of all detections). We also provide a catalogue containing stacked observations, 4XMM-DR9s, where at least two observations overlap. This method has the advantage of reaching fainter fluxes than for single observations and therefore provides additional detections with respect to the same fields in the catalogue of detections. It also provides long term variability of the detected sources.

Key words. Catalogs – Astronomical data bases – Surveys – X-rays: general

1. Introduction

The sky is constantly being surveyed by many different telescopes exploiting the full range of the electromagnetic spectrum, in addition to gravitational wave, neutrino and cosmic ray observatories. Each observation can provide a clue to the nature of the source and the physical processes underway. In addition, many objects are known to be highly variable in time, requiring many observations to fully understand the nature of the variability. Whilst dedicated observations can be necessary to answer some science questions, frequently, catalogues can provide the required information. Catalogues can also provide homogeneous

datasets for classes of objects and can also reveal previously unknown objects.

The *European Space Agency's* second cornerstone mission from the *Horizon 2000 programme*, *XMM-Newton* (Jansen et al. 2001) was launched twenty years ago on 10th December 1999. It has the largest effective area of any X-ray satellite (Ebrero 2019), thanks to the three X-ray telescopes aboard, each with ~1500 cm² of geometric effective area. This fact, coupled with the large field of view (FOV) of 30' diameter, means that a single pointing on average detects 50 to 100 serendipitous X-ray sources (Watson et al. 2009). The *XMM-Newton Survey Science Centre*¹ (SSC), a consortium of ten European Institutes (Watson et al. 2001), has developed much of the *XMM-Newton Science Analy-*

[★] Based on observations obtained with XMM-Newton, an ESA science mission with instruments and contributions directly funded by ESA Member States and NASA.

¹ <http://xmmssc.irap.omp.eu/>

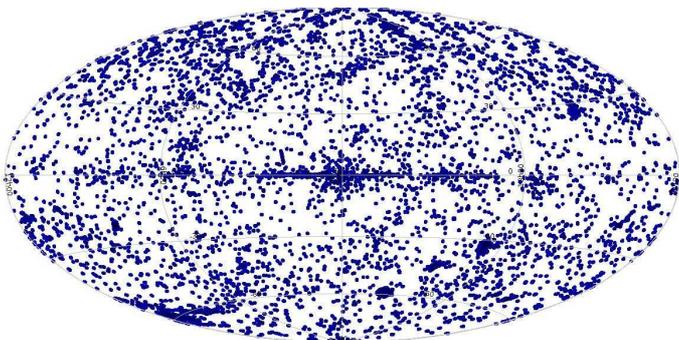


Fig. 1. Hammer-Aitoff equal area projection in Galactic coordinates of the 11204 4XMM-DR9 fields.

sis System (SAS) (Gabriel et al. 2004) for reducing and analysing *XMM-Newton* data and created pipelines to perform standardised routine processing of the *XMM-Newton* science data. The XMM-SSC also produces catalogues of all of the detections made with *XMM-Newton*. The catalogues of X-ray detections made with the three EPIC (Strüder et al. 2001; Turner et al. 2001) cameras that are placed at the focal point of the three X-ray telescopes have been designated 1XMM, 2XMM and 3XMM (Watson et al. 2009), with incremental versions of these catalogues indicated by successive data releases, denoted -DR in association with the catalogue number. This paper presents the latest version of the XMM catalogue, 4XMM, which spans 19 years of observations made with *XMM-Newton* and includes many improvements with respect to previous *XMM-Newton* catalogues. The most notable change between 3XMM and 4XMM is the methodology used for background modelling (see Sec. 3.4).

2. Catalogue observations

4XMM-DR9 is made from data drawn from 11204 *XMM-Newton* EPIC observations that were publicly available as of 1st March 2019 and that processed normally. The repartition of data modes for each camera and observation can be found in Table 1. The Hammer-Aitoff equal area projection in Galactic coordinates of the 4XMM-DR9 fields can be seen in Fig. 1. All of those observations containing > 1 ks clean data (> 1 ks of good time interval) were retained for the catalogue. Fig. 2 shows the distribution of total good exposure time (after event filtering) for the observations included in the 4XMM-DR9 catalogue and using any of the *thick*, *medium* or *thin* filters, but not the *open* filter. Open filter data were processed but not used in the source detection stage of pipeline processing. The same *XMM-Newton* data modes were used as in 2XMM (Watson et al. 2009) and are included in Table 2 of this paper, for convenience. The data in 4XMM-DR9 include 322 observations that were publicly available at the time of creating 3XMM-DR8, but were not included in that version due to high background or processing problems. Due to changes in the pipeline and in the background modelling, these problems have been overcome and thus the data could be included in 4XMM-DR9.

3. Data processing

Data processing for the 4XMM-DR9 catalogue was based on the SAS version 18 and carried out with the pipeline version 18² and

² <https://www.cosmos.esa.int/web/xmm-newton/pipeline-configurations>

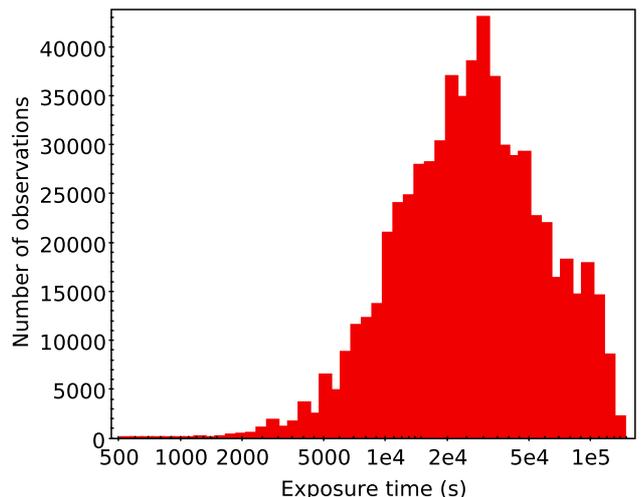


Fig. 2. Distribution of MOS 1 good exposure time (after event filtering) for the observations included in the 4XMM-DR9 catalogue.

Table 2. Data modes of XMM-Newton exposures included in the 4XMM catalogue.

Abbr.	Designation	Description
<i>MOS cameras:</i>		
PFW	Prime Full Window	covering full FOV
PPW2	Prime Partial W2	small central window
PPW3	Prime Partial W3	large central window
PPW4	Prime Partial W4	small central window
PPW5	Prime Partial W5	large central window
FU	Fast Uncompressed	central CCD in timing mode
RFS	Prime Partial RFS	central CCD with different frame time ('Refreshed Frame Store')
<i>pn camera:</i>		
PFWE	Prime Full Window Extended	covering full FOV
PFW	Prime Full Window	covering full FOV
PLW	Prime Large Window	half the height of PFW/PFWE

the latest set of current calibration files at the time of processing (February/March 2019).

The main data processing steps used to produce the 4XMM data products were similar to those outlined in Rosen et al. (2016); Watson et al. (2009) and described on the SOC web-pages³. For all the 4XMM data, the observation data files were processed to produce calibrated event lists. The optimised background time intervals were identified and using them, the “useful” exposures (taking into account exposure time, instrument mode, etc.); multi-energy-band X-ray images and exposure maps were generated. For the catalogue of detections (4XMM-DR9), source detection and parameterisation followed before cross-correlation of the source list with a variety of archival catalogues, image databases and other archival resources. The creation of spectra and light curves for the brightest sources followed. Automatic and visual screening procedures were carried out to check for any problems in the data products.

For the stacked catalogue (4XMM-DR9s), the ensuing procedure is described in Traulsen et al. (sub.). It is very similar to

³ https://xmm-tools.cosmos.esa.int/external/xmm/_products/pipeline/doc/17_40_20181123_1545/modules/index.html

Table 1. Characteristics of the 11204 XMM-Newton observations included in the 4XMM-DR9 catalogue.

Camera	Modes			Filters			Total
	Full ^a	Window ^b	Other ^c	Thin	Medium	Thick	
pn	8462	683	1344	5640	4011	838	10489
MOS1	8681	1950	373	5080	4943	981	11004
MOS2	8728	1981	348	5120	4971	966	11057

^a Prime Full Window Extended (PFWE) and Prime Full Window (PFW) modes; ^b pn Prime Large Window (PLW) mode and any of the various MOS Prime Partial Window (PPW) modes; ^c other MOS modes (Fast Uncompressed (FU), Refresh Frame Store (RFS)).

the procedure used to produce the catalogue of detections, except that the source detection is carried out on all overlapping observations simultaneously and the results per input image are stored in an intermediate source list, before determining the final source list and parameterisation, see also Sec. 7. The catalogue is cross-matched with 4XMM-DR9, and long-term lightcurves and auxiliary images of the sources are produced.

The data from this processing have been made available through the XMM-Newton Science Archive⁴ (XSA), but see also Sec. 10.

3.1. Exposure selection

The same criteria used for selecting exposures for 3XMM were retained for 4XMM. Almost 50% more exposures (or 10.7 Gs) were available for 4XMM-DR9 compared to 3XMM-DR5.

3.2. Event list processing

Much of the pipeline processing that converts raw ODF event file data from the EPIC instruments into cleaned event lists has remained unchanged from the pre-cat9.0 pipeline and is described in section 4.2 of Watson et al. (2009). A number of improvements have been made since the 2XMM (Watson et al. 2009) and 3XMM (Rosen et al. 2016) catalogues, which can be found in the SAS release notes⁵. These include source spectra and light curves created for pn Timing mode and small window data, source detection on pn small window data, energy dependent Charge Transfer Inefficiencies (CTI) and double event energy corrections applied, time and pattern dependent corrections of the spectral energy resolution of pn data, X-ray loading and rate dependent energy (PHA) and CTI corrections for EPIC pn Timing and Burst modes, binning of MOS spectra changed from 15 eV to 5 eV and filtering with XMMEA_EM, which is a bit-wise selection expression, automatically removing “bad events” such as bad rows, edge effects, spoiled frames, cosmic ray events (MIPs), diagonal events, event beyond threshold, etc, instead of XMMEA_SM (which removed all flagged events except those flagged only as CLOSE_TO_DEADPIX). Other modifications are background regions for EPIC spectra and light curves selected from the same EPIC chip where the source is found, observations of solar system objects processed such that X-ray images and spectra correctly refer to the moving target, pileup diagnostic numbers for EPIC sources included (see also Sec. 6.4.1), and footprints for EPIC observations based on combined EPIC exposure maps provided as ds9 region files. Other changes carried out specifically for the production of 4XMM include a revised systematic position error (see Sec. 3.3), the modelling of the EPIC background (see Sec. 3.4) and finer binning of EPIC

lightcurves (see Sec. 4.1). A small rotation of $\sim 0.4^\circ$ was noted in 3XMM fields, but analysis of 4XMM data shows that the recent improvements to calibration have resolved this issue. Below we describe some of the more recent developments specifically implemented for 4XMM.

3.3. Systematic position error

The astrometry of the X-ray detections is improved by using the *catcorr* task to cross-correlate the X-ray detections with the USNO B1.0, 2MASS or SDSS (DR8) optical/IR catalogues. Using pairs of X-ray and optical/infrared detections that fall within $10''$ of each other, the astrometry for the field is corrected using a translational shift in the right ascension (RA) and declination (DEC) directions, together with the rotational error component. A systematic error on the position (SYSERRRCC) is then calculated using the 1σ errors on the shifts in the RA ($\Delta\alpha_{error}$) and DEC ($\Delta\delta_{error}$) directions and the rotational error component in radians ($\Delta\theta_{error}$), derived from the catalogue that yields the ‘best’ solution, using $SYSERRRCC = \sqrt{(\Delta\alpha_{error}^2 + \Delta\delta_{error}^2 + (r * \Delta\theta_{error})^2)}$, where r is the radial off-axis angle of the detection from the spacecraft boresight in arcsecs. However, where *catcorr* fails to obtain a statistically reliable result (*poscorrok*=false), a systematic error of $1.5''$ was used to create the 3XMM catalogue.

Following analysis of the astrometry in 3XMM, a systematic error of $1.5''$ was shown to be too high. In order to determine an improved systematic error, we identified fields in 3XMM-DR8 where *catcorr* failed. We used sources from the *Sloan Digital Sky Survey Data Release 12* quasar (SDSS DR12 QSO) catalogue (Pâris et al. 2017) with good quality spectra (ZWARNING=0) and point-like morphology (SDSS_MORPHO=0). To avoid mismatches between targets and matched photometry⁶ we chose non-empty OBJ_ID values. We then cross-matched with the SDSS DR9 photometry catalogue (Ahn et al. 2012) in Vizier⁷ with a maximum distance of $5''$. This step provided the uncertainty in the astrometric position of SDSS. We adopted the radially-averaged uncertainty in the SDSS positions to which we had already added a systematic $0.1''$ in quadrature, $\Delta S = \sqrt{(\Delta\alpha^2 + \Delta\delta^2)/2 + 0.1^2}$. We then discarded all quasars with more than one SDSS DR9 counterpart within that distance. Out of the 256107 “clean” quasars, we selected the potential counterparts to the 3XMM DR8 sources, but also discarded those which could be counterparts of more than one 3XMM DR8 source. We used the “slim” catalogue for this purpose, since multiple detections of the same physical source appear only once. We cross-matched the SDSS DR9 positions of “clean” QSOs with the positions of the sources in the slim catalogue out to a distance of $r = 30''$. For each of the resulting pairs we estimated the combined positional error as $\sigma = \sqrt{\Delta S^2 + \Delta X^2/2}$, where $\Delta X \equiv SC_POSERR$

⁴ <https://nxsas.esac.esa.int/nxsas-web>

⁵ <https://www.cosmos.esa.int/web/xmm-newton/sas-release-notes/>

⁶ see <https://www.sdss.org/dr12/algorithms/match/>

⁷ <http://cdsarc.u-strasbg.fr/viz-bin/cat/V/139>

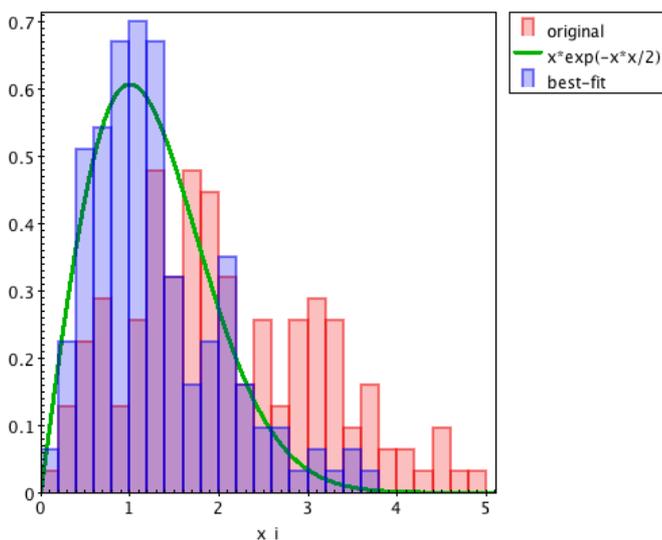


Fig. 3. 157 XMM-Newton-SDSS quasar pairs as a function of normalised distance x before adding a systematic uncertainty (red histogram) and after its addition (blue histogram), along with the Rayleigh distribution (green solid line).

and discarded all quasars that had more than one counterpart out to $r/\sigma = 6$, leaving 7205 suitable QSO (there were 26 QSO with more than one counterpart out to that limit). There were no pairs of quasars that corresponded to the same X-ray source.

Since each instance of an X-ray source in the 3XMM DR8 detection catalogue is an independent measurement, we cross-matched the sample of suitable quasars with the detection catalogue where `poscorrok=false`, out to $r = 30''$ again, filtering the latter with `SUM_FLAG=0` and `EP_EXTENT=0`, to keep only the cleanest sample of secure point-like X-ray sources. At this point we have 178 quasar-X-ray detection pairs. Similarly to that which was done for the slim catalogue, we define the combined positional error as $\sigma = \sqrt{\Delta S^2 + \Delta X^2/2}$, where $\Delta X \equiv \text{POSERR}$, and $x = r/\sigma$. Our final filtering retained only the 157 QSO-X-ray pairs with $x < 5$.

The expected probability density distribution of x should follow the Rayleigh distribution $P(x) = xe^{-x^2}$. Since this was not the case for the 157 pairs of sources found above, we added an additional positional uncertainty, Σ , in quadrature, so that the total positional uncertainty is now $\sigma' = \sqrt{\sigma^2 + \Sigma^2}$, looking for the value of Σ that minimizes the difference between the distribution of the $x' \equiv r/\sigma'$ and the Rayleigh distribution. We found $\Sigma = 1.29 \pm 0.01''$, where the uncertainty (1σ) has been calculated by bootstrap with replacement. The improvement can be seen in Figure 3. This value was then used to replace the $1.5''$ systematic error when `poscorrok=false`.

3.4. Modelling the EPIC background

For each input image to the source detection, the background is modelled by an adaptive smoothing technique. The method was initially applied to the proto-catalogue from overlapping XMM-Newton observations and described by Traulsen et al. (2019). Since the proto-catalogue was based on a selection of clean observations, the smoothing parameters were revised for the 4XMM catalogues, which cover observations of all qualities. The three parameters of the smoothing task are the cut-out radius to excise sources, the minimum kernel radius of the adaptive smoothing, and the requested signal-to-noise ratio in the

map. Their best values were determined in a three-fold assessment which involved real observations, randomized images, and visual screening.

656 observations were chosen which cover positions of cluster candidates identified by Takey et al. (2013) to involve a considerable number of extended and of point-like sources. Their background was modelled using different combinations of the smoothing parameters, and source detection was performed. The number of detections and recovered clusters, and the source parameters of the clusters and point-like detections were compared, opting for a reasonable compromise between total number of detections and potentially spurious content and for reliable fluxes and extent radius of the clusters. The source parameters of point-like detections were largely unaffected by the parameter choice in the tested parameter range.

The optimisation was then re-run on ninety observations, in which the background was replaced by a Poissonian randomisation. Finally, the two best combinations of smoothing parameters and the previously used spline fit were compared in a blind test. The detection images were inspected in randomised order, so the screeners could not know which source-detection results were based on which background model. The three parts of the assessment confirmed the preference for the adaptive smoothing approach over a spline fit and the final parameters: a brightness threshold for the source cut-out radius of 2×10^{-4} counts arcsec $^{-2}$, a minimum smoothing radius of 10 pixels ($40''$ in default image binning), and a signal-to-noise ratio of 12.

3.5. Updated flagging procedures

A single change to the flags provided for each detection has been introduced. Flag 12 now indicates if the detection falls on a region of the detector that can show hot pixels that can be misinterpreted as a source. Further information is provided in Section 3.5.1.

3.5.1. Hot areas in the detector plane

Warm pixels on a CCD (at a few counts per exposure) are too faint to be detected as such by the automatic processing, but can either push faint sources above detection level, or create spurious sources when combined with statistical fluctuations. This is an intrinsically random process, not visible over a short period of time, but which creates hot areas when projecting all sources detected over 18 years onto the detector plane.

We addressed this by projecting all sources onto CCD coordinates `PN/M1/M2_RAWX/Y`, keeping only sources above the detection threshold with the current instrument alone. In that way, we can distinguish hot areas coming from different instruments, see Figure 4.

We proceeded to detect hot pixels or columns in each CCD, using a similar method to the SAS task `embadpixfind`. Because the localization precision of faint sources is several arcseconds (larger than the MOS pixel size of $1.1''$) the detection was carried out for MOS after binning the image to 3×3 pixels (and testing all 9 single-pixel shifts). We flagged hot pixels with a probability less than $10^{-2}/N_{\text{trials}}$ to be compatible with a Poisson distribution at the local average (estimated from the local median plus 1). The trial factor N_{trials} was set to the image size (64×200) for PN and three times the binned image size (3×200^2) for MOS, accounting approximately for the fact that the shifted binned images are correlated.

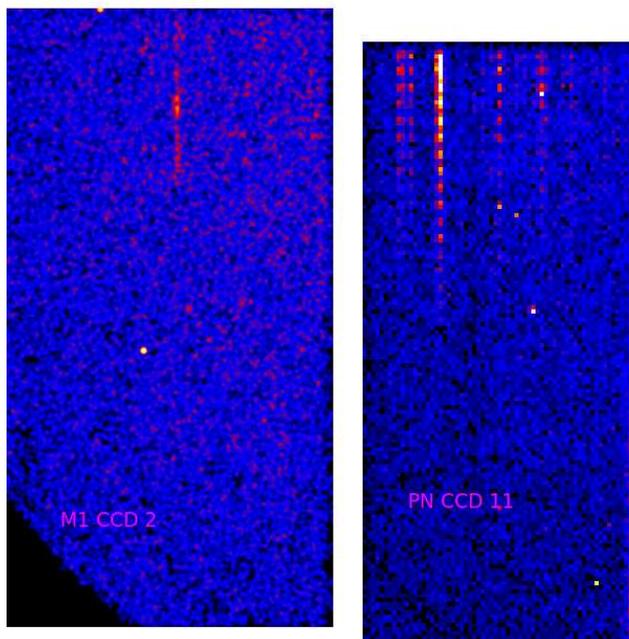


Fig. 4. Same area of the detector (lower left in PN detector coordinates) viewed by PN (CCD 11) and MOS1 (CCD 2). The maps are in CCD coordinates, but offset and zoomed so that they are approximately aligned. All point-like 4XMM sources with $\log(\text{likelihood}) > 6.5$ in the total band for the current instrument are accumulated on each map. The MOS1 map is smoothed with a 3×3 boxcar average. The colour scale is square root between 0 and 3 sources/pixel in MOS1, 0 and 100 in PN. Obvious hot areas are visible. MOS2 is omitted because it shows no hot area in that region.

Hot columns are detected in the same way after projecting the images (with hot pixels masked) onto *RAWX*. A column was considered bright when it was too high at the 7σ level applying the likelihood ratio test for Poisson counts (Li & Ma 1983) with respect to its surroundings (excluding immediate neighbours). This very high threshold was chosen such that subtle increases not obvious by eye were not detected (there are hundreds of sources per column, so that method is very sensitive).

It often occurs that only a piece of a column is bright. In order to identify such occurrences we compared the distribution of sources along *RAWY* in the hot column with that in the same neighbouring columns used in the column detection, using the Kolmogorov-Smirnov (hereafter KS) test. If the probability of compatibility was less than 10^{-4} , we looked for the bright interval with repeated KS tests on restricted lengths on each side of the *RAWY* value where the maximum distance between the two distributions occur, until we reached a probability of compatibility larger than 10^{-2} . The remainder on each side was considered normal or hot depending on the result of a Li & Ma test at the 3σ level with respect to the neighbouring columns.

We defined contiguous hot areas after reprojecting all the hot pixels and segments of columns onto the CCD (at the full pixel resolution for MOS). Many of those warm pixels were not present at the beginning of the mission, and some appear for a short amount of time. So we tested each hot area for variability using revolution number, and the same KS-based algorithm used to detect segments of bright columns, compared to the reference established over all sources on all CCDs and all instruments. This resulted in a revolution interval for each hot area. The distribution of remaining sources is shown in Figure 5.

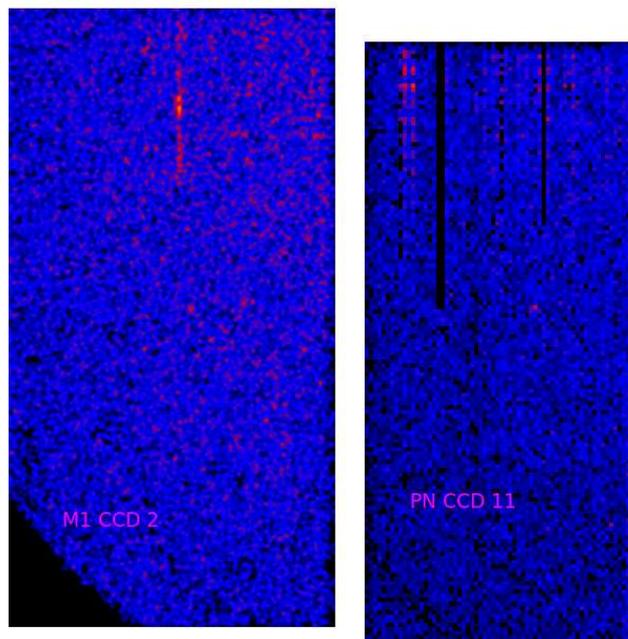


Fig. 5. Same area of the detector and color scale as in Figure 4. Sources on a hot area and inside the associated revolution interval are rejected. The remaining features cannot be distinguished from statistical fluctuations with the current algorithm.

Sources on top of a hot area for a particular instrument and within the corresponding revolution interval are flagged with flag 12. PN hot areas result in 16,503 flagged sources, MOS1 in 6,245 and MOS2 in 1,382 flagged sources.

4. Source-specific product generation

In order to minimise any contribution from soft proton flares, Good Time Interval (GTI) filtering is carried out. The light curves of the background data used to establish the count rate threshold for excluding background flares are extracted in the 0.5-7.5 keV energy range, with the full procedure described in Rosen et al. (2016). It is optimised for each detection where possible, but if no lightcurve can be generated, a general filtering for the observation is carried out. Image data are extracted from events using GTIs determined from when the pointing direction is within $3'$ of the nominal pointing position for the observation.

Following the source detection process, detections identified with at least 100 EPIC counts have their spectra extracted. If the number of counts not flagged as 'bad' (in the sense adopted by *XSPEC*) is still greater than 100 counts, a spectrum and a time series are extracted using an aperture around the source whose radius is automatically determined to maximise the signal-to-noise of the source data. This is done with a curve-of-growth analysis, performed by the SAS task, *eregionanalyse*. The algorithm then searches for a circular background region on the same CCD where the source is located, excluding regions where sources have been detected, as described in Rosen et al. (2016). The exception is in the case when the source falls on the central CCD of a MOS observation in SmallWindow mode (PrimePartialW2/3). In that case the background is estimated from an annulus (inner radius of $5.5'$ and outer radius of $11'$) centered on the source. The background is therefore estimated from the peripheral CCDs and the central CCD is completely excluded. For EPIC-pn sources, the algorithm avoids the same *RAWY* column

as the source in order to exclude out-of-time events from the background estimation. The background region always has a radius larger than 3 pixels, otherwise no background is calculated. Response files (.rmf and .arf) are then created using the SAS tasks *rmfgen* and *arfgen*.

The pile-up is estimated as described in Section 6.4.1 and written to the header.

4.1. Lightcurve generation

Lightcurves are corrected using the SAS task, *epiclccorr*, to take into account events lost through inefficiencies due to vignetting, bad pixels, chip gaps, PSF and quantum efficiency, dead time, GTIs and exposure. *epiclccorr* also takes into account the background counts, using the background lightcurve, extracted over the identical duration as the source lightcurve. The time bin size for the pn lightcurves was previously a minimum of 10 s and could be as poorly sampled as tens of thousands of seconds for the faintest sources. To exploit the high time resolution and high throughput of the pn, for 4XMM we now extract the pn lightcurve such that each bin is 20 times the frame time, usually 1.46 s. The binning of the MOS data remains as it was for 3XMM.

4.2. Variability characterisation

As in 3XMM, the χ^2 test was used to determine if a source is variable during a single observation. Variability was defined as $P(\chi^2) \leq 10^{-5}$. We also gave the fractional variability, F_{var} , to provide the scale of the variability (Rosen et al. 2016). These values are still provided in the 4XMM catalogue. The χ^2 statistic can be applied to binned data sets where the observed number of counts in a bin deviates from expectation approximately following a Gaussian distribution. Cash (1979) showed that when the number of counts per bin falls below ~ 10 -20, the χ^2 statistic becomes inaccurate. Therefore, as the pn lightcurves are now binned to $20 \times \text{frame time}$ (Section 4.1), these data are rebinned to contain 20 counts per bin before applying the variability tests. Future versions of the catalogue are expected to exploit the high time resolution of the pn lightcurves using a Kolmogorov-Smirnov test which can be carried out on finely binned data.

As in previous catalogue versions, we still provide columns with the minimum EPIC source flux (and error) and the maximum EPIC source flux (and error). This allows the user to find sources variable between observations. Alternatively, the fluxes from each observation, along with the observation date can be seen directly as a table when querying a source on the catalogue server⁸. Whilst the majority of sources do not vary in flux, the maximum variability in the catalogue is a factor 3.5×10^5 in flux (Markarian 421).

Variability between observations is also provided in the stacked catalogue, see Section 7.

5. Screening

Visual inspection of each detection in every observation that was included in 4XMM was carried out, as has been done for previous versions of the catalogue (Rosen et al. 2016). The aim of the screening is to visually validate the new methodology employed in the pipeline, ensure that the pipeline processing has run correctly, and to flag detections that are likely to be spurious and that have not been automatically identified as possibly spurious

⁸ xmm-catalog.irap.omp.eu

in the pipeline processing. Whilst the source detection process is very robust, some spurious detections can still occur in the wings of the PSF of a bright source, in reflection arcs caused by a bright source outside of the field of view, in very extended diffuse emission in the field of view, or because of anomalous noise in a region of the detector, for example. The regions affected are masked and any detections in such regions are subsequently assigned a manual flag (flag 11) in the flag columns (*PN_FLAG*, *M1_FLAG*, *M2_FLAG*, *EP_FLAG*). The fraction of the field of view that is masked is characterised by the observation class (*OBS_CLASS*) parameter. The definition of the *OBS_CLASS* parameter is given in the Table 3, along with the percentage of the catalogue (4XMM-DR9 and 3XMM-DR8 for comparison) with that particular *OBS_CLASS* value.

Table 3. 4XMM observation classification (*OBS_CLASS*, first column), the percentage of the field considered problematic (second column) and the percentage of fields that fall within each class for 3XMM-DR8 and 4XMM-DR9 (third and fourth columns respectively)

OBS CLASS	masked fraction	3XMM-DR8	4XMM-DR9
0	bad area = 0%	18%	30%
1	0% < bad area < 0.1%	17%	30%
2	0.1% < bad area < 1%	16%	17%
3	1% < bad area < 10%	26%	13%
4	10% < bad area < 100%	14%	9%
5	bad area = 100%	4%	1%

There has been a marked improvement in the number of spurious sources within each observation from the 3XMM to the 4XMM catalogue. This can be seen in Table 3 which gives the area of each observation containing spurious detections. 77% of fields have less than 1% of the field containing spurious sources, compared to only 51% in 3XMM-DR8. Only 1% of the fields in 4XMM-DR9 have no good sources in the field of view, compared to four times this value in 3XMM-DR8. This clearly shows the improvement in source detection, primarily due to the new background methods employed in the pipeline for 4XMM.

In 3XMM, flag 12 was not officially used. In 4XMM it indicates whether the source maybe spurious due to being on or close to warm/flickering pixels identified through stacking all of the detections in the 4XMM catalogue (see Sec. 3.5).

6. Catalogue construction

6.1. Unique sources

The 4XMM detection catalogue contains multiple detections (up to 69 times in the most extreme case) of many X-ray sources, due to partial overlap between fields of view as well as repeated observations of the same targets. As has been done in previous versions of the catalogue (Rosen et al. 2016), we assign a common unique source identifier, SRCID, to individual detections that are considered to be associated with the same X-ray source on the sky. The procedure used to perform associations is the same (and therefore subject to the same caveats) as the one outlined in section 6 of Rosen et al. (2016).

6.2. Naming convention for the DETID and the SRCID

Starting in 3XMM-DR5, the procedure for attributing the detection identification number (DETID) and the unique source identification number (SRCID), both being unique to each detection and each unique source respectively, has been modified. Previously, identification numbers were re-computed for each cata-

logue version leading to supplementary columns added to the catalogue with the DETID and SRCID from previous releases.

The DETID is now constructed from the OBS_ID, which always remains the same for an observation, coupled with the source number SRC_NUM⁹ as follow:

DETID = “1” + OBS_ID + SRC_NUM

where the “+” sign indicates string concatenation and where SRC_NUM is zero-padded to form a 4 digit number. The SRCID of a unique source is then determined from the first DETID attributed to that source (i.e. in the observation where the source was first detected) and replacing the first digit “1” by “2”.

Despite the new naming convention that aims at preserving SRCID numbers across catalogue versions, a certain number of SRCID can disappear from one catalogue version to another. This is a normal consequence of the algorithm that groups detections together into unique sources (see section 6 of Rosen et al. 2016). When new data are added and statistics are improved, the algorithm might find a better association of detections into unique sources. As an example, a total of 134 SRCIDs listed in 3XMM-DR7 are absent in 3XMM-DR8.

6.3. Missing detections and DETID change

In addition, the pipeline reprocessing of the full public dataset for the 4XMM version of the source catalogue led to significant modifications of the detection list. There are 10 214 observations that are common between the 3XMM-DR8 and 4XMM-DR9 catalogues, resulting in 773 241 detections in 3XMM-DR8 and 726 279 detections in 4XMM-DR9. Of these, there are 608 071 point-like detections with a SUM_FLAG ≤ 1 in 3XMM-DR8 and 607 196 in 4XMM-DR9. However, amongst these observations, there are $\sim 128\,000$ detections that appear in 3XMM-DR8 that are not matched with a detection in the same observation in the 4XMM-DR9 catalogue within a 99.73% confidence region (i.e. $2.27 \times \text{POSERR}$). About 67 000 of these were classified as the cleanest (SUM_FLAG ≤ 1), point-like sources in 3XMM-DR8 – these are referred to as missing 4XMM detections in what follows. It should be noted that in reverse, there are $\sim 164\,000$ detections in the 4XMM-DR9 catalogue that are in common observations but not matched with a detection in 3XMM-DR8 within 99.73% confidence region, approximately 107 000 of which are classed as being clean and point-like. This is an expected consequence of the reprocessing which was already encountered in the transition from 2XMM to 3XMM (see Section 8 and Appendix D in Rosen et al. 2016). The number of missing 4XMM detections is consistent with the number of missing 3XMM detections, where there were $\sim 25\,700$ good detections that appeared in 2XMMi-DR3 that were not matched with a detection in the same observation in the 3XMM-DR5 catalogue (Rosen et al. 2016). This amounts to $\sim 4.5\%$ which is of the same order as the number of missing sources in 4XMM (8.3%). The origin of these source discrepancies between the two catalogues are the improvements made to the pipeline and in particular the new background estimation. The majority of the detections present in 3XMM-DR8 that are not present in 4XMM-DR9 are from the lowest maximum likelihoods, see Figure 6. A small change in the parameters can cause a source with a maximum likelihood close to the cut-off of 6, but none the less slightly

⁹ SRC_NUM is the source number in the individual source list for a given observation; Sources are numbered in decreasing order of count rate (i.e. the brightest source has SRC_NUM = 1).

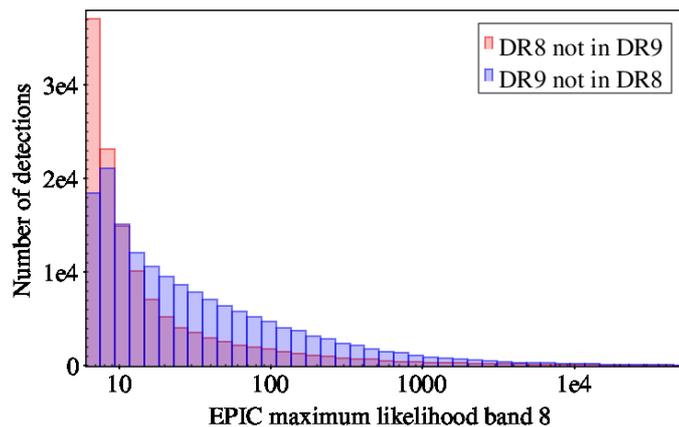


Fig. 6. A histogram showing the detections present in 3XMM-DR8 and not present in 4XMM-DR9 as a function of maximum likelihood (red) and those in 4XMM-DR9 and not in 3XMM-DR8 (blue).

above, to have a value slightly below the cut-off and therefore be excluded from the catalogue. Conversely, the changes in the pipeline for sources just below the maximum likelihood cut-off of 6 and therefore not in 3XMM-DR8 can mean that they will then have a higher maximum likelihood and be present in 4XMM-DR9. As discussed in Section 5, fewer obviously spurious detections are found in 4XMM-DR9 than in 3XMM-DR8, which is also reflected in Figure 6, where the detections found in 4XMM-DR9 and not in 3XMM-DR8 are generally more reliable (higher maximum likelihood).

A related consequence is that the source numbering within a given observation (i.e. the SRC_NUM) has been altered in 4XMM-DR9 by the detections added and those removed. Therefore, amongst the detections that are matched between 3XMM-DR8 and 4XMM-DR9, the majority of them have different DETIDs in 4XMM-DR9 and 3XMM-DR8 (since the DETID is constructed from SRC_NUM). To minimise this effect, for the detections matched between the two catalogues, we have chosen to keep the original 3XMM-DR8 DETIDs instead of the newly generated ones for 4XMM-DR9. However, in doing so, we ended up with $\sim 36\,000$ DETID duplicates due to unmatched 4XMM-DR9 detections having the same DETID as matched 3XMM-DR8 detections. In such cases, we added 5000 to the DETID of the unmatched detection to create a new unique DETID.

6.4. New and revised data columns in 4XMM

We have taken the opportunity of this major release version to revise some data columns and introduce new ones to the catalogues of detections and unique sources (the slim version).

- A pileup evaluation per instrument for each detection is now provided as three new columns: PN_PILEUP, M1_PILEUP and M2_PILEUP, see Section 6.4.1.
- In 3XMM-DR8 and earlier version, the extent likelihood EP_EXTENT_ML was provided only for sources detected as extended. We now provide the extent likelihood for all sources, see Section 6.4.2.
- Small changes have been made to the calculation of the source extent of unique sources (SC_EXTENT) to be aligned with the documentation.
- We now provide the error on the total band extent of a unique source: SC_EXT_ERR. It is calculated in the same way as

the errors on the other unique source parameters (e.g. the SC_EP_FLUX_ERR or the SC_HRn_ERR) namely, as the error on the weighted mean:

$$\text{SC_EXT_ERR} = \sqrt{\frac{1}{\sum_i \frac{1}{\text{EP_EXTENT_ERR}_i^2}}}$$

where EP_EXTENT_ERR_{*i*} is the total band error on the extent of the *i*th detection of the unique source.

6.4.1. Pile up information

As of 4XMM we provide three new columns (PN_PILEUP, M1_PILEUP and M2_PILEUP) quantifying whether each detection may be affected by pile-up in any instrument. A value below 1 corresponds to negligible pile-up (less than a few % flux loss) while values larger than 10 denote heavy pile-up. Pile-up is dependent on time for variable sources. We neglect that here, but note that a variable source is more piled-up than a constant one for the same average count rate, so our pile-up level can be viewed as a lower limit. We also neglect the slight dependence on the source spectrum due to the event grade dependence of pile-up.

Our pile-up levels are not based on a fit of the full images using a pile-up model (Ballet 1999). For point sources, they are based on the measured count rates reported in the catalogue over the full energy band, transformed into counts per frame. The thresholds (at which the pile-up level is set to 1) are set to 1.3 cts/frame for MOS and 0.15 cts/frame for PN (Jethwa et al. 2015).

For extended sources, the pile-up level is based on the measured count rate per CCD pixel at the source position, and therefore refers to the peak brightness, assuming this can be considered uniform at the pixel scale (4.1'' for PN). The threshold is set for all instruments to 5×10^{-3} cts/frame/pixel, such that the flux loss is also a few % when the pile-up level is 1.

6.4.2. Extent likelihood

All detections are tested for their potential spatial extent during the fitting process. The instrumental point-spread function (PSF) is convolved with a β extent model, fitted to the detection, and the extent likelihood EP_EXTENT_ML is calculated as described by Section 4.4.4 of Watson et al. (2009). A source is classified as extended if its core radius (of the β -model of the PSF), $r_c > 6''$ and if the extended model improved the likelihood with respect to the point source fit such that it exceeded a threshold of $L_{ext,min}=4$. In the 4XMM catalogues, EP_EXTENT_ML is included for all detections, while it was set to undefined for point-like detections in previous catalogues. $L_{ext,min} \geq 4$ indicates that a source is probably extended, whilst negative values indicate a clear preference of the point-like over the extended fit. As in the previous catalogue, a minimum likelihood difference of four has been chosen to mark a detection as extended. This threshold makes sure that the improvement of the extended over the point-like fit is not only due to statistical fluctuations but from a more precise description of the source profile.

7. The stacked catalogue

A second independent catalogue is compiled in parallel by the XMM-Newton SSC, called 4XMM-DR9s, where the letter 's' stands for stacked. This catalogue lists source detection results

on overlapping XMM-Newton observations. The construction of the first version of such a catalogue, 3XMM-DR7s, is described in Traulsen et al. (2019). The construction of 4XMM-DR9s essentially follows the ideas and strategies described there with a few important changes that are described in full detail in the accompanying paper Traulsen et al. (sub.). The two main changes concern the choice of input observations and event-based astrometric corrections before source detection. Also it was found necessary to perform some visual screening of the detections, whose results are reported in the source catalogue.

Observations entering 3XMM-DR7s were filtered rather strictly. Only observations with OBS_CLASS < 2, with all three cameras in full-frame mode, and with an overlap area of at least 20% of the usable area were included. All those limitations were relaxed for the construction of 4XMM-DR9s which resulted in a much larger number of observations to be included and potentially much larger stacks (more contributing observations). Before performing simultaneous source detection on the overlapping observations, individual events were shifted in position using the results from the previous catcorr positional rectification of the whole image processed for 4XMM-DR9. This led to a clear improvement of the positional accuracy in stacked source detection.

All sources found by stacked source detection are listed in 4XMM-DR9s, including those from image areas where only one observation contributes. One may expect some differences between these same sources in 4XMM-DR9 and DR9s, because their input events were treated differently. More information is given in Traulsen et al. (sub.).

4XMM-DR9s is based on 1329 stacks (or groups) with 6604 contributing observations. Most of the stacks are composed of 2 observations, the largest has 352. The catalogue contains 288191 sources, of which 218283 have several contributing observations. Auxiliary data products comprise X-ray and optical images and long term X-ray light curves. Thanks to the stacking process, fainter objects can be detected and 4XMM-DR9s contains more sources compared to the same fields in 4XMM-DR9.

8. Catalogue properties

The 4XMM-DR9 catalogue contains 810795 detections, associated with 550124 unique sources on the sky, extracted from 11204 public XMM-Newton observations. Amongst these, 121792 unique sources have multiple detections, the maximum number of repeat detections being 69, see Fig. 7. 76999 X-ray detections in 4XMM-DR9 are identified as extended objects, i.e. with a core radius parameter, r_c , as defined in section 4.4.4 of Watson et al. (2009), $> 6''$ and EP_EXTENT_ML ≥ 4 , with 74163 of these having $r_c < 80''$.

8.1. Astrometry

To evaluate the astrometry in 4XMM-DR9, we followed a process similar to that described in Section 3.3, that was used to update the systematic error in the case that poscorrok=false. However, here we use all of the detections in 4XMM-DR9 and any value of poscorrok. The systematic astrometric uncertainty of the 4XMM DR9 detection catalogue has been estimated empirically using the SDSS DR14 QSO catalogue (Pâris et al. 2018), following similar steps as those detailed in section 3.3. The sources in the latter catalogue have been filtered (good quality spectra and to avoiding mismatches between targeting and matched pho-

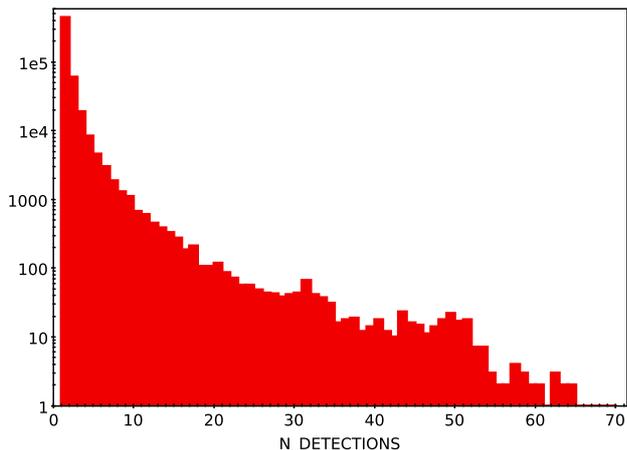


Fig. 7. The number of 4XMM-DR9 unique sources plotted as a function of the number of detections.

tometry¹⁰) and cross-matched with the SDSS DR9 photometry catalogue with a maximum distance of 5 arcsec, at this stage there are 487845 SDSS QSO in our sample. We discarded all quasars with more than one SDSS DR9 counterpart within the radially averaged 4XMM-DR9 uncertainty in the SDSS positions $\Delta S = \sqrt{(\Delta\alpha^2 + \Delta\delta^2)/2} + 0.1^2$, keeping only pointlike objects ($c1=6$). We cross-correlated the 402291 “clean” quasars with the “slim” catalogue out to a distance of $r = 30''$. For each of the resulting pairs we have estimated the combined positional error as $\sigma = \sqrt{\Delta S^2 + \Delta X^2/2}$, where $\Delta X \equiv \text{SC_POSERR}$, which is two dimensional, and discarded all quasars that had more than one counterpart out to $x = r/\sigma = 6$, making 11640 suitable quasars (there were 43 quasars with more than one counterpart out to that limit).

Filtering as described in Section 3.3 leaves 15001 quasar-X-ray pairs with $x < 5$. To follow the Rayleigh distribution $P(x) = xe^{-x^2/2}$, we have added an additional positional uncertainty Σ in quadrature, so that the total positional uncertainty is now $\sigma' = \sqrt{\sigma^2 + \Sigma^2}$, looking for the value of Σ that minimizes the log-likelihood of the $x' \equiv r/\sigma'$ and the Rayleigh distribution. We find $\Sigma = 0.961 \pm 0.008$ arcsec for the uncorrected 4XMM-DR9 X-ray positions, where the uncertainty (1σ) has been calculated by bootstrap with replacement. This can be seen in Figure 8.

To directly compare the quality of the astrometry in 3XMM-DR8 and 4XMM-DR9, we matched each catalogue of detections with the DR14 release of the SDSS quasar catalogue. Cross-matching was performed without restrictions on the types of XMM-Newton and SDSS sources considered, but we kept only those matches within a matching radius of $15''$. This yielded a total of 16530 3XMM-QSO pairs and 18002 4XMM-QSO pairs. Figure 9 shows a scatter plot and associated histograms of the RA and Dec offsets between the XMM sources and SDSS quasars. We see that the general astrometric quality of the 4XMM-DR9 catalogue is very good, with mean RA and Dec offsets of $-0.01''$ and $0.005''$ respectively with corresponding standard deviation of $0.70''$ and $0.64''$. No significant improvement is observed when comparing with the 3XMM-DR8 - SDSS match.

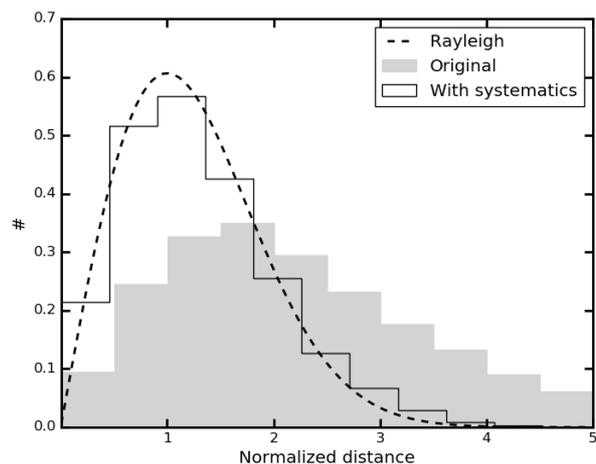


Fig. 8. Fraction of XMM-Newton-SDSS quasar pairs as a function of normalised distance x , before adding a systematic uncertainty (grey histogram) and after its addition (black solid line), along with the Rayleigh distribution (black dashed line).

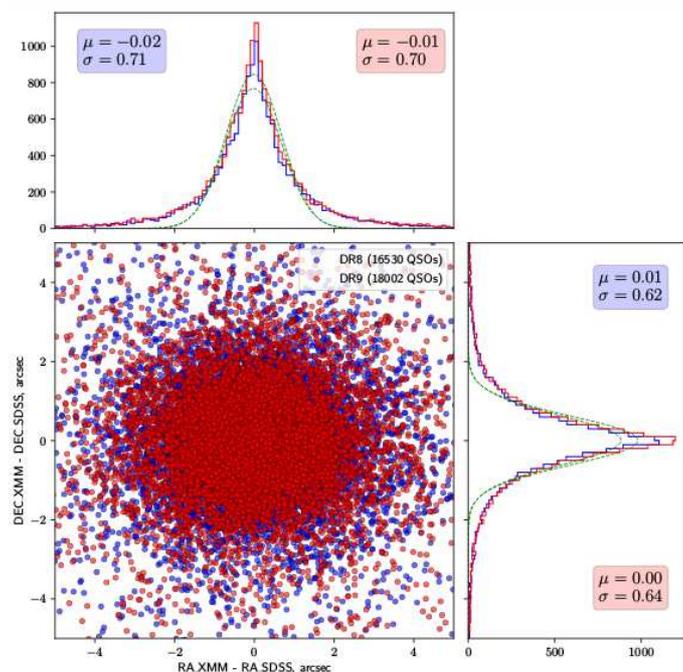


Fig. 9. Scatter plot and associated distribution of the RA and Dec offsets between the XMM sources and the SDSS optical quasars. Two versions of the XMM catalogues are compared: 4XMM-DR9 (red) and 3XMM-DR8 (blue). The dashed green curves in the histogram plots represent gaussian fits to the distributions. The derived mean μ and standard deviation σ for each fit are shown in the coloured boxes respectively.

8.2. Extended sources

Only 76999 4XMM-DR9 detections (9.50%) are identified as extended, compared to 91111 in 3XMM-DR8 (11.75% of the catalogue). However, of the extended sources in 4XMM-DR9, 30464 have the best quality flag (SUM_FLAG=0, 40% of extended sources), whereas only 12256 of the 3XMM-DR8 extended sources (13%) have this flag. This implies that the detection of extended sources is more reliable in the new version of

¹⁰ see <https://www.sdss.org/dr15/spectro/caveats/>

the catalogue, with fewer spurious extended sources. This is due to the improved background modelling used for 4XMM-DR9.

9. External catalogue cross-correlation

Cross-correlation with archival catalogues is performed by a distinct pipeline module running at the Observatoire Astronomique de Strasbourg and referred to as the Astronomical Catalogue Data Subsystem (ACDS). For each individual EPIC detection the ACDS lists all possible multi-wavelength identifications located within a 3σ combined XMM and catalogue error radius from the EPIC position. Finding charts and overlays with ROSAT all-sky survey images of the field are also produced. A detailed description of the ACDS is given in Rosen et al. (2016).

We took the opportunity of the reprocessing of the entire XMM-Newton archive to update the list of archival catalogues and image servers entering the cross-correlation process and finding chart generation. In ACDS version 10.0, a total of 222 catalogues are queried, of which 53 are new with respect to ACDS version 9.0. Among the catalogues providing the largest sky coverage are; GALEX GR6+7 (Bianchi et al. 2017), UCAC4, SDSS DR12 (Alam et al. 2015), panStarrs-DR1 (Chambers et al. 2016), IPHAS DR2 (Barentsen et al. 2014), Gaia DR2 (Gaia Collaboration et al. 2018), 2MASS, AllWISE, Akari, NVSS, FIRST and GLEAM (Hurley-Walker et al. 2017). The XMM-OM Serendipitous Source Survey Catalogue XMM-SUSS4.1 (Page et al. 2012), XMM-Newton slew survey Source Catalogue v. 2.0, the 3XMM-DR8 catalogues, Chandra V2.0 catalogue and the second ROSAT all-sky survey are also queried. Apart from the Chandra Catalogue Release 2.0 whose entries are extracted from the CXC server, all other ACDS catalogues are queried using the VizieR catalogue server.

As for previous releases, 4XMM ACDS tentative identifications are not part of the catalogue proper but are distributed to the community by the XSA and through the XCAT-DB (Michel et al. 2015)¹¹. Finding charts are extracted from several imaging surveys with the following decreasing priority order. First the Sloan digital sky survey (Alam et al. 2015) with colour images made from the g, r and i images extracted from the SDSS server. Second the Pan-STARRS-DR1 (Chambers et al. 2016) with colour images based on the z, g and z+g surveys, third, the MAMA/SRC-J and MAMA/POSS-E plate collections and as a last choice the DSS2 photographic plates. For the one colour photographic surveys, we select the blue image at Galactic latitude $> 20^\circ$, while the red images are preferred in the Galactic plane. Apart from the SDSS, all images are extracted in HEALPix format from Hierarchical Progressive Surveys (HiPS) Aladin server (Fernique et al. 2014).

9.1. Methodology for producing multi-wavelength Spectral Energy Distributions

Spectral energy distributions (SEDs) are provided for each of the unique 4XMM sources. For that purpose, we use basically the same tools as those developed in the framework of the ARCHES project (Motch et al. 2017). The ARCHES algorithm (Pineau et al. 2017) cross-matches in a single pass all selected archival catalogues and for each combination of catalogue entries, computes the cross-match probability. Probabilities are computed from the likelihood that sources in the different catalogues have exactly the same position on the sky, considering their astrometric uncertainties. In particular, the resemblance of the derived

SED with that of any given class of objects do not enter in the computation of the probability. The association probability eventually rests on the prior probability that a given X-ray source has a true counterpart in the longer wavelength catalogue considered. This prior is estimated from the observed distribution of X-ray - longer wavelength catalogue associations taking into account the expected rate of spurious matches. In the original ARCHES project, X-ray sources were grouped by XMM observations with similar exposure times, corresponding to similar limiting sensitivities. Although this grouping method offers a clean and relatively easy way to build X-ray source instalments, it still has the disadvantage of mixing bright and faint X-ray sources that will not have the same a priori probability to have a counterpart in the longer wavelength catalogues considered. In order to cope with this potential statistical bias, we designed a method aimed at grouping X-ray sources by range of X-ray flux instead. Accordingly, the ARCHES cross-matching tool had to be modified so as to read the sky area covered by the sample as an input instead of computing it from the list of observations given in entry.

Source detection area requires building EPIC sensitivity maps for each of the XMM observations. In order to compute sensitivity maps, we first tried the approach proposed by e.g. Carrera et al. (2007). The method consists of equating the probability of existence of a given source as provided by EP_8_DET_ML with that derived from an excess of counts above a given background assuming Poissonian statistics. Although good fits can be obtained for EP_8_DET_ML higher than ~ 15 , we found that best fit background areas are highly dependent on off-axis angle and background values when approaching the threshold of EP_8_DET_ML = 6, used as a criteria for a detection to be included in the 4XMM catalogue. Such a discrepancy is not unexpected since the existence probabilities given by the *emldetect* algorithm also depend on the resemblance of the distribution of photons to that of the PSF. In addition, *emldetect* relies on the Cash statistics (Cash 1979) and on the approximation of the Wilks theorem to derive probabilities. Instead, we built sensitivity maps by computing at each pixel location the total EPIC broad band count rate that would yield a mathematical expectation of EP_8_DET_ML equal to 6. For that purpose we assume a power law input source spectrum ($\Gamma = 1.42$; $N_H = 1.7 \times 10^{20} \text{ cm}^{-2}$) similar to that of the unresolved sources contributing to the extragalactic background (Lumb et al. 2002). The source spectrum is then folded through the exposure maps and filter responses so as to obtain the source counts in each band and camera in operation. EP_8_DET_ML is then computed taking into account the background maps and the varying shape of the PSF with telescope and off-axis angle.

We estimated the overlap of the 4XMM-DR9 catalogue with 26 archival catalogues selected to cover the largest sky coverage and widest span in wavelength from UV to radio. The Multi-Order-Coverage map (MOC) (Fernique et al. 2015) of each XMM observation was computed with a resolution of $12.8''$ (order 14) and compared to the MOC footprint of each catalogue using a python code developed at CDS (Baumann & Boch 2019). Table 4 lists the pre-selected catalogues sorted by 4XMM coverage. In the optical band, catalogues were prioritized according to their depth, astrometric quality and range of colours in the following order, SDSS12, PanStarrs DR1 and Skymapper, so as to cover the entire sky. Whenever a GAIA DR2 match was found within $1.4''$ from the catalogue entry, the GAIA position was assigned to the merged source. APASS9 photometry was added to the merged source if found within a $1.4''$ distance so as to extend the photometric measurements to brighter objects. The $1.4''$ radius was derived from the shape of the Rayleigh dis-

¹¹ <http://xcatdb.unistra.fr>

tribution of the distances between matching sources and guarantees a low rate of false cross-identification. In a similar manner, we cross-matched the ALLWISE and 2MASS catalogues keeping the 2MASS position whenever the difference of position was lower than $3.5''$ at $|b| \geq 20$ deg and $1.5''$ at $|b| \leq 20$ deg. Special sky regions such as M31 and the LMC were discarded due to their high optical source density. For each unique source, we only kept the observation offering the highest detection area. 4XMM sources were then grouped into 4 EPIC (0.2-12.0 keV) ranges of flux with boundaries at 1.4, 3.1 and 7.2×10^{-14} erg $\text{cm}^{-2} \text{s}^{-1}$. This grouping yields a nearly even number of sources in each flux band.

The statistical ARCHES cross-match procedure was applied to 5 catalogues or group of catalogues: XMM, Gaia, SUSS-OM, merged optical and merged infrared. Due to the different areas of the non all-sky catalogues (Gaia, SDSS12, PanStarrs and Skymapper) we splitted the XMM observations into groups having homogeneous catalogue coverages. In addition, the galactic plane region was treated separately. Finally, a simple cross-match between the ARCHES result and both the AKARI and merged FIRST/NVSS compiled by Mingo et al. (2016) was made. However, their matching likelihoods do not enter in the computation of the overall SED probability provided by the ARCHES tool.

A standard table at CDS¹² allows us to convert magnitudes into flux. The resulting SEDs are available as individual FITS files and graphical output for the 3 highest probability SEDs.

The sensitivity maps, individual observation MOCs and total 4XMM MOCs are available on the XMM-SSC website¹³.

10. Catalogue access

The catalogue of detections is provided in several formats. A Flexible Image Transport System (FITS) file and a comma-separated values (CSV) file are provided containing all of the detections in the catalogue. For 4XMM-DR9 there are 810795 rows and 336 columns. A separate version of the catalogue (the slim catalogue) with only the unique sources is provided, i.e. 550124 rows, and has 45 columns, essentially those containing information about the unique sources. This catalogue is also provided in FITS and CSV format. We also provide SQL CREATE statements to load the data in CSV format. These can be found on the *XMM-Newton* Survey Science Centre webpages¹⁴. The stacked catalogue is provided in FITS format only. Ancillary tables to the catalogue also available from the *XMM-Newton* Survey Science Centre webpages include the table of observations incorporated in the catalogue.

The *XMM-Newton* Survey Science Centre webpages provide access to the 4XMM catalogue, as well as links to the different servers distributing the full range of catalogue products. These include, the ESA *XMM-Newton* archive (XSA), which provides access to all of the 4XMM data products, and the ODF data, the XCat-DB¹⁵ produced and maintained by the *XMM-Newton* SSC, which contains possible EPIC source identification produced by the pipeline by querying 222 archival catalogues, see Section 9. Finding charts are also provided for these possible identifications. Other source properties as well as images, time series and

Table 4. Overlapping area between photometric catalogues and 4XMM observations. The last column shows the way the catalogue was processed, either using the ARCHES multi-catalogue statistical cross-match (s) or using a simple positional cross-match (x)

Catalogue	Total area covered (deg ²)	Overlap with 4XMM (deg ²)	Band	Xmatch mode
AllWISE	all-sky	1152	ir	s
Gaia DR2	all-sky	1152	opt	s
UCAC4	all-sky	1152	opt	
2MASS	all-sky	1152	ir	s
APASS	all-sky	1126	opt	s
Akari	39406	1108	farir	x
GMRT	36996	1000	radio	
NVSS	34069	927	radio	x
PanStarrs DR1	32134	881	opt	s
GalexGR67	26249	696	uv	s
GLEAM	25423	657	radio	
SkyMapper	19585	550	opt	s
SDSS12	14520	504	opt	s
FIRST	10847	427	radio	x
VHS	13670	364	ir	
XMM-OM-SUSS41	348	343	uv	s
SUMSS	8354	216	radio	
UKIDSS LAS	3695	174	ir	
VST	3988	86	opt	
Galex MIS	1880	83	uv	
VPHAS	670	77	opt	
UKIDSS GPS	1366	76	ir	
WBH2005 20	614	72	radio	
Glimpse	471	70	ir	
IPHAS	1888	59	opt	
WBH2005 6	164	35	radio	

spectra are also provided. Multi-wavelength data taken as a part of the XID (X-ray identification project) run by the SSC over the first fifteen years of the mission are also provided in the XIDresult database¹⁶. The XMM-SSC catalogue server¹⁷ provides access to each source and regroups information concerning all of the detections for a unique source. It also provides the XMM-Newton lightcurves and spectra and permits the user to undertake simple spectral fitting, as well as overlays of the same region of sky in all wavelengths. The catalogue can also be accessed through HEASARC¹⁸ and VIZIER¹⁹. The results of the external catalogue cross-correlation carried out for the 4XMM catalogue (section 9) are available as data products within the XSA or through the XCat-DB. The *XMM-Newton* Survey Science Centre webpages also detail how to provide feedback on the catalogue.

Where the 4XMM catalogue is used for research and publications, please acknowledge their use by citing this paper and including the following:

¹² [http://vizier.u-strasbg.fr/viz-bin/VizieR-3?](http://vizier.u-strasbg.fr/viz-bin/VizieR-3?source=METAfltr)

[-source=METAfltr](#)

¹³ http://xmmssc.irap.omp.eu/Catalogue/4XMM-DR9/4XMM_DR9.html

¹⁴ <http://xmmssc.irap.omp.eu/>

¹⁵ <http://xcatdb.unistra.fr/4xmm/>

¹⁶ <http://xcatdb.unistra.fr/xidresult/>

¹⁷ <http://xmm-catalog.irap.omp.eu/>

¹⁸ <http://heasarc.gsfc.nasa.gov/db-perl/W3Browse/w3table.pl?table-head=name%3Dxmmssc&Action=More+Options>

¹⁹ <http://vizier.u-strasbg.fr/cgi-bin/VizieR>

This research has made use of data obtained from the 4XMM XMM-Newton serendipitous source catalogue compiled by the 10 institutes of the XMM-Newton Survey Science Centre selected by ESA.

It is important to note that the 4XMM catalogue of detections, as for previous versions of this catalogue, contains detections with a significance as low as $\sim 3\sigma$ (Maximum likelihood of 6), along with sources that have been flagged as possibly spurious. Statistically some of these sources will be spurious. In order to create the cleanest catalogue possible, where statistically almost all sources are real, it is necessary to filter the catalogue to include only EPIC sources with for example a 5σ significance (Maximum likelihood of ~ 14) and to keep only those with with no flags, for example,

```
EP_8_DET_ML > 14 && SUM_FLAG < 1
```

11. Upper limits for observed regions of the sky

The XMM-SSC provides an upper limit server for the user to determine an upper limit for the flux given a non-detection in a region observed by *XMM-Newton*. The server is known as FLIX (Flux Limits from Images from XMM-Newton). This upper limit can be calculated for any of the standard *XMM-Newton* bands for a user defined statistical significance and sky region. A single region or many regions may be queried at the same time. This upper-limit flux is determined empirically using the algorithm described by Carrera et al. (2007). A link to the FLIX upper limit server is provided on the XMM-SSC webpages and the ESA SOC webpages²⁰.

12. Limitations of the catalogue

12.1. Maximum extent of extended detections

When dealing with extended detections, the software determines the radius of the detection, up to a limit of $80''$ to optimise processing time. Whilst this may appear restrictive, only 0.007% of the catalogue detections are clean and extended, with a radius of $>80''$.

12.2. Error values on counts, rate and flux

Should a detection fall close to a chip gap or the edge of the field of view on one or more cameras, only a small fraction of the point spread function will be recorded for that camera. The fraction is given by the `XX_MASKFRAC` columns, where `XX` refers to EP (EPIC), PN (pn), M1 (MOS 1) or M2 (MOS 2). Where the `XX_MASKFRAC` value is low, the error on the counts, rate or flux may be very high, compared to the value of the counts, rate or flux, as these quantities are derived for the whole PSF. Note, detections which have less than 0.15 of their PSF covered by the detector are considered as being not detected.

13. Future catalogue updates

Incremental releases (data releases) are planned to augment the 4XMM catalogue. At least one additional year of data will be included with each data release. Data release 10 (DR10) will provide data becoming public during 2019 and should be released during 2020. These catalogues will be accessible as described in Section 10.

²⁰ <https://www.cosmos.esa.int/web/xmm-newton/xsa>

14. Summary

This paper describes the improvements made to the software and calibration used to produce the new major version of the *XMM-Newton* catalogue, 4XMM. Source detection has been shown to be much improved, with fewer spurious sources and in particular, many fewer spurious extended sources. In addition, we provide lightcurves and spectra for a much larger fraction of the catalogue than in previous versions (36% of detections in 4XMM-DR9 compared to 22% of detections in 3XMM-DR8). These spectra and lightcurves benefit from finer binning (MOS spectra and pn lightcurves). We also provide a new version of the stacked catalogue, which covers a much larger fraction of the sky compared to the previous version (485 square degrees in 4XMM-DR9s compared to 150 square degrees in 3XMM-DR7s). The catalogue benefits from extra complementary products, such as multi-wavelength spectral energy distributions for each source, sensitivity maps and catalogue footprint maps. We provide information on how to access the catalogue as well as how to retrieve upper limits for non-detections in the catalogue footprint. The catalogue is ideal for quick access to data products (fluxes, spectra, images, etc), searching for new objects, population studies of homogenous samples and cross correlation for multi-wavelength studies.

Acknowledgements. We are grateful for the strong support provided by the XMM-Newton SOC. We also thank the CDS team for their active contribution and support. The French teams are grateful to Centre National d'Études Spatiales (CNES) for their outstanding support for the SSC activities. SSC work at AIP has been supported by Deutsches Zentrum für Luft- und Raumfahrt (DLR) through grants 500X1701 and 500X1901, which is gratefully acknowledged. FJC acknowledges financial support through grant AYA2015-64346-C2-1P (MINECO/FEDER). MTC and FJC acknowledge financial support from the Spanish Ministry MCIU under project RTI2018-096686-B-C21 (MCIU/AEI/FEDER/UE) cofunded by FEDER funds and from the Agencia Estatal de Investigación, Unidad de Excelencia María de Maeztu, ref. MDM-2017-0765. This paper used data from the SDSS surveys. This research has made use of the VizieR catalogue access tool, CDS, Strasbourg, France (DOI : 10.26093/cds/vizieR). The original description of the VizieR service was published in 2000, A&AS 143, 23. This paper made use of the topcat software (Taylor 2005).

References

- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, *ApJS*, 203, 21
 Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, *ApJS*, 219, 12
 Ballet, J. 1999, *A&AS*, 135, 371
 Barentsen, G., Farnhill, H. J., Drew, J. E., et al. 2014, *VizieR Online Data Catalog*, II/321
 Baumann, M. & Boch, T. 2019, *Astronomical Society of the Pacific Conference Series*, Vol. 523, *New Python Developments to Access CDS Services*, ed. P. J. Teuben, M. W. Pound, B. A. Thomas, & E. M. Warner, 253
 Bianchi, L., Shiao, B., & Thilker, D. 2017, *ApJS*, 230, 24
 Carrera, F. J., Ebrero, J., Mateos, S., et al. 2007, *A&A*, 469, 27
 Cash, W. 1979, *ApJ*, 228, 939
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, *arXiv e-prints*, arXiv:1612.05560
 Ebrero, J. 2019, *XMM-Newton Users Handbook*, Tech. Rep. 2.17, ESA: XMM-Newton SOC
 Fernique, P., Allen, M. G., Boch, T., et al. 2015, *A&A*, 578, A114
 Fernique, P., Boch, T., Pineau, F., & Oberto, A. 2014, in *Astronomical Society of the Pacific Conference Series*, Vol. 485, *Astronomical Data Analysis Software and Systems XXIII*, ed. N. Manset & P. Forshay, 281
 Gabriel, C., Denby, M., Fyfe, D. J., et al. 2004, in *Astronomical Society of the Pacific Conference Series*, Vol. 314, *Astronomical Data Analysis Software and Systems (ADASS) XIII*, ed. F. Ochsenbein, M. G. Allen, & D. Egret, 759
 Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, 616, A1
 Hurley-Walker, N., Callingham, J. R., Hancock, P. J., et al. 2017, *MNRAS*, 464, 1146
 Jansen, F., Lumb, D., Altieri, B., et al. 2001, *A&A*, 365, L1
 Jethwa, P., Saxton, R., Guainazzi, M., Rodriguez-Pascual, P., & Stuhlinger, M. 2015, *A&A*, 581, A104

- Li, T. P. & Ma, Y. Q. 1983, *ApJ*, 272, 317
- Lumb, D. H., Warwick, R. S., Page, M., & De Luca, A. 2002, *A&A*, 389, 93
- Michel, L., Grisé, F., Motch, C., & Gomez-Moran, A. N. 2015, in *Astronomical Society of the Pacific Conference Series*, Vol. 495, *Astronomical Data Analysis Software and Systems XXIV (ADASS XXIV)*, ed. A. R. Taylor & E. Rosolowsky, 173
- Mingo, B., Watson, M. G., Rosen, S. R., et al. 2016, *MNRAS*, 462, 2631
- Motch, C., Carrera, F., Genova, F., et al. 2017, in *Astronomical Society of the Pacific Conference Series*, Vol. 512, *Astronomical Data Analysis Software and Systems XXV*, ed. N. P. F. Lorente, K. Shorridge, & R. Wayth, 165
- Page, M. J., Brindle, C., Talavera, A., et al. 2012, *MNRAS*, 426, 903
- Pâris, I., Petitjean, P., Aubourg, É., et al. 2018, *A&A*, 613, A51
- Pâris, I., Petitjean, P., Ross, N. P., et al. 2017, *A&A*, 597, A79
- Pineau, F. X., Derriere, S., Motch, C., et al. 2017, *A&A*, 597, A89
- Rosen, S. R., Webb, N. A., Watson, M. G., et al. 2016, *A&A*, 590, A1
- Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, 365, L18
- Takey, A., Schwobe, A., & Lamer, G. 2013, *A&A*, 558, A75
- Taylor, M. B. 2005, *Astronomical Society of the Pacific Conference Series*, Vol. 347, *TOPCAT & STILTS: Starlink Table/VOTable Processing Software*, ed. P. Shopbell, M. Britton, & R. Ebert, 29
- Traulsen, I., Schwobe, A. D., Lamer, G., et al. 2019, *A&A*, 624, A77
- Turner, M. J. L., Abbey, A., Arnaud, M., et al. 2001, *A&A*, 365, L27
- Watson, M. G., Auguères, J.-L., Ballet, J., et al. 2001, *A&A*, 365, L51
- Watson, M. G., Schröder, A. C., Fyfe, D., et al. 2009, *A&A*, 493, 339