The *XMM* Cluster Survey: X-ray analysis methodology

E. J. Lloyd-Davies,[1][*] A. Kathy Romer,[1] Nicola Mehrtens,[1] Mark Hosmer,[1]
Michael Davidson,[2] Kivanc Sabirli,[3] Robert G. Mann,[2] Matt Hilton,[4,5]
Andrew R. Liddle,[1] Pedro T. P. Viana,[6,7] Heather C. Campbell,[1,8] Chris A. Collins,[9]
E. Naomi Dubois,[1] Peter Freeman,[10] Craig D. Harrison,[11] Ben Hoyle,[8,12,13,14]
Scott T. Kay,[15] Emma Kuwertz,[1] Christopher J. Miller,[11] Robert C. Nichol,[8]
Martin Sahlén,[16] S. A. Stanford[17,18] and John P. Stott[9]

[1]*Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QH*
[2]*SUPA, Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ*
[3]*Department of Physics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*
[4]*School of Mathematical Sciences, University of KwaZulu-Natal, Private Bag X54001, Durban 4000, South Africa*
[5]*School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD*
[6]*Centro de Astrofísica da Universidade do Porto, Rua das Estrelas, 4150-762 Porto, Portugal*
[7]*Departamento de Física e Astronomia, Faculdade de Ciências, Universidade do Porto, 4169-007 Porto, Portugal*
[8]*Institute of Cosmology and Gravitation, Dennis Sciama Building, Burnaby Road, Portsmouth PO1 3FX*
[9]*Astrophysics Research Institute, Liverpool John Moores University, Twelve Quays House, Egerton Wharf, Birkenhead CH41 1LD*
[10]*Department of Statistics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*
[11]*Department of Astronomy, University of Michigan, Ann Arbor, MI 48109, USA*
[12]*Institute of Sciences of the Cosmos (ICCUB-IEEC), Department of Physics, University of Barcelona, Barcelona 08024, Spain*
[13]*CSIC, Consejo Superior de Investigaciones Cienticas, Serrano 117, Madrid 28006, Spain*
[14]*Helsinki Institute of Physics, University of Helsinki, PO Box 64, FIN-00014 Helsinki, Finland*
[15]*Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL*
[16]*The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, SE-106 91 Stockholm, Sweden*
[17]*Department of Physics, University of California, Davis, CA 95616, USA*
[18]*Institute of Geophysics and Planetary Physics, Lawrence Livermore National Laboratory, Livermore, CA 94551, USA*

ABSTRACT

The *XMM* Cluster Survey (XCS) is a serendipitous search for galaxy clusters using all publicly available data in the *XMM–Newton* Science Archive. Its main aims are to measure cosmological parameters and trace the evolution of X-ray scaling relations. In this paper we describe the data processing methodology applied to the 5776 *XMM* observations used to construct the current XCS source catalogue. A total of $3675 > 4\sigma$ cluster candidates with >50 background-subtracted X-ray counts are extracted from a total non-overlapping area suitable for cluster searching of $410 \, \text{deg}^2$. Of these, 993 candidates are detected with >300 background-subtracted X-ray photon counts, and we demonstrate that robust temperature measurements can be obtained down to this count limit. We describe in detail the automated pipelines used to perform the spectral and surface brightness fitting for these candidates, as well as to estimate redshifts from the X-ray data alone. A total of 587 (122) X-ray temperatures to a typical accuracy of $<40 \, (<10)$ per cent have been measured to date. We also present the methodology adopted for determining the selection function of the survey, and show that the extended source detection algorithm is robust to a range of cluster morphologies by inserting mock clusters derived from hydrodynamical simulations into real *XMM* images. These tests show that the simple isothermal β-profiles is sufficient to capture the essential details of the cluster population detected in the archival *XMM* observations. The redshift follow-up of the XCS

[*]E-mail: E.Lloyd-Davies@sussex.ac.uk

cluster sample is presented in a companion paper, together with a first data release of 503 optically confirmed clusters.

Key words: surveys – galaxies: clusters: intracluster medium – cosmology: observations – X-rays: galaxies: clusters.

1 INTRODUCTION

Clusters of galaxies are massive objects ($10^{13.5-15}$ M$_\odot$) composed of galaxies, hot ionized gas and dark matter. The gravitational potential is dominated by dark matter, with the mass ratio of the three components being roughly 3:10:87, respectively, although with a strong mass dependence in the ratio of gas to stars (Gonzalez, Zaritsky & Zabludoff 2007). Clusters provide us with the opportunity to obtain information about the underlying cosmological model and important insights into the processes that govern structure formation (see Voit 2005; Allen, Evrard & Mantz 2011, for reviews).

While detailed studies of individual clusters are extremely important, especially for obtaining an insight into the small-scale processes that influence the evolution of their baryonic components, a full understanding of the complex nature of cluster formation and evolution requires the study of the galaxy cluster population as a whole. This is best achieved, in practice, by undertaking cluster surveys. The first large cluster surveys were carried out via eyeball searches for galaxy overdensities on optical photographic plates (Abell 1958; Zwicky, Herzog & Wild 1968), but, nowadays, cluster finding uses sophisticated automated techniques.

In this paper we describe automated cluster finding at X-ray wavelengths; the hot ionized gas (or intracluster medium/ICM) emits soft X-ray radiation in proportion to the square of the electron density. However, this is not the only way new clusters are being discovered. For example, the effect of cluster-sized gravitational potentials can be seen in the optical/infrared, via strong or weak gravitational lensing (e.g. Wittman et al. 2003). Increasing numbers of clusters are also being discovered at millimetre wavelengths (e.g. Staniszewski et al. 2009; Marriage et al. 2011; Menanteau et al. 2010; Vanderlinde et al. 2010; Ade et al. 2011; Foley et al. 2011; Williamson et al. 2011) using the Sunyaev–Zel'dovich (SZ) effect (Sunyaev & Zeldovich 1972): the inverse Compton scattering of photons from the cosmic microwave background (CMB) by the hot ICM. At longer wavelengths still, one can discover clusters out to high redshifts using radio telescopes, via the unusual signature of head–tail galaxies (Blanton et al. 2003). Due to the advent of large-format CCD detectors, cluster finding using galaxy overdensities is also currently undergoing a renaissance (e.g. Gladders & Yee 2000; Miller et al. 2005; Koester et al. 2007; Wilson et al. 2009).

Cluster surveys have already revolutionized our understanding of the physics of the ICM (e.g. Ponman, Cannon & Navarro 1999; Arnaud et al. 2010) and delivered cosmological constraints independent of, and competitive with, those derived from observations of the CMB (e.g. Dunkley et al. 2011; Larson et al. 2011) and Type 1a supernovae (e.g. Kessler et al. 2009). When combined with these other cosmological probes, clusters are playing an important role in the quest to understand the nature of dark energy (e.g. Vikhlinin et al. 2009b; Mantz et al. 2010; Rozo et al. 2010; Sehgal et al. 2011; see Sahlén et al. 2009 for a review of earlier cluster cosmology studies dating back to Frenk et al. 1990; Oukbir & Blanchard 1992). Clusters are also being used to test general relativity on large scales (e.g. Rapetti et al. 2010), constrain the properties of neutrinos (e.g. Mantz, Allen & Rapetti 2010), and search for evidence of

non-Gaussian primordial density fluctuations (e.g. Hoyle, Jimenez & Verde 2010). Future cluster surveys will be wider, more sensitive and better calibrated than ever before, and so are sure to deliver significantly improved constraints compared to these existing works (e.g. Majumdar & Mohr 2004; Predehl et al. 2006; Cunha, Huterer & Frieman 2009; Wu, Rozo & Wechsler 2010).

In this paper we present the *XMM* Cluster Survey (XCS), a search for serendipitous galaxy clusters in archival *XMM–Newton* observations. The original XCS concept and motivation is described in Romer et al. (2001). The main goals of the survey are (i) to measure cosmological parameters; (ii) to measure the evolution of the X-ray luminosity–temperature scaling relation (L_X–T_X relation, hereafter); (iii) to study the galaxy properties in clusters to high redshifts; and (iv) to provide the community with a high-quality, homogeneously selected X-ray cluster sample. The XCS follows a rich tradition of X-ray cluster surveys dating back almost 30 years using earlier satellites: Piccinotti et al. (1982, *HEAO I*), Gioia et al. (1990, Einstein), and several derived from the *ROSAT* All Sky Survey (RASS; Ebeling et al. 1998; Böhringer et al. 2000; Ebeling et al. 2000; Ebeling, Edge & Henry 2001; Cruddace et al. 2002; Ebeling, Mullis & Tully 2002; Gioia et al. 2003; Böhringer et al. 2004; Henry et al. 2006), and the ROSAT pointed observations archive (Rosati et al. 1998; Romer et al. 2000; Perlman et al. 2002; Burke et al. 2003; Mullis et al. 2003; Burenin et al. 2007; Horner et al. 2008).

The XCS is not the only project currently exploiting the *XMM–Newton* (*XMM* hereafter) archive for new detections of clusters. Other projects include: XDCP (Mullis et al. 2005; Fassbender et al. 2008; Santos et al. 2009; Fassbender et al. 2011; Schwope et al. 2010; Suhada et al. 2011); *XMM*-LSS (Bremer et al. 2006; Pierre et al. 2006; Pacaud et al. 2007; Adami et al. 2011); SEXCLAS (Kolokotronis et al. 2006); COSMOS (Finoguenov et al. 2007); *XMM*-BSC (Suhada et al. 2010); SXDS (Finoguenov et al. 2010); and the one being carried out by the members of the *XMM* Survey Science Center (Schwope et al. 2004; Lamer et al. 2008). This intense international interest stems from the fact that *XMM* has several features advantageous to cluster searching: in essence it combines sensitivity, and a large field of view, with spectral imaging capabilities.

The *XMM* image quality does not match that of *Chandra*, but it is still good enough to allow one to differentiate between point-like and extended sources over the whole field of view: given that clusters dominate the extended X-ray source population, this then allows us to identify cluster candidates efficiently, despite the fact that clusters only comprise $\simeq 10$ per cent of the total X-ray source population. Moreover, the spectral capabilities of *XMM* allow the measurement of the temperature of the hot ICM directly from the discovery data. These T_X measurements allow us to then estimate cluster masses, something of a vital importance to cosmological studies. Finally, the mission has been in operation for over 10 years, and has built up a large archive of observations distributed across the sky. By now there are several hundred square degrees available that are suitable for a serendipitous cluster survey, already exceeding that of the largest deep ROSAT survey (Burenin et al. 2007). Serendipitous cluster surveys have also been conducted using the *Chandra* archive

(e.g. Barkhouse et al. 2006), although the available area for cluster searching is significantly smaller in comparison to the *XMM* archive.

As predicted in Romer et al. (2001), and now demonstrated below, XCS will deliver the largest number of cluster temperature measurements to date. Importantly, these clusters will form a homogeneous sample (both in terms of selection and analysis) and have a well-understood selection function. In a companion paper (Mehrtens et al. 2001), we present our first data release (XCS-DR1) and this includes 402 T_X measurements. By comparison, the largest previous compilations of T_X values from homogeneous samples contain less than 100 clusters each, e.g. Reiprich & Böhringer (2002, 63 clusters), Henry (2004, 25 clusters), Pacaud et al. (2007, 29 clusters), Pratt et al. (2009, 31 clusters), Vikhlinin et al. (2009a, 85 clusters), and Mantz et al. (2010, 96 clusters). Larger compilations of clusters with heterogeneous selection do exist, and some have significantly better per cluster T_X precision than XCS, but even so the largest published collection is still only 115 strong (Maughan 2007; a larger sample, of 273 low-redshift clusters, was put together by Horner 2001, but was not made public).

XCS highlights to date include the detection and subsequent multiwavelength follow-up of a $z = 1.46$ cluster (XMMXCS J2215.9−1738; Stanford et al. 2006; Hilton et al. 2007, 2009, 2010), which for several years held the record for the highest redshift spectroscopically confirmed cluster (recent discoveries of higher redshift X-ray clusters include Henry et al. 2010; Papovich et al. 2010; Tanaka, Finoguenov & Ueda 2010; Gobat et al. 2011). XCS clusters have also been used in compilation studies of galaxy evolution in high-redshift clusters (Collins et al. 2009; Stott et al. 2010). Conservative forecasts of the performance of XCS for cosmological parameter estimation and cluster scaling relations can be found in Sahlén et al. (2009): we expect to measure (at 1σ and from clusters alone, i.e. not in combination with CMB or supernovae observations) Ω_m

to ±0.03 (and Ω_Λ to the same accuracy assuming flatness), and σ_8 to ±0.05, whilst also constraining the normalization and slope of the L_X–T_X relation to ±6 and ±13 per cent, respectively.

In this paper, we present an overview of the XCS data analysis strategy, from acquiring the data to producing a catalogue. A schematic of our approach is shown in Fig. 1; however, note that components indicated with dashed outlines are discussed in Mehrtens et al. (in preparation). The paper is comprised of three main sections. Section 2 describes data acquisition, reduction and image generation. Section 3 describes source detection, the compilation of candidate lists and simulations of the survey selection function. Section 4 describes how we use *XMM* data to measure X-ray redshifts, temperatures and luminosities for the candidates.

2 *XMM* DATA REDUCTION

The *XMM* archive contains thousands of public observations suitable for conducting the *XMM* Cluster Survey. Such a large volume of data means we have to carry out most of the XCS in a fully automated manner – the only parts that are not automated are the mask making (Section 2.4.1), optical follow-up, and quality control (Mehrtens et al., in preparation). While this automation presents a number of challenges, in terms of handling the variety and complexity of the archival data, it also has a number of benefits: not only has the entire data set been treated in a consistent and systematic way, but we are also able to run realistic simulations of our selection function.

In this section we describe how the raw *XMM* archive is manipulated into science-grade image files. First the data are downloaded from the remote storage facility at the European Space Astronomy Centre (ESAC), near Madrid, to the University of Sussex (Section 2.2). Then the data are calibrated and cleaned of periods of high background contamination (Section 2.3). Next, images are
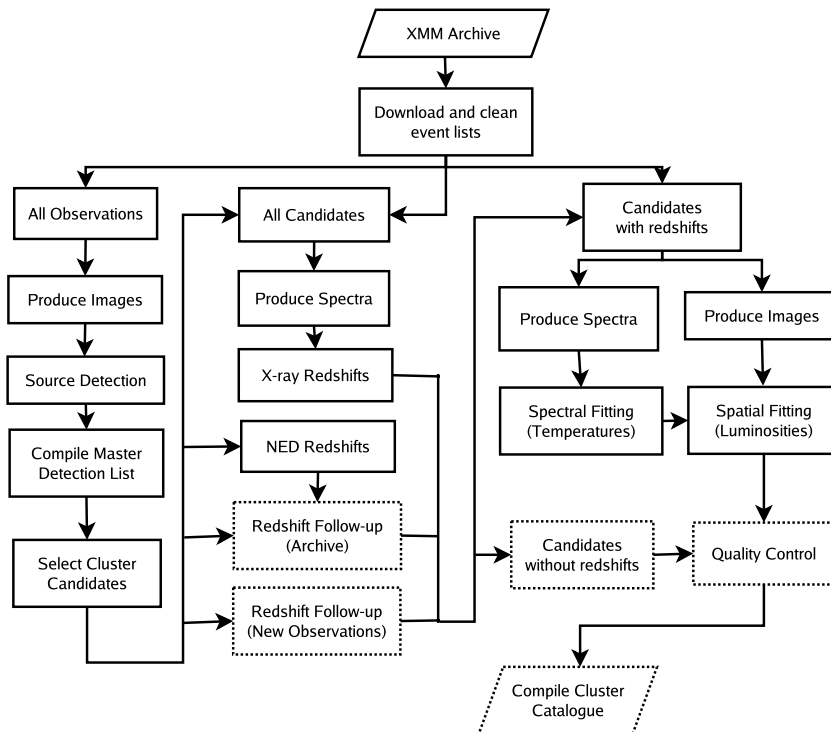

Figure 1. Flow chart showing an overview of the XCS analysis methodology. This illustrates the sequence by which data from the *XMM* archive are used to create a catalogue of galaxy clusters.

produced (Section 2.4) and flux conversion factors calculated (Section 2.5). We begin this section with an overview of some of the salient features of the *XMM* mission.

2.1 The *XMM–Newton* mission

The *XMM* mission (Jansen & Laine 1997) consists of three co-aligned Wolter Type I (Wolter 1952b,a) X-ray telescopes mounted on the same spacecraft. The mission was undertaken by the European Space Agency (ESA) and the spacecraft was launched on 1999 December 10. The mission configuration, with three separate telescopes simultaneously illuminating three cameras, means that most exposures generate data with potential for serendipitous cluster finding: by comparison, Chandra (Weisskopf 1999) has a single telescope that illuminates only one of several instruments at any given time, and not all those instruments are suitable for cluster finding.

The European Photon Imaging Camera (EPIC: Villa et al. 1996) consists of three separate cameras, each in the focal plane of a separate X-ray telescope. Each camera consists of an array of charge-coupled devices (CCDs: Boyle & Smith 1970) in different configurations. Two cameras, the EPIC-mos1 and 2, consist of arrays of seven metal oxide semiconductor CCDs illuminated by 44 per cent of the light from their respective telescopes (the rest is redirected to the Reflection Grating Spectrometers). The EPIC-pn camera consists of 12 back-illuminated CCDs. These CCDs are not only more sensitive than those in the EPIC-mos cameras, but the EPIC-pn receives all the light from its respective telescope. Thus, the EPIC-pn camera has more than twice the sensitivity of the EPIC-mos cameras.

One disappointing aspect of both *XMM* and *Chandra* has been the unexpectedly high background in their CCD cameras. Both these missions are in similar, highly elliptical orbits, and it was only after their launch that it was realized that these orbits intersect a population of low-energy protons trapped in the Earth's magnetosphere. The lower energy protons can be funnelled by the grazing incidence mirrors on to the detectors and this has resulted in a significantly higher background than was expected before launch. Consequently, certain aspects of XCS have proved to be more challenging than was anticipated in our pre-launch predictions (Romer et al. 2001). In addition to the enhanced background, there have been a number of incidents of damage to the EPIC cameras while in orbit, but in only one case has this resulted in a significant loss of detector area (Abbey et al. 2006).

2.1.1 XMM–Newton point spread function

A crucial issue for the detection of extended sources by XCS is the treatment of the *XMM* point spread function (PSF). The PSF is a strong function of off-axis angle and photon energy (where off-axis angle is the angle between the source location and the centre of the field of view). As the off-axis angle increases, the PSF shape morphs from being circularly symmetric to ellipsoidal and finally bow-tie shaped. There have been a number of attempts to characterize the *XMM* PSF including: simulations based on measurements of the shape of the mirrors (Gondoin et al. 1996); measurements taken on the ground by passing X-ray beams from synchrotron sources through *XMM* mirror modules (Stockman et al. 1997; Gondoin et al. 1998); and fitting one-dimensional profiles to observations of bright X-ray sources (Gondoin et al. 2000; Ghizzardi 2001, 2002; Read 2004). Unfortunately, thus far, this has not resulted in a complete

and reliable characterization of the *XMM* PSF. Currently four PSF models are available: the Low, Medium, High and Extended Accuracy Models (Altieri et al. 2004). Of these, only the Medium Accuracy Model (MAM) is two-dimensional, but as it is based on simulations that relied on pre-launch measurements of the mirrors, it suffers from a number of deficiencies. The Extended Accuracy Model (EAM) is a one-dimensional model based on in-orbit measurements of real sources, and is considered the most accurate but obviously it does not encapsulate the complex two-dimensional structure observed in the PSF at large off-axis angles. Currently in XCS, we use the EAM when measuring source extents for both real sources and the simulated ones used to create the survey selection function (Section 3.2.3), and when carrying out spatial fits to cluster surface brightness profiles (Section 4.3.2), and we use the MAM when creating simulated sources for the selection function (Section 3.4.3). In the future we hope to include the new 2D model under development by the *XMM* Science Operations Centre (Read et al. 2010). This improved model will more accurately encode the off-axis, azimuthal and energy dependencies of the PSF.

2.2 Data acquisition

In Fig. 2 we illustrate how the non-overlapping area in the public *XMM* archive has grown over the past 10 years, both in terms of total area and in terms of the area suitable for the discovery of clusters, i.e. outside the Galaxy ($|b| > 20°$) and Magellanic Clouds [>6° (3°) of the Large (Small) Magellanic Clouds]. We note that these calculations take into account other, smaller, regions deemed by XCS to be unsuitable for serendipitous source detection (see Section 2.4.1). By now there are over 600 deg^2 of the sky covered by *XMM*, but of that, only ≃50, 280 and 410 deg^2, at >40, >10 and >0 ks depths, respectively (exposure times are those after flare cleaning, Section 2.3.3), are in regions suitable for cluster searching. This area is distributed across the sky (Fig. 3) rather than as a contiguous region.

As shown in Fig. 2, new data enter the archive almost every day, but due to practical constraints we have only processed the data in a small number of large batches, corresponding to all the public
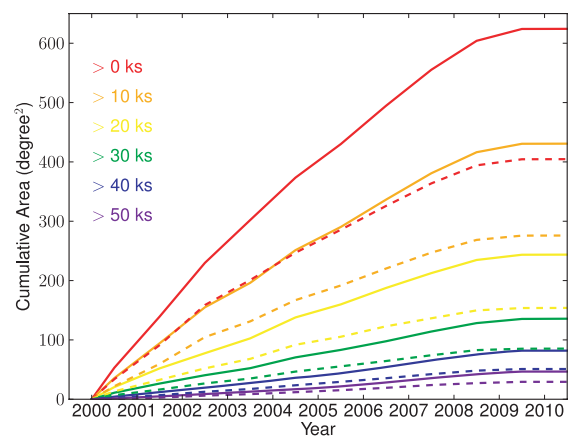

Figure 2. Cumulative sky area covered by public data in the *XMM* archive as a function of exposure time, for the whole sky (solid) and excluding the Galactic plane and Magellanic Clouds (dashed) and for a variety of different exposure time cuts, at the time of the most recent XCS download in 2010 July. The flattening of the curves midway through 2009 reflects the fact that proprietary observations only become public a year after they are completed, so that most data taken after that time were still proprietary at the time of the download.
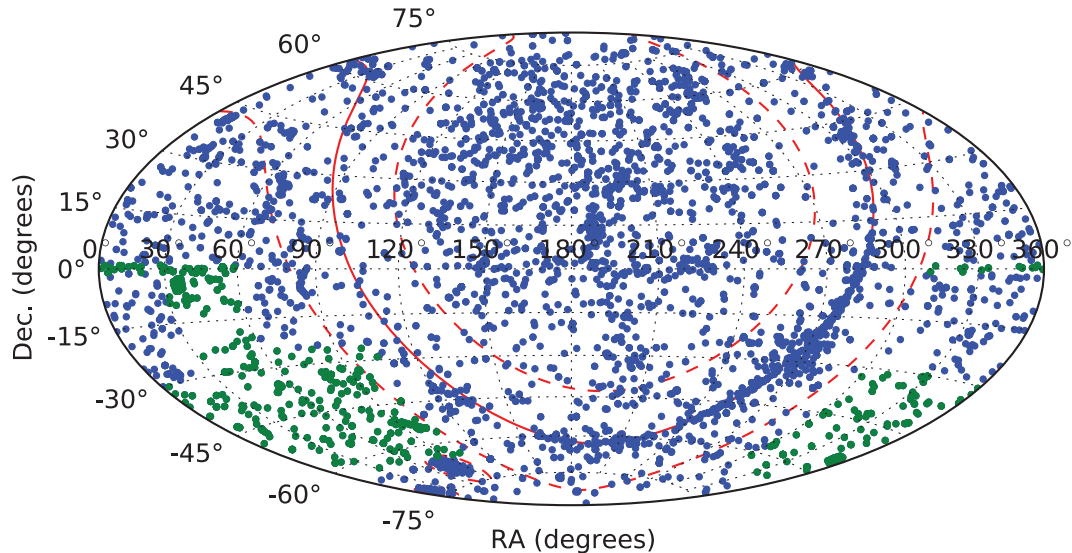
Figure 3. The distribution on the sky of the 5776 ObsIDs in the *XMM* archive as of 2010 July 21. Locations in green (blue) are inside (outside) the proposed footprint of the Dark Energy Survey (darkenergysurvey.org). The Galactic plane and locations of the Magellanic Clouds are highlighted by the red dashed line (we do not carry out cluster searches within those regions).

EPIC data available at that particular time. The downloads take advantage of the Archive InterOperability System (AIO: Arviset, Osuna & Salgado 2004); this protocol allows the *XMM* Science Archive (XSA: Clavel 1998; Arviset et al. 2002) to be searched in an automated fashion. At the time of writing, the most recent download was completed on 2010 July 21, corresponding to 5776 separate *XMM* observations. Their locations are shown in Fig. 3. Each of these observations (including those broken down into multiple exposures) has a unique identification number, or ObsID. In the following, we use the term ObsID to refer to the set of Observation Data Files (ODF) that contains all the observation-specific data. We note that, even with appropriate compression etc., the XCS archive, of raw and processed data products, amounts to of the order of 4 terabytes.

2.3 Data reduction

The data reduction was carried out in a fairly standard manner (see for instance Section 3 of Read & Ponman 2003). Only events with patterns (characterizations of how many CCD pixels are involved in an event) 0–4 were used for the EPIC-pn and 0–12 for the EPIC-mos. A schematic of the data reduction procedure is shown in Fig. B1.

2.3.1 Calibration

The reduction and analysis of *XMM* data requires calibration information detailing how the telescopes and instruments behave, e.g. the effective area of the *XMM* telescopes and the detection efficiency of the instruments (both being functions of photon energy and detector position), plus the instrumental uncertainty associated with measuring photon energies. The most up-to-date version of the *XMM* Current Calibration Files (CCF), as of 2010 July 21, were used for the analysis presented herein.

2.3.2 Software versions

Several different software packages are deployed for XCS analysis: version 10.0.0 of the Science Analysis Software (SAS : Gabriel et al.

2004); version 6.9 of HEASOFT (Blackburn 1995); version 4.2 of CIAO (Doe, Noble & Smith 2001; Deponte Evans et al. 2008); and version 12.6.0i of XSPEC (Arnaud 1996). In order for these packages to be used in the automated batch manner needed for XCS, several different wrapper programmes were written in scripting languages. For the work described in Sections 2 and 4, version 2.6.4 of PYTHON (docs.python.org) was used to write these wrapper programmes, whereas version 7.1 of IDL (www.ittvis.com) was used for the work presented in Section 3.

2.3.3 Flare cleaning

One important aspect of our pipeline reduction was the treatment of background flares. It is well documented (Lumb et al. 2002; Read & Ponman 2003; Pradas & Kerp 2005) that *XMM* observations often suffer from periods of enhanced particle background, caused mostly by variations in solar activity in conjunction with the position of the spacecraft in its orbit. To increase the signal-to-noise ratio of the data, we have designed an automated procedure to remove periods of high background. This was achieved by creating a light curve, divided into 50-s bins. The bin size was chosen to balance a reasonable time resolution with minimizing shot noise. The light curve was first generated, and cleaned, using the high-energy events (12–15 keV for the EPIC-pn and 10–12 keV for the EPIC-mos cameras), because these events are more likely to be from the particle background than from astronomical sources. The cleaning process is then repeated, using a soft-energy light curve (0.2–1.0 keV), to account for periods of elevated background coming from soft protons.

The cleaning process for each energy band involved an iterative 3σ clipping procedure that selected which 50-s bins to exclude. The mean and standard deviation of the light curve were calculated and bins more than $\pm 3\sigma$ from the mean were removed. The 3σ limits were then re-calculated and the process repeated up to 50 times or until a stable state is reached, whereby the bins that are being excluded are not changing (note that previously excluded bins can be reinstated in subsequent iterations if the 3σ limits become larger). The maximum of 50 iterations was set to avoid
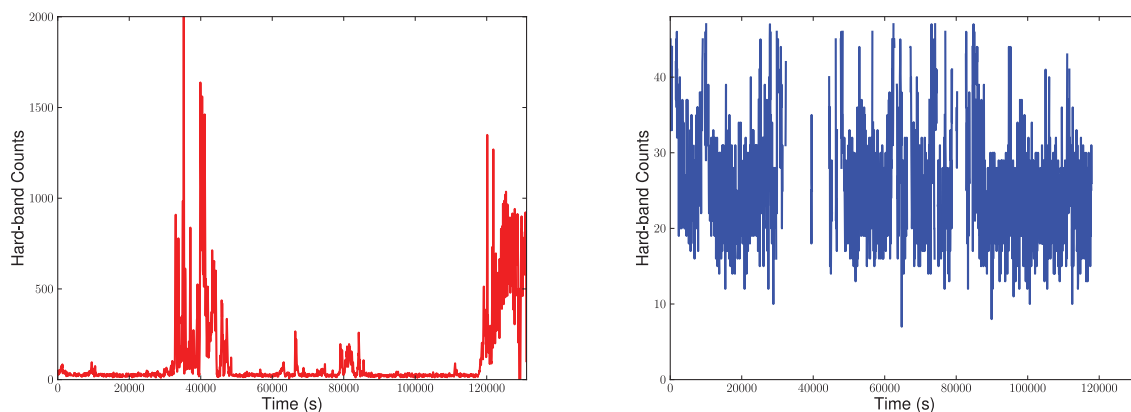
Figure 4. EPIC-pn example hard-band light curve with 50-s bins. Left panel: raw events before cleaning. Right panel: cleaned events with periods of high background removed.

cases where the stable solution oscillates between two or more similar states.

We note that before the first 3σ clipping took place, an initial maximum-rate threshold is used to 'clip' the light curve. This threshold is the greater of either 50 counts per bin for the EPIC-pn (and half this for the EPIC-mos cameras) or 125 per cent of the highest value in the lowest 5 per cent of the bins. This initial filtering was found to improve the flare cleaning results when flares accounted for a large fraction of the total exposure time. A flow chart illustrating the flare-cleaning steps is shown in Fig. B2. Fig. 4 shows an example hard-band light curve before and after cleaning.

The combination of the excluded bins for the hard and soft-background light curves is then used to define the good time intervals (GTI) used to filter the raw event files. Fig. 5 shows the distribution of ObsID exposure times before and after the process of flare cleaning. The filtered event files are used several times during XCS analysis. They are used to produce the images (Section 2.4), used for the initial XCS source detection (Section 3.1), and then again to determine spectroscopic (Sections 4.2 and 4.4) and spatial parameters (Section 4.3) for the cluster candidates.

2.4 Image production

Starting with the cleaned event lists described above (Section 2.3.3), the individual camera exposures were spatially binned, with a pixel size of 4.35 arcsec, to generate images. This pixel size was chosen because it is smaller than the PSF, at all detector locations and photon energies. Images were produced in two bands, soft (0.5–2.0 keV) and hard (2–10 keV). Exposure maps were also created for each image. The exposure maps encode the impact of vignetting on the image sensitivity and also record the locations of chip gaps, bad rows, etc.

The EPIC cameras do not have shutters, so events received while an observation is reading out, the so-called out-of-time events, will be assigned incorrect positions and energies. For XCS, only EPIC-pn images were corrected for out-of-time events, because the EPIC-mos cameras have a much lower readout rate and negligible out-of-time events. The EPIC-pn corrections were done in the standard way, i.e. the event file was recreated assuming all the events are out-of-time and assigning them new positions along the CCD column at random. These are then used to create out-of-time images that can be subtracted off the true images (with the appropriate correction for the fraction of out-of-time events).

The images and exposure maps for the individual cameras were merged to create a single image and exposure map per ObsID. For this, the pixel values in the EPIC-mos maps were scaled to that of the EPIC-pn camera using the previously calculated ECFs (Section 2.5). Examples of XCS-generated exposure maps and images can be seen in Figs A1 and 6. A total of 5642 image files have been generated from the 5776 *XMM* ObsIDs that make up the current XCS data
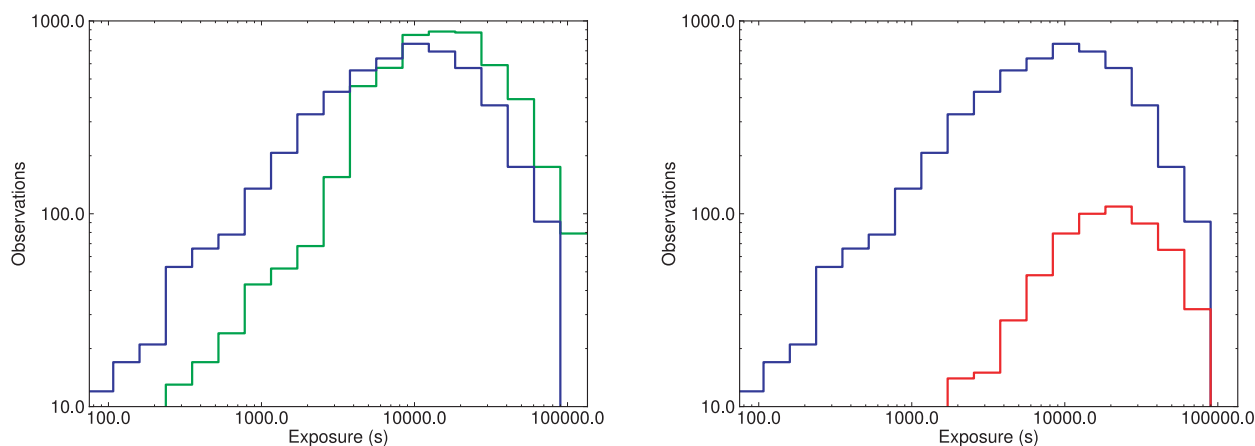

Figure 5. The distribution of ObsID exposure times. Left panel: the number of ObsIDs before (green) and after (blue) the process of flare cleaning. Right panel: the number of ObsIDs in which extended XCS sources with 300 or more counts were detected (red), compared with all ObsIDs (after flare cleaning).
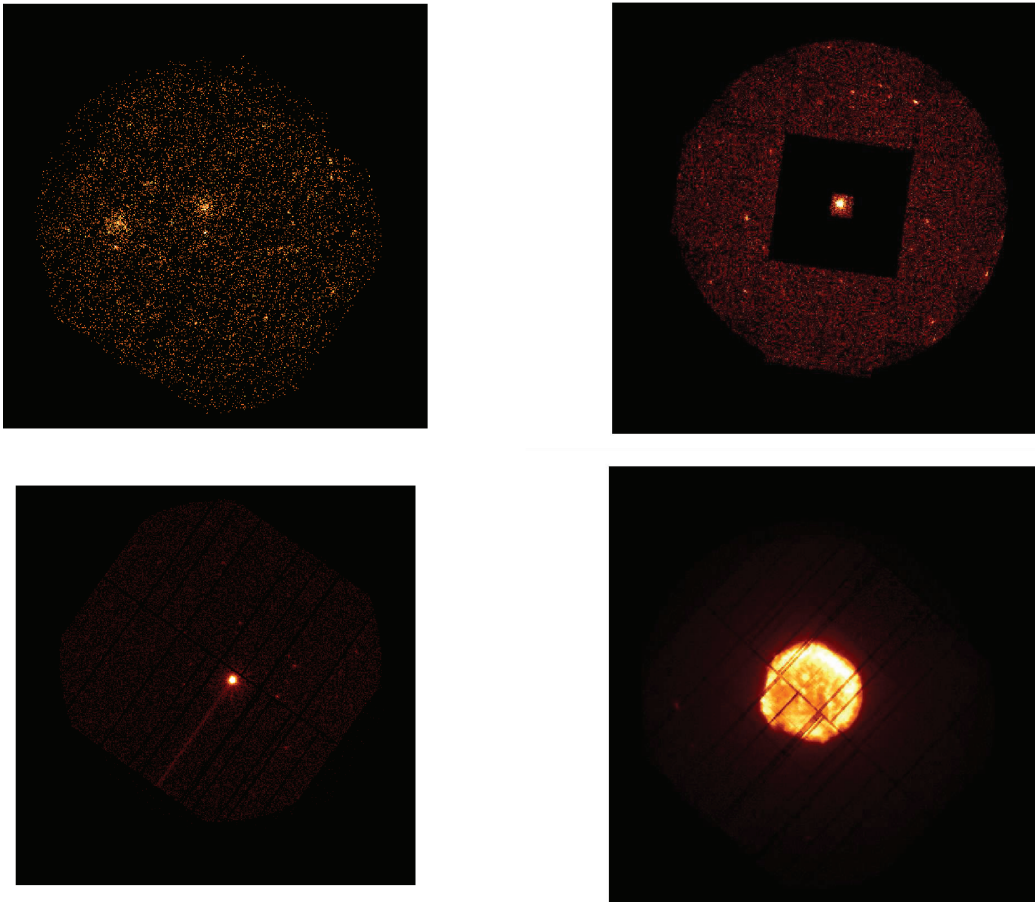
Figure 6. Examples of reduced and merged *XMM* images with a variety of different target types.

set (a small number of ObsIDs in the archive are not suitable for automated image generation for a variety of technical reasons such as telemetry and calibration issues, etc.).

2.4.1 Image masking

The production of images is an automated process; however, they do need to be checked by eye before passing them to the source detection pipeline (Section 3.1). This is because we download all public data, regardless of the intended (by the PI) target. As a result, the XCS image archive includes ObsIDs with very extended targets (such as low-redshift clusters or Galactic supernova remnants) and ObsIDs with very bright targets (such as luminous AGN). The very extended targets will enhance the background level over the majority of the *XMM* field of view, and thus reduce our ability to make serendipitous detections of sources. The very bright sources will generate artefacts in the images, such as radial spikes and out-of-time bleed trails; those artefacts could then be falsely identified as additional sources. The eyeballing process identifies ObsIDs that should be completely excluded from the other stages of the XCS pipelines. It also allows us to mask out regions of ObsIDs that are only partially afflicted by bright/extended targets. Approximately one-third of ObsIDs require some degree of masking, with the median area lost being around 4 per cent (though this can be as high as 80 per cent in extreme cases). The mask files are of the same dimensions as the image files and are used during the source detection and also when creating backgrounds for the spectral and spatial fitting. We show some examples of XCS images that require full or partial masking in Fig. 7.

2.5 Energy conversion factors

In order to be able to convert image source counts into energy fluxes, energy conversion factors (ECFs) need to be calculated. These are necessarily model dependent and are affected not only by the source and instrument properties but also by the H I column, n_H hereafter, along the line of sight. In our survey, the source properties are not known in advance, so a generic model has to be assumed. Since the vast majority of the sources detected by XCS are point sources, and point sources are likely to have power-law spectra, the model used to calculate the conversion is an absorbed power law with a canonical AGN index of 1.7 (Mushotzky, Done & Pounds 1993). The photoelectric absorption is set to the appropriate n_H value for the field (Section 2.5.1). The ECFs were calculated, using the XSPEC spectral fitting package and the on-axis spectral responses, for each camera exposure related to a particular ObsID. For the specified model, the ratio of the resulting flux and count-rate is stored as the ECF for that exposure. ECFs are not exposure time dependent, but due to variations in n_H, the choice of optical blocking filter and the effective area of the instrument, ECFs in XCS still vary from exposure to exposure and from ObsID to ObsID. They generally range from 4.4 to 6.6 for the EPIC-pn and from 1.6 to 2.0 for the EPIC-mos cameras (in units of 10^{-11} count cm^{-2} erg^{-1}). Even though the ECFs are calculated for the on-axis aim point, they can

Figure 7. Examples of masks created for the two lower images in Fig. 6.

still be used for sources detected anywhere in the field of view, by correcting them using the exposure map.

We also calculate, for each ObsID, a further set of conversions using the MEKAL model (Mewe, Lemen & van den Oord 1986). The MEKAL model is the standard model used to describe thermal and line emission from clusters of galaxies. The MEKAL conversions are done over a grid of n_H, temperature and redshift; however, the metal abundance is kept fixed at $Z = 0.3 \times$ the solar values in Anders & Grevesse (1989). [This choice of metallicity is standard in the field because previous work, such as the one by Maughan et al. (2008), has shown that abundances vary little from this value over a wide range of redshifts.] The gridded MEKAL conversions can be used to convert count-rates to bolometric luminosities and vice versa (and we refer to these conversions as LCFs hereafter). The LCFs are used to calculate synthetic cluster count-rates for the survey selection function (Section 3.4.3) and to estimate luminosities for XCS candidates during the literature redshift search (Section 4.1). The LCFs, like the ECFs, are calculated for the on-axis aim point, but can be adjusted to another location using the exposure map.

2.5.1 Galactic H I column

X-ray photons are absorbed by material along the line of sight, and in particular by helium and oxygen for photons above ~ 0.5 keV (Wilms, Allen & McCray 2000). One can predict the level of absorption if n_H is known, so, for XCS, we estimated the n_H values for each source using the compilation of Dickey & Lockman (1990), which combines the Bell Labs H I Survey (Stark et al. 1992) data with other surveys for the all-sky coverage. We use n_H to calculate ECFs and LCFs (see above), but also at other points in the XCS pipeline, e.g. when analysing X-ray spectra (Section 4.2). We note that self-shielding of molecular hydrogen, from ambient ultraviolet radiation, can occur when $n_H > 5 \times 10^{20}$ cm^{-2} (Arabadjis & Bregman 1999). This molecular gas absorbs X-rays and thus distorts flux conversions that are based only on n_H values. For this reason, XCS fluxes derived when $n_H > 5 \times 10^{20}$ cm^{-2} should be regarded as lower estimates.

3 GENERATION OF THE XCS SOURCE CATALOGUE

In this section, we provide details of our source detection algorithm, known as the XCS Automated Pipeline Algorithm or XAPA. In Section 3.1, we explain how XAPA applies wavelets to the pipeline

generated images (Section 2.4) to generate a source list per ObsID. In Section 3.2, we describe the parameters that are measured by XAPA for each detected source. In Sections 3.3 and 3.4 we demonstrate the quality of the XAPA data products for point and extended sources, respectively. In Appendix B we provide the related flow charts.

3.1 Source detection

XAPA source detection is based upon the mission-independent source detection package WAVDETECT (Freeman et al. 2002, F02 hereafter), which is available as a part of the CIAO software package. F02 have shown that WAVDETECT 's wavelet-based algorithm is more sensitive than standard sliding-cell algorithms (e.g. CELLDETECT from CIAO, Fruscione et al. 2006) and is considerably faster than event-list-based algorithms such as CIAO 's VTPDETECT. Before deciding to use WAVDETECT as the basis for the XAPA algorithm, we also examined the *XMM* SAS EWAVELET program and the SEXTRACTOR package (Bertin & Arnouts 1996), finding both of them to be inadequate for our purposes (see Davidson 2006, for a discussion).

The F02 version of WAVDETECT consists of two components, `wtransform` and `wrecon`. The former convolves binned images with Mexican Hat (Slezak, Bijaoui & Mars 1990) wavelet functions with various user-specified scale sizes and then identifies pixels that are significantly above the background. In XAPA, we use the F02 version of `wtransform` as part of an automated pipeline known as `md_detect`,[1] as illustrated by the flow chart of Fig. B3. We use a set of nine wavelet scales, numbered according to increasing size, and corresponding to $\sqrt{2}$, 2, $2\sqrt{2}$, 4, $4\sqrt{2}$, 8, $8\sqrt{2}$, 16 and 32 image pixels. On each scale, the convolved image is compared with a threshold image. Convolved image pixels with values greater than their corresponding threshold image pixels are assumed to be associated with astronomical sources ('significant pixels' hereafter). For those pixels, we reject the null hypothesis that they are consistent with the measured background. We then generate a set of support images, which record the significant pixels at each wavelet scale.

In order to enhance the detectability of faint extended emission, `md_detect` performs the wavelet analysis in two stages (or 'Runs'). In Run 1 (scales 1–2), bright compact sources are located first. These are then masked out before performing Run 2 (scales 3–9). The masking step was found to be necessary because bright

---

[1] Where the md_ prefix acknowledges the architect of the routine, Michael Davidson.

point sources can pollute the wavelet signal on large scales, and hence mimic extended sources. Unfortunately, this masking can occasionally result in genuine extended sources being excluded from the candidate list, so an extra step was added to XAPA to mitigate this effect (Section 3.1.1).

The second component of the F02 version of WAVDETECT is `wrecon`. This generates a source list for each image, by grouping collections of significant pixels together into source regions, or 'cells'. A drawback of the F02 version of `wrecon` is that it uses the instrument PSF to define the size of the cells. This means that extended sources can be broken up into multiple contiguous 'sources' (because a single PSF-sized cell is not big enough to enclose all the flux). To overcome this problem, we wrote a modified version of `wrecon`, called `md_recon`, for XAPA. Unlike `wrecon`, `md_recon` does not assume a priori the size of the detected sources, and is consequently considerably better at fitting ellipses to extended sources. The operation of `md_recon` is as illustrated by the flow chart in Fig. B4. At each wavelet scale, `md_recon` first combines lists of significant source pixels into source cells. Multiscale objects, i.e. those detected by `md_detect` on multiple scales, are then filtered using a 'vision model' (Section 3.1.2). The vision model is a set of rules for combining the support images derived for different wavelet scales. The vision model is able to recognize when a point source is embedded in an extended source. It also fits elliptical regions to the recovered sources (the region enclosed by a source ellipse is referred to as ϵ_f in the following descriptions).

3.1.1 Extended sources with central cusps

The two-step (Run 1, Run 2) procedure adopted by `md_detect` for source detection works well, in that it prevents bright point sources from contaminating the extended source list. However, it has the disadvantage that when a genuine extended source is detected in Run 1, it will be excluded from Run 2. This means that its size will be underestimated by the vision model, and it will not appear in an extended source list. Extended sources with cuspy brightness profiles will be particularly affected by this, e.g. clusters with cool cores. We have therefore devised a 'cuspiness test' that is carried out between Run 1 and Run 2. This involves generating a grid of 5-by-5 pixels, Q, centred on the position of each source detected in Run 1. A quantity, C, representing the cuspiness of the central region, is then calculated as follows:

$$C = \frac{Q_{\text{max}} - Q_{\text{min}}}{Q_{\text{max}}}. \tag{1}$$

Tests showed that real point sources have $C \geq 0.85$, so if a Run 1 source is found to have $C < 0.85$ – i.e. it possesses a flatter central

profile than a real point source – it is removed from the list of Run 1 detections, resulting in it being available to be detected again in Run 2. This situation is illustrated by Fig. 8.

3.1.2 The XAPA vision model

Here we give more details about the vision model used to filter sources detected at multiple scales by `md_recon`. To describe our vision model we introduce the following two terms: a 'structure' is a connected set of pixels in the support image for a particular scale; and an 'object' is a set of connected structures from different scales. The steps are as follows.

(i) For each structure, comprising a set of pixels $\{(x, y)\}$ in S_i which is the support image for scale i, determine whether the structure defines the 'root' of an object, i.e. whether $S_j(\{(x, y)\}) = 0$ for all $j < i$.

(ii) For each such root, check to see if there is a structure in the scale above at this position, i.e. if $\exists (x', y'), (x', y') \in \{(x, y)\}, S_{i+1}(x', y') \neq 0$.

(iii) If such a structure exists, and its maximum pixel value lies within $\{(x, y)\}$, then these two structures are linked, such that the image pixels belonging to the object comprise the union of the pixels in the linked structures from scales i and $i + 1$.

(iv) The process of upward linking continues until the condition in step (ii) is not satisfied, at which time the object is terminated. When each scale has been scanned for root structures and they have been propagated in the 'tree-like' fashion, then for each object created there exists a set of image pixels belonging to it. An ellipse can then be fitted to these regions and a source list created.

This vision model can handle both point and extended sources. Crucially, it can also cope with point sources embedded in extended sources, and with close pairs of points which should be separated rather than blended. A schematic to illustrate how the vision model works when a point source is embedded in an extended source can be seen in Fig. 9.

3.2 Source properties

Once `md_recon` has been run on a given image, the source list is passed on to the next part of the XAPA pipeline, `find_srcprop`. The two-stage operation of `find_srcprop` is illustrated by the flow chart in Fig. B5. In the first stage, `find_srcprop` determines the significance of each detected source (Fig. B6). In the second, a subroutine known as `find_srcprop_final` computes other source properties (such as the count-rate and probability of extent); it is
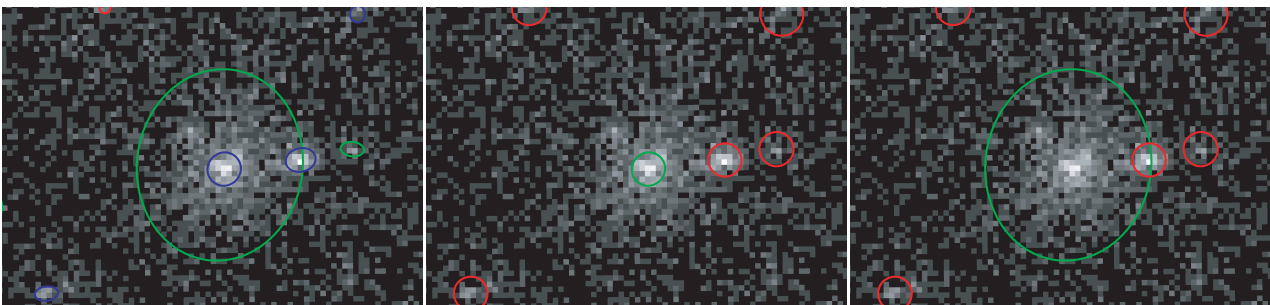

Figure 8. Illustration of the effect of extended source cuspiness. Left: the original (before the cuspiness test was introduced) Run1 (blue) and Run2 (green) detections. Middle: the final source list if the cuspiness test is not performed. Right: the final source list (after the cuspiness test was introduced). Extended and point sources have green and red outlines, respectively.
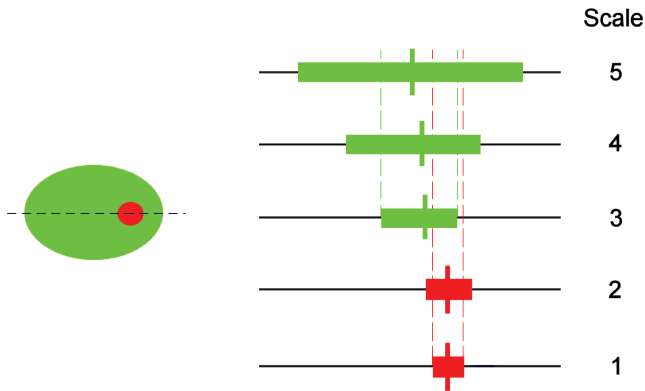
Figure 9. Illustration of the 'tree' vision model. Left: the source configuration showing a point source embedded in a larger source. The dashed line indicates a 1D cut through the sources. Right: a schematic of the significant pixels at each scale showing how the structures are connected to form objects. The vertical bars denote the position of the maximum coefficient at each scale. The maximum of scale 3 lies outside of the structure of scale 1 hence a new object is started.

the results from the `find_srcprop_final` that appear in the XCS data tables (Section 3.2.7).

3.2.1 Measuring source and background counts

Here we describe how background-corrected source counts were calculated in XAPA by `find_srcprop` and by its subroutine `find_srcprop_final`. Tests during the development of XAPA showed that the best results were obtained using different aperture sets for each stage. The aperture set comprises the region for source flux determination, the region for background flux determination and a masked region (which is not used for either). In Table 1 we note the configuration for both the aperture sets. In specifying these, we denote by ϵ_f the ellipse as fitted to the object region, so that $3\epsilon_f$ is the ellipse with major and minor axes three times those fitted to the source by the vision model. We use Uniq(X) to denote, for a particular source, those pixels that lie only within region X defined relative to that source: e.g. Uniq($3\epsilon_f$) defines, for a particular source, the set of pixels which lie within the

Table 1. Mask and aperture configurations for source and background flux determination used in `find_srcprop` and `find_srcprop_final`.

Type	Configuration (`find_srcprop`)
Run 1	Mask: Run 1 sources masked at $2\epsilon_f$
	Flux: $1\epsilon_f$+Uniq($3\epsilon_f$)
	Background: Inner radius at $2\epsilon_f$, min. area = 400 pix
Run 2	Mask: All sources masked at $3\epsilon_f$
	Flux: $1\epsilon_f$+Uniq($3\epsilon_f$)
	Background: Inner radius at $3\epsilon_f$, min. area = 2000 pix

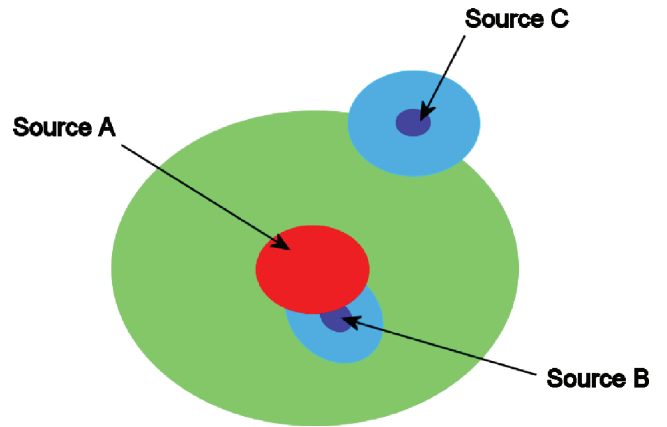| Type	Configuration (`find_srcprop_final`)
Point	Mask: Point sources masked at $2\epsilon_f$
	Flux: $1\epsilon_f$
	Background: Inner radius at $2\epsilon_f$, min. area = 400 pix
Extended	Mask: All sources masked at $3\epsilon_f$
	Flux: $1\epsilon_f$+Uniq($3\epsilon_f$), with internal point sources masked at $1\epsilon_f$
	Background: Inner radius at $3\epsilon_f$, min. area = 2000 pix

Figure 10. A diagram showing how the aperture used to measure source flux is created. The source to be measured is *Source A* and there are also two other objects nearby (*Source B* and *Source C*). Both the $1 * \epsilon_f$ and $3 * \epsilon_f$ ellipses are shown for each source (red and green respectively for *Source A* and dark blue and light blue for *B* and *C*. Hence, the area used to calculate the flux for *Source A* is the red plus the green region.

$3\epsilon_f$ region for that source and for no other source (as illustrated in Fig. 10).

The expected background contribution is computed locally. An elliptical annulus is placed around the source position: the inner edge varies but is usually at $3\epsilon_f$ and the outer edge is increased until there are at least 2000 background pixels, or no more area is available. The background count-rate, b_{pix}, is then calculated as $b_{pix} = B/\bar{E}' \times a'$, where B is the total number of counts in the annulus, $\bar{E}'$ the mean exposure in the annulus and a' is the number of pixels in the annulus. The expected number, B_a, of background counts within the source aperture is then computed as $B_a = b_{pix} \times \bar{E} \times a$, where $\bar{E}$ is the mean exposure in the source aperture and a the number of pixels in the aperture.

3.2.2 Removing low-significance sources

The first task is to remove any sources which are statistically of low significance, because they will not yield accurate properties. The source and background apertures used to determine this significance must be chosen carefully (Section 3.2.1), but once the expected number of background counts, B_a, within the source aperture, ϵ_f, is known, it is possible to assess the significance of the detected source. This is done by computing the probability that the background could, by chance, produce the detected number of counts in the source aperture, assuming a Poisson distribution for the background counts, with mean B_a. Those sources with a probability higher than 0.000 032 are removed from the source list: this probability is equivalent to a 4σ threshold for a Gaussian distribution. In addition, detections comprised of only a single significant pixel are excised from the source list, regardless of their significance. These are likely to be hot pixels or sources that are too faint to be accurately parametrized.

3.2.3 Measuring source extents

After low-significance sources have been removed, the `find_srcprop` routine is run again on the sources above the $\geq 4\sigma$ threshold, in order to classify them as point-like or extended. For

this, we need to compare the sources to the instrument PSF. Unfortunately, no satisfactory 2D PSF model for *XMM* exists (Section 2.1.1), so for XCS we adopted the best publicly available 1D (radially averaged) model – the EAM. This, in turn, necessitated the development of a source classification criterion based on a 1D source property. For XCS, we used the encircled energy fraction (EEF). The EEF records the fraction of the total energy of a source as a function of increasing (circular) aperture size. We note that even though the shape of the PSF changes considerably towards large off-axis angles, its radial average, the EEF profile, is only a weak function of the off-axis angle (Davidson 2006), making it a good basis for a classification criterion to be applied across the full field of view.

Our extent classification is based on testing the null hypothesis that the measured EEF for a source is consistent with the PSF EEF, at the appropriate off-axis angle. This is implemented using a Kolmogorov–Smirnov (K-S) test, using the EEF profile of the source and a model-merged PSF EEF. The PSF EEF is derived from EAM EEFs produced by the SAS task CALVIEW from the Current Calibration Files (CCF) for each camera. This is weighted by the Energy Conversion Factor (ECF, Section 2.5) appropriate for that ObsID. We adopt for $P(point)$ the probability that the source is point-like, the maximum value of the probability returned by the K-S test run on a 3×3 pixel grid (with spacing ± 0.5 pixels in x and y) around the source position (in Section 3.3.1 we show that the typical positional accuracy of XCS source centroids are good, even better than 1 pixel).

The reliability of the $P(point)$ values is a function of several factors, including the position on the field of view, the background level, the number of source counts and the proximity of neighbouring sources. For that reason, choosing a fixed threshold in $P(point)$ for our classification would be inappropriate. Instead we are forced to conduct a series of Monte Carlo (MC) simulations for every source: this is computationally expensive, but it is vital to prevent misclassification. This simulation process involves generating 200 realizations of the appropriate PSF EEF model and populating them with the same number of counts as measured in the data. Each of the 200 realizations is compared to the model and an empirical distribution of the K-S d values is established. If none of the simulated distributions returns a d value as great as the measured value, we classify the source as being extended. With this procedure, the statistical probability of misclassifying an isolated point source as extended is 0.005 or less. However, we note that this does not take into account systematics, such as when two or more point sources have been blended by XAPA into a single source profile. These can only be removed a posteriori, by eyeballing the extended sources that make it through to the cluster candidate list. This eyeballing, or quality control (see Fig. 1), process is described in more detail in Mehrtens et al. (in preparation).

3.2.4 Correcting artefacts

After the second pass of `find_src_prop` has been completed, we have a preliminary list of sources (classified as extended or point-like) for a given ObsID. Initial tests showed that these preliminary lists include a number of artefacts. These must be corrected for before the inclusion in an XCS source catalogue (see below). The corrections are not foolproof, as not all genuine clusters make it through to the candidate list and not all contaminating sources are excluded, but because the corrections are folded into the survey selection function (Section 3.4.3), they should not
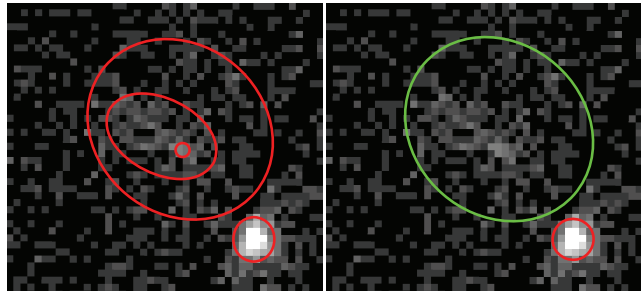

Figure 11. An example where several initial detections of an extended source are subsequently merged by XAPA to improve the derived properties.

impact our ability to use XCS cluster catalogues for statistical studies.

XAPA's `md_recon` algorithm successfully detects sources within sources (see Fig. 9). However, one unintended consequence is the occasional multiple detection of a single source that has got split into two or more overlapping sources. This more often happens with extended sources, but can also occur with point sources at the edge of the field of view. Therefore, where there are incidences of two sources with overlapping cells, the sources are merged and source properties recalculated by `find_srcprop` (see Fig. 11). This refinement ensures that in most cases the source flux and morphology are recovered well.

When a bright compact source lies on the outskirts of the field of view, it can produce a significant number of counts in the asymmetric outer regions of the PSF. We term these objects as 'point-sources-with-lobes'. The core of such sources are detected in Run 1 of `md_detect`, and hence the core counts will be masked from Run 2 (Section 3.1), but the remaining outer counts might still yield a Run 2 detection (see Fig. 12). Removing these point-sources-with-lobes, without also removing clusters with cuspy cores (Section 3.1.1), proved to be one of the most difficult problems to overcome with XAPA. After extensive tests, we arrived at the following compromise: an extended source is excised from the source list, as a suspected point-source-with-lobe, if it is both located within the $3\epsilon_f$ region of a Run 1 source, and has less than one-fifth of the counts of that source. This removes the majority of the lobe artefacts, but can
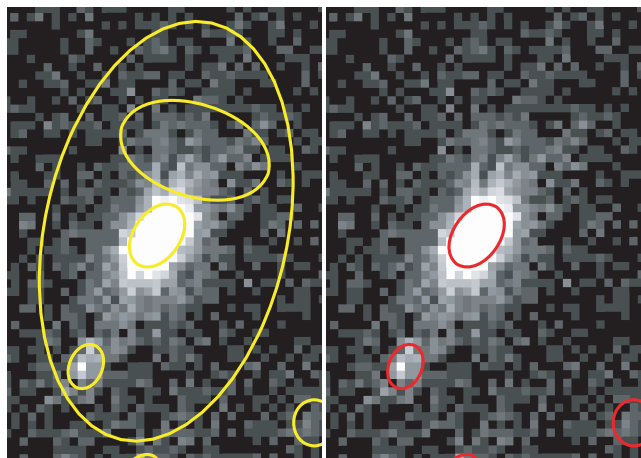

Figure 12. Source ellipses defined by XAPA for a bright, off-axis, point source. Left: before the lobe removal step was included. Right: after the lobe removal step was included: note that the two point sources have still been recovered, but there is no erroneous large (extended) ellipse enclosing both of them.

unfortunately also result in some genuine faint extended sources being excluded from the XCS cluster candidate list.

3.2.5 Extended source flags

When developing XAPA, we had to find a compromise between contamination and completeness, i.e. between effective and over-cleaning of the extended source list. Therefore, rather than removing from the extended source list every object that could be erroneous, we have flagged certain sources that, conservatively, we view as suspicious. Our aim is to use the survey selection function (Section 3.4.3) to help us understand whether flagged sources should be included in statistical studies or not, but to date we have taken a conservative approach and not included them in cluster candidate lists, or as targets for optical follow-up (Mehrtens et al., in preparation). The source flags are as follows.

(i) *Extended sources that are PSF-sized.* At large off-axis angles it is not infrequent for the flaws in the PSF model to cause an obvious, bright, point source to be classified as extended. Therefore, any source that is just extended (i.e. that has a size very close to the PSF at the respective off-axis angle) is flagged as being 'PSF-sized' by XAPA.

(ii) *Extended sources with internal point sources.* Even with the inclusion of the point-source-with-lobe test (Section 3.2.4), the XAPA vision model (Section 3.1.2) will occasionally misclassify flux from the outskirts of a point source (or flux from a collection of neighbouring point sources) as an erroneous extended source. We can mitigate this by flagging up likely incidences. Therefore, any extended source region that encloses one or more point sources that contribute ≥ 1.3 times the extended source flux is flagged as being 'point contaminated' by XAPA.

(iii) *Extended sources with internal Run1 sources.* The final flag is similar to the 'point contaminated' case, but covers the incidences of genuine point sources, detected in Run 1 by `md_detect`, being erroneously passed on to Run 2 by the cuspiness test (Section 3.1.1). Therefore, any extended source region that encloses one or more Run 1 detection regions that contribute at least half the extended source flux is flagged as being 'Run 1 contaminated' by XAPA .

3.2.6 Source parameters

Once the source list per ObsID has been cleaned of artefacts, a file is generated that saves all the relevant data. This file is then interrogated when the survey-wide data base is being generated (Section 3.2.7). The following attributes are saved per source:

(i) the centroid location in image coordinates;

(ii) the centroid location in sky coordinates (J2000);

(iii) the centroid location in radial coordinates, i.e. the off-axis angle (arcminutes) and the azimuthal angle (degrees);

(iv) the major axis, minor axis and orientation of the source ellipse;

(v) the average exposure time at the source location (seconds);

(vi) the 0.5–2.0 and 2–10 keV background-subtracted source counts (in the merged image and in the individual camera exposures);

(vii) the 0.5–2.0 and 2–10 keV background-subtracted count-rates and 1σ count-rate uncertainties (in the merged image and in the individual camera exposures);

(viii) the source significance and extent probability;

(ix) the value of the source flags (see Section 3.2.5).

3.2.7 Master detection list

XAPA produces a source list for each of the input ObsIDs, then these lists are concatenated to form a Master Detection List (MDL). Present in the *XMM* archive are many areas that have been observed multiple times. As a result, some sources will have been detected by XAPA multiple times. When duplicates are found, only the detection with the most soft-band counts is passed to the MDL. To remove duplicates, it is necessary to set an appropriate matching radius. The positional accuracy of the survey is higher for point sources than for extended sources, so it makes sense to use a different radius for each type. The accuracy for point sources varies as a function of off-axis and azimuthal angles (amongst other parameters). However, for simplicity we use a single value for the radius of 5 arcsec. The case for extended sources is less straightforward because of the variety of source types and morphologies. The positional accuracy for large diffuse objects, such as low-redshift clusters, can be very poor, making it hard to pick an appropriate radius. Fortunately, the largest diffuse sources should have already been masked from their host ObsID. So, for XCS, we use a fixed matching radius of 30 arcsec for extended sources. This radius is large enough to allow reliable source matching, but small enough to minimize the removal of genuine cluster candidates.

As of 2011 May 01, XAPA had run on 4029 ObsIDs, resulting in 114 711 point sources and 12 582 extended sources being included in the MDL. Of the 12 582 extended sources, roughly half were flagged (Section 3.2.5) and these were removed from the list of potential cluster candidates (leaving 6983 sources). Additional cuts to this list included the removal of sources within $20°$ of the Galactic plane and $6°$ [$3°$] of the Large [Small] Magellanic Cloud. Those cuts were made because it can be hard to carry out effective optical follow-up in regions of high projected stellar density. Moreover, the closer one gets to the Galactic plane, the higher the hydrogen column (large n_H values impact our ability to recover accurate source fluxes). Further cuts, see below, are then imposed to ensure that the vast majority of XCS cluster candidates are genuinely serendipitous detections, rather than the intended target of the ObsID. A final cut, on minimum source count (>50), is then applied, leaving 3675 sources drawn from 1533 different ObsIDs; when we use the term 'candidate' hereafter, we are referring to these 3675 sources. The candidates have a range of counts, from 50 to several thousand. Of particular interest to the cosmology and evolution studies we plan with XCS are the 993 with more than 300 counts, because these should deliver, once redshift information is available, reliable temperature estimates (Fig. 17).

As mentioned above, filters were applied to exclude non-serendipitous or 'target' objects from the candidate list. The targets in question are primarily clusters, but other types of extended X-ray sources should also be excluded (e.g. low-redshift galaxies). It is also important to identify extended sources that are physically associated with the target, e.g. if they both belong to the same supercluster. The target filters involved both checks of the ObsID file headers and automated queries to the NASA Extragalactic Data base (NED).[2] The filters were run separately on each ObsID that a particular extended source was detected in. A given extended source (that passed the other cuts described above) was included in the candidate list if it was classed as being a serendipitous detection in at least one of those ObsID (even if it was classed as being the target of one or more of the others). We acknowledge
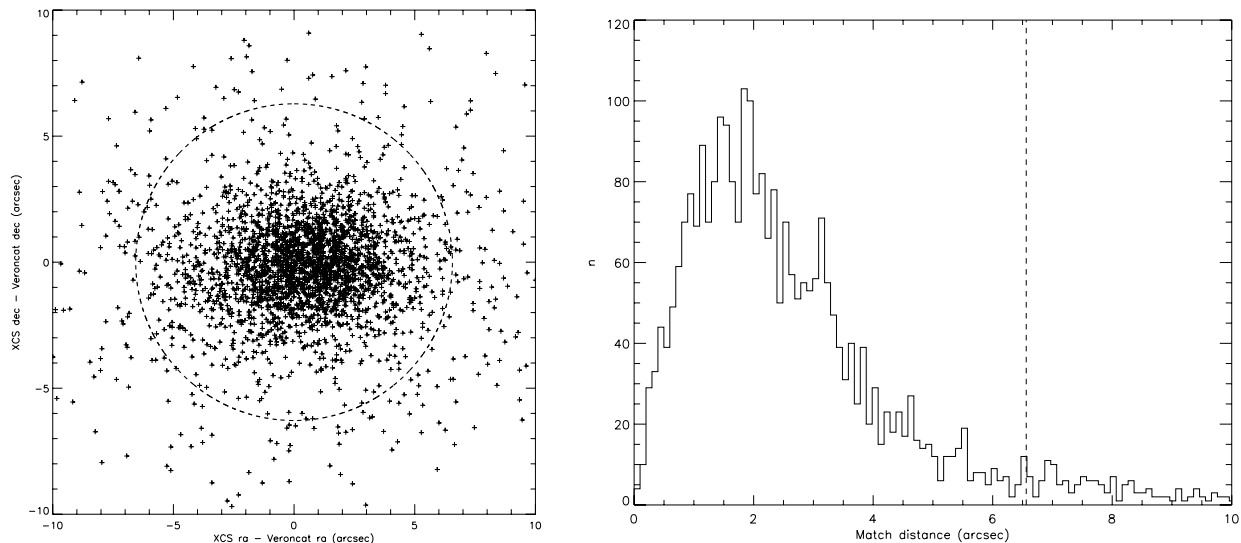
---

[2] http://nedwww.ipac.caltech.edu

Figure 13. The relative position of the matches between XCS and VeronCat (Véron-Cetty & Véron 2006) source positions. The dashed line represents the 95 per cent matching radius.

that some *XMM* targets, and some sources associated with targets, do make it through into our candidate list. However, as shown in Mehrtens et al. (in preparation), these are straightforward to remove at the quality control stage (eight such examples were removed from XCS-DR1).

Extended sources were excluded from the candidate list if they met one or more of the following criteria.

(i) Their XAPA centroid fell within 2 arcmin of the aim point of an ObsID with an object classification (as listed in the header) of 'cluster' or 'group'.

(ii) Their XAPA ellipse overlaps the aim point of an ObsID with a target name (as listed in the header) that has been associated with a cluster or group in NED. (This filter is necessary either when the pointing type is not included in the header, or is incorrect.)

(iii) Their XAPA ellipse overlaps the centroid of a cluster or group in NED, when the aim point of the ObsID falls within 2 arcmin of that cluster or group. (This filter is necessary because sometimes non-standard target names are listed in the header.)

(iv) Their redshift is within $5000\,\mathrm{km\,s^{-1}}$ of the redshift of an object in NED, when the target name (as listed in the header) has been associated with that object. Both redshifts are automatically extracted from NED. [This filter is necessary because some ObsID targets are deliberately positioned off-axis. This filter also reduces the number of sources entering the candidate list that are physically associated with targets (including with non-cluster targets, such as AGN).]

(v) Their XAPA ellipse overlaps the aim point of an ObsID with a target name (as listed in the header) that has been associated with a known galaxy in NED. (This filter was found to be the most effective way to exclude low-redshift galaxies from the candidate list.)

3.3 XAPA verification: point sources

As mentioned above (Section 3.2.7), XAPA has catalogued to date in excess of 100 000 unique point sources. In this section we test XAPA astrometry and flux measurements using these point sources, finding both measures to be robust.

3.3.1 Positions

To determine the positional accuracy of the XCS point sources, it is desirable to use a catalogue that has a high spatial resolution and a high astrometric precision. It would also need to have a significant overlap with the *XMM* archive. A natural choice for this is the Sloan Digital Sky Survey (SDSS,[3] Abazajian et al. 2009); the data are of high quality and contains many objects that would be expected to have X-ray counterparts, e.g. quasars and AGN. A cross-match of XCS point sources against the SDSS Quasar Catalog IV (Schneider et al. 2007) using a radius of 10 arcsec produces 1131 matches. This was extended further using the catalogue of Véron-Cetty & Véron (2006, VeronCat hereafter). VeronCat is a compilation of all the known AGNs and QSOs (including those in the SDSS). A 10-arcsec matching radius returns 2807 matches, the distribution of which can be seen in Fig. 13. We have determined the chance of false association between the XCS and VeronCat with a 10-arcsec matching radius to be 1 per cent. The mean matching distance is 2.6 arcsec, and 95 per cent of the matches fall within 6.6 arcsec. This level of precision is consistent with previous determinations of the positional accuracy obtainable with *XMM* data (Watson et al. 2009).

3.3.2 Fluxes

To assess the accuracy of the point source fluxes measured by XAPA, we have compared the XCS point source list to the *XMM* Serendipitous Survey 2XMM catalogue (Watson et al. 2009). This catalogue is the ideal counterpart to XCS because it is also based on automated pipeline analysis of the entire *XMM* archive. A 10-arcsec matching radius has been used to compare the samples. Fig. A2 shows the flux comparison from the individual cameras aboard *XMM*, using a 0.5–2.0 keV band. There is clear consistency between the two surveys, with no significant systematic offsets. It is important to note that the default XAPA fluxes for extended sources are not equally reliable. This happens for two reasons: first, the ECFs used to generate the fluxes relate to power-law spectra

---

[3] http://www.sdss.org

(whereas extended sources are more likely to have thermal spectra), and second, the fluxes have not been properly corrected for any source flux lying outside the XAPA -defined ellipse. In Section 4.3, we describe how aperture-corrected energy fluxes are determined for the candidates.

3.4 XAPA verification: extended sources

As mentioned above (Section 3.2.7), XAPA has catalogued to date in excess of 10 000 unique sources that have been statistically classified as extended. XAPA is not infallible however, and some of the objects in the candidates list will be erroneous – because they are blends of point sources or other artefacts of the data reduction – and a small fraction will be other types of genuinely extended X-ray sources (such as nearby galaxies or supernova remnants). Nevertheless, most of them will be clusters. In this section we first compare the XAPA -determined extents for the clusters in the XCS-DR1 sample to those in the 2XMM catalogue (Section 3.4.1). We then compare the candidate list to the cluster sample of the *XMM*-LSS survey in the same ObsIDs (Section 3.4.2). We then describe how we quantify the completeness level using simulations of our selection function (Section 3.4.3). We note that it is harder to quantify the contamination (due to blends and artefacts) level than the completeness level. In XCS we do not use simulations for this, but rather examine each source (and its optical counterpart) by eye (Mehrtens et al., in preparation).

3.4.1 Comparison with 2XMM

To investigate the quality of the XAPA -determined source extent, we have used the 394 XCS-DR1 clusters with matches to extended sources in the 2XMM catalogue. For these clusters, we have compared the XAPA major axis to the 2XMM extent measure. In the latter case, the quoted value is equivalent to the core radius of a β-profile (equation 2), so is always smaller than the XAPA value, typically by a factor of 5. The two measures are correlated (correlation coefficient of 0.514), with a $\simeq 30$ per cent scatter about the best-fitting relation. Therefore, despite the very different methods by which extents are measured by the two surveys, both descriptions are useful when determining source sizes. We note that only 49 (11 per cent) of the XCS-DR1 clusters in ObsIDs processed by 2XMM were not classified as extended sources in the 2XMM catalogue. Of course there might well be other clusters that 2XMM has detected as extended, but XAPA has not.

3.4.2 Comparison with the XMM-LSS

The *XMM* Large Scale Structure (*XMM*-LSS) Survey is reported in Pierre et al. (2006) and Pacaud et al. (2007). It covers a single contiguous region of roughly 6 deg^2, comprised of 51 ObsIDs, in which the authors have undertaken a dedicated cluster survey, accompanied by a detailed selection function. In this region they detected 33 'Class 1' extended objects. This class is designed to be uncontaminated by misclassified point sources. A more detailed examination of these objects (including optical overlays, photometry and spectroscopy) has confirmed 28 of these to be genuine clusters; the remaining five were shown to be nearby X-ray emitting galaxies. 29 of the 33 Class 1 objects have counterparts in XCS that were classified as extended by XAPA. This includes two of the non-cluster objects. Three of the remaining four Class 1 objects were detected by XAPA, but classified as point-like. The final

object (XLSS J022210.7−024048) was detected by XAPA, but subsequently removed from the source list because it did not meet our 4σ significance requirement.

The radius used in the matching of XAPA sources to the *XMM*-LSS was typically 10 arcsec. However, for XLSS J022433.8−041405 a radius of 24 arcsec was required to get a match; this source is large and elliptical, hence there is some uncertainty at the source centre, though the extent of the XCS source and that of its *XMM*-LSS counterpart are overlapping.

The *XMM*-LSS also have a C2 class of clusters with slightly less conservative selection criteria. This sample has yet to be published, but the authors report this class to contain ∼60 sources. Within the *XMM*-LSS ObsIDs, XCS detects 82 extended objects without flags (Section 3.2.5), so the overlap is likely to be substantial.

3.4.3 Selection function: method

Pioneering work by Adami et al. (2000), and later by Burenin et al. (2007), demonstrated the impact of complex selection effects on cluster samples derived from X-ray surveys. Pacaud et al. (2007) have shown, using the *XMM*-LSS Class 1 sample described above, that the measured evolution in the normalization of the $L_X − T_X$ relation is significantly affected by selection biases. In another X-ray study, Mantz et al. (2010) provide an in-depth discussion of Malmquist and Eddington biases and their effect on measurements of scaling relations. Optical and SZ cluster surveys are also increasingly supported by selection function simulations (Melin, Bartlett & Delabrouille 2005; Koester et al. 2007).

The ability to measure selection functions for XCS was embedded at the outset in XAPA. Indeed, one of the main reasons that drove us to design our own source detection pipeline, rather than using the excellent data products available from the *XMM* Survey Science Centre (Watson et al. 2009), was the requirement that we needed to be able to quantify the extended source selection function using synthetic clusters. In the following, we describe how the selection functions are carried out and present some results.

Our approach follows a general method in which synthetic cluster profiles are added to EPIC merged images, which are then run through XAPA. The angular size of the synthetic cluster profile is determined from the angular diameter distance at the chosen input redshift. The profile is then randomly positioned into a blank *XMM* 'image', with a uniform probability across the field of view, and then convolved with the appropriate PSF model. For this purpose we use the two-dimensional MAM (Section 2.1.1). This is a natural choice of PSF model for the selection functions because it accounts for the azimuthal variation in the PSF, and also because the alternative model (EAM, Section 2.1.1) is implemented in XAPA for source classification (Section 3.2.3): to keep the simulations fair, we cannot use the same model for blurring as we do for extent classification. The convolution with the PSF creates a probability density function (PDF) for the synthetic cluster profile. We note that the shape of the synthetic cluster profile depends on the user's specific requirements, and we will discuss some examples in Sections 3.4.4 and 3.4.5.

Next, an ObsID is chosen for the synthetic cluster to be placed in. The choice of ObsID will depend on the particular test being undertaken. For example, one might want to know the detection sensitivity in a particular ObsID, or one might want to know the detection sensitivity for a set of ObsIDs, e.g. those with similar n_H or exposure times. The synthetic cluster is added to the chosen ObsID as follows.

(i) The absorbed count-rate of the cluster profile is determined from the gridded LCFs (Section 2.5) for that ObsID, so that it matches the synthetic cluster's luminosity, temperature and redshift.

(ii) The cluster PDF is normalized to the LCF-predicted count-rate, thus creating a count-rate image.

(iii) The count-rate image is converted into a count-image by multiplying by the appropriate exposure map.

(iv) The synthetic count-images for the individual cameras are then added to the respective real images (Section 2.4).

(v) The individual images are then added to make a merged image.

The resulting merged image, containing the synthetic cluster, is then processed by XAPA in the standard way. There are two criteria that must be met in order for an input synthetic cluster to be deemed successfully 'recovered' by XAPA: the detection software must identify a source at the synthetic cluster location, and that source must be classified as extended. This has to be a new source; if the synthetic cluster happens to have been placed at random close to a previously detected real extended source, then the synthetic cluster is not classed as having being recovered (even if its 'counts' dominate those from the real source). Depending on the application, we might further require that the new detection not be flagged (see Section 3.2.5). It is not sufficient to perform the synthetic cluster recovery test only once, rather one must perform it multiple times to ensure an accurate measurement of the recovery probability for a given set of input parameters. There is so much parameter space to be tested (see below) that the number of selection function tests can run into the millions for certain applications. Determining the survey selection function is by far the most computationally demanding part of XCS.

3.4.4 Selection function: results (analytical models)

The simplest profile type that we have studied is that of an isothermal β-profile cluster (equation 2). Using this profile we have tested the selection function dependency on cluster parameters (e.g. redshift, temperature, luminosity, core radius, profile slope and ellipticity); on image parameters (e.g. exposure time, off-axis angle, azimuthal angle); and on cosmological parameters (e.g. k and Ω_m, the curvature and present mean mass density of the Universe, respectively). Some results from the β-profile selection function runs have already been published (Sahlén et al. 2009). In Figs A3, A4 and 14 we show some additional results.

Figs A3 and A4 show how the selection function depends on cluster luminosity and redshift. We show results for 3 and 6 keV clusters with a range of luminosities. In all cases, the input profiles were spherically symmetric, with $r_c = 160$ kpc and $\beta = 2/3$. The profiles were placed randomly in a subset of ObsIDs chosen to be a representative sample of the whole archive, i.e. to have a similar distribution of exposure times and Galactic latitudes. This simple test confirms that bright clusters can be consistently detected out to redshifts of at least 1, whilst fainter clusters can only be found with a reasonable certainty at lower redshifts. We note that typical L_X values at 3 and 6 keV are roughly 1 and 10×10^{44} erg s^{-1}, respectively, based on the low-redshift $L_X - T_X$ relation of Arnaud & Evrard (1999). Therefore, we can expect to detect roughly 60 and 85 per cent of 3- and 6-keV clusters, respectively, at $z = 0.6$, but only 10 and 75 per cent at $z = 0.9$ (assuming no evolution in the $L_X - T_X$ relation).

Fig. 14 shows how the selection function depends on cluster angular size. Here we have run a set of 3 keV clusters through
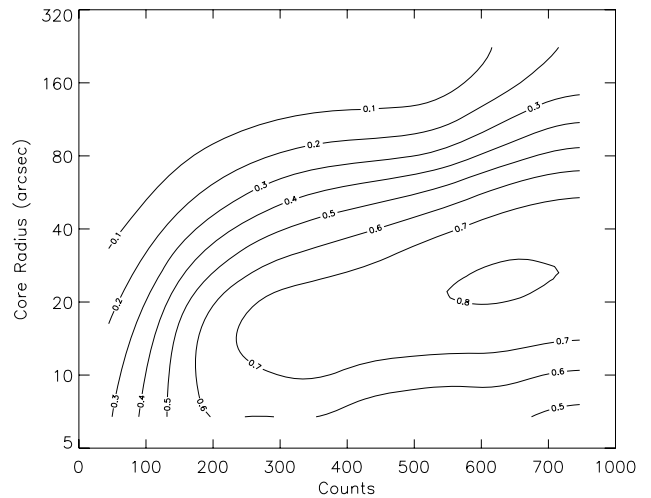

Figure 14. Predicted recovery efficiency for 3-keV clusters as a function of core radius and recovered counts. The synthetic clusters used for this test had circularly symmetric β-profiles ($\beta = 2/3$).

the selection function process with the physical core radii varying according to the findings of Jones & Forman (1984), i.e. in the range of 50–400 kpc. Over the range of redshifts probed by XCS ($0.1 \leq z \leq 1.0$), these core radii have an angular size in the range of 217–7 arcsec. Fig. 14 shows the fraction of clusters recovered by XAPA as a function of both angular size and the number of input synthetic cluster source counts. For clusters with more than 300 counts, the cluster recovery rate is good (≥ 70 per cent) when the extent is in the range $\simeq 10$–20 arcsec. These limits roughly translate to $0.1 < z < 0.6$ for $r_c = 50$ kpc and $z > 0.3$ for (more typically) $r_c = 160$ kpc.

3.4.5 Selection function: results (numerical models)

We have also investigated the recovery of clusters with profiles drawn from cosmological hydrodynamic simulations. For this purpose we have used clusters from the CLEF (CLuster Evolution and Formation in Supercomputer Simulations with Hydro-dynamics) simulation (Kay et al. 2007). The use of the CLEF clusters has enabled an investigation into the effects of more realistic (than β-profiles) cluster shapes on the XCS selection function, because CLEF includes clusters with cool cores and substructure.

The process by which the CLEF profiles are input into the *XMM* images is the same as for the analytical models. For simplicity, we use the CLEF catalogued mean cluster temperature, rather than the full temperature map, when calculating the total count-rates using the gridded LCFs (Section 2.5). These count-rates are then distributed using the emission measure profile (see Onuora, Kay & Thomas 2003 for details) as a probability map. The emission measure maps fully encode variations in temperature and density and so this approach will preserve any substructure in the surface brightness and also the presence of any central peak caused by a cool core.

The selection function work using CLEF has shown that strong peaks in the surface brightness profile, either due to substructure, or to a cool core, make clusters easier to detect than β-profiles. However, these clusters do not always make it into the 'recovered list', because they are either misclassified as point sources or flagged as being PSF-sized extended sources (Section 3.2.5). This misclassification trend is mitigated if the total number of detected counts is
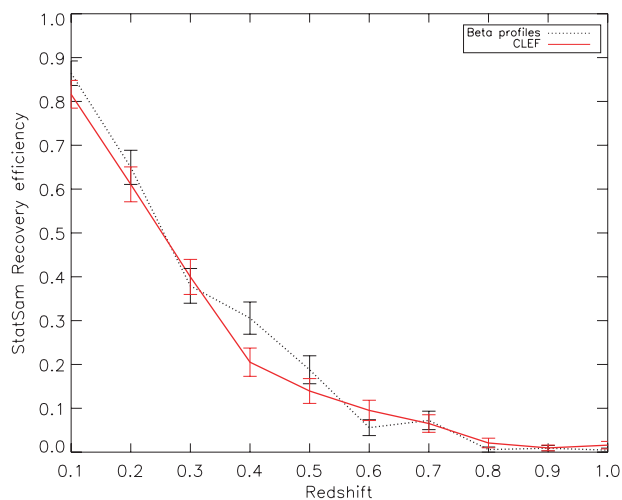
Figure 15. Predicted recovery efficiency of CLEF (Kay et al. 2007) and β-profile clusters as a function of redshift. A 500-count cut has been imposed, where the counts are as measured by the XAPA pipeline. The β-profile clusters are paired with a CLEF counterpart, in that they have the same redshift, temperature, luminosity and location in the respective ObsID.

large enough to sample more of the extended profile, and is almost completely resolved at the 500-counts-per-source level (Hosmer 2010).

The CLEF investigation has further shown that symmetrical β-profiles are an acceptable approximation to the XCS selection function. This is important because CLEF, and most other hydrodynamical samples, are only available for a single assumed underlying cosmology. In order to use XCS to measure the underlying cosmology, we need to know the selection function across a broad range of cosmological parameters. The suitability of the β-profiles was demonstrated by comparing the results of two duplicate selection function runs. The first used the CLEF cluster profiles, each input multiple times to determine recovery efficiencies. The second run replaced the CLEF clusters with isothermal β-profiles ($r_c = 160$ kpc, $\beta = 2/3$), whilst keeping all other aspects the same (ObsID, location, luminosity, temperature, redshift and input cosmology). The results, after a 500-count detection limit has been imposed, are shown in Fig. 15. Plotted in red is the average recovery efficiency obtained using the CLEF cluster sample, and overplotted are the data from using the β-profiles (dotted-black line).

4 ANALYSIS OF XCS CLUSTER CANDIDATES

In this section we describe a further set of XCS analysis pipelines. These re-examine the *XMM* observations of the candidates delivered by XAPA. The first pipeline examines each of the candidates in turn, jointly interrogating the respective ObsID and NED in the search for published redshifts (Section 4.1). The second pipeline carries out batch X-ray spectroscopy on candidates with redshift measurements and delivers measurements of the X-ray temperature (Section 4.2). The third outputs total luminosities, by fitting spatial profiles to those candidates with T_X measurements (Section 4.3). The fourth pipeline has been designed to estimate redshifts directly from the X-ray data (Section 4.4). The methodology of all of these pipelines is described below, together with a range of verification tests. Here, all quantities are calculated assuming a standard concordance cosmology ($H_0 = 70$ km s^{-1} Mpc^{-1}, $\Omega_m = 0.27$ and $\Omega_\Lambda = 0.73$) and all quoted luminosities are bolometric and within

radii where the overdensity is 500 relative to the critical density (R_{500}). See Fig. B7 for an overview of the process.

4.1 Automated NED queries: literature redshifts

Redshift measurements are essential if we are to use the candidates for science applications. However, with 3675 candidates in our current catalogue (2011 May 01), we need to find automated ways to derive as much redshift information as possible. In order to automatically identify redshifts that are already in the literature, we have constructed an algorithm that searches the NASA Extragalactic Data base (NED). These 'literature redshifts', or z_{lit}, are only available for a small fraction of the candidates, but they are still extremely important, in that they allow us to check our other redshift estimation techniques, in particular the X-ray redshifts described below (Section 4.4) and the red-sequence redshifts described in our companion catalogue paper (Mehrtens et al., in preparation).

The NED search was carried out for all the candidates. An initial search extracts all sources, classified as either a galaxy or a cluster, within a 30-arcmin search radius of the candidate centroid. Then, for every extracted object with a catalogued redshift, we calculate a crude placeholder luminosity, $L_{X,\text{ph}}$. The $L_{X,\text{ph}}$ is derived using the gridded LCFs (Section 2.5) and the soft-band XAPA count-rate. From the $L_{X,\text{ph}}$, we then estimate a corresponding placeholder temperature, $T_{X,\text{ph}}$, for the candidate using the L_X–T_X relation of Arnaud & Evrard (1999). From the $T_{X,\text{ph}}$ we can then estimate a placeholder $R_{500,\text{ph}}$ value (R_{500} is the radius from the cluster centre that represents an overdensity of 500 times the critical density), using the prescription in Arnaud, Pointecouteau & Pratt (2005), and a corresponding redshift appropriate angular search radius, $\theta_{500,\text{ph}}$. The velocity dispersion–temperature relation of Bird, Mushotzky & Metzler (1995) is used in a similar way to estimate a placeholder velocity dispersion $\sigma_{v,\text{ph}}$ for the candidate.

Any NED objects that lie outside their respective $\theta_{500,\text{ph}}$ are discounted as a true match. If any lie inside, then those classified in NED as clusters are then considered as a potential match. Should there be only one such object, then that is chosen as the best match. If there is more than one, then the object with the smallest positional offset is chosen. If no objects classified in NED as clusters fall inside the search radius, but some galaxies do, we then look for groupings of galaxy redshifts within the ($\theta_{500,\text{ph}}$, $\sigma_{v,\text{ph}}$) volume. If more than one grouping of galaxies is found, then the one with the smallest positional offset is chosen as the best match. When the query was run (2011 May 01), a total of 493 candidates were associated with published redshifts for clusters (412) or galaxy groupings (81) in NED.

The NED redshifts were then passed to the 'Redshift Followup (Archive)' stage of the XCS pipeline (Fig. 1) and individually checked. In doing so, it was discovered that some matches are wrong, i.e. the XCS source is not associated with the selected NED cluster. This is especially true at low NED redshifts (where the allowed matching radius is large). We found that imposing a redshift limit of $z = 0.08$ was effective in removing the erroneous matches, although this reduced the number of NED redshifts available to 218. Of these 218 candidates, 127 passed the quality control stage and made it into the XCS-DR1 sample. That is not to say that the remainder are not clusters, but rather that they cannot be confirmed as being so using the currently available optical and X-ray data (see Mehrtens et al., in preparation). Of these 219 only 121 are listed in XCS-DR1 with the automatically selected NED redshift. This is because, in the other cases, alternative redshifts were available. The alternative redshifts mostly came from either our own observations

or from our analysis of the SDSS archive, but in eight cases they came from the literature. In those eight cases, the NED redshift listed for the cluster had not been updated to reflect more recent optical follow-up. The tendency for NED to retain outdated redshift information first became apparent to us when we compared the NED redshift ($z = 1.2$) for the highest redshift XCS cluster (XMMXCS J2215.9−1738) to the value we published ($z = 1.46$) based on 31 secure spectroscopic redshifts (Hilton et al. 2010). The $z = 1.2$ value had been taken from (Olsen et al. 2008) and was based on single (*i*) band photometry.

4.2 Spectral fitting: X-ray temperatures

In this section we describe our pipeline to measure X-ray temperatures for candidates with a secure redshift measurement; the T_X-pipeline hereafter. In Section 4.2.1 we explain how spectra are extracted and corrected for background contamination. Next we describe how these spectra are fitted to X-ray models and how parameter uncertainties are calculated (Section 4.2.2). Both of these tasks are carried out in an automated fashion so, to assess their efficacy, we have carried out a series of tests. These tests are described in Section 4.2.3.

4.2.1 Generating the spectra

Spectra are generated for every candidate with an associated redshift measurement. The first step is to establish all the ObsIDs in which the candidate was observed and all the exposures within them were observed as well. We need to do that to ensure we have the maximum number of source counts available to carry out the fit. In the simplest case, the candidate will have only been observed in the ObsID listed in the MDL (Section 3.2.7), and there will be only three subexposures to exploit (one each for EPIC-mos1, EPIC-mos2 and EPIC-pn). However, in other cases the candidate might be covered by multiple ObsIDs (only the one generating the most soft counts is listed in the MDL). Moreover, there can also be multiple exposures within an ObsID, especially if the exposure time is long and had to be broken up over several satellite orbits. Finally, in some cases, one or more of the cameras might have been turned off, so fewer than three exposures are available.

When all the exposures have been gathered then the cleaned event lists, described earlier (Section 2.3.3), are used to generate spectra. Only photons in the 0.3–7.9-keV band are used for this (the telescope is poorly calibrated at softer energies and the spectra are background dominated at higher energies). The regions used to extract the source spectra are the ellipsoidal regions, ϵ_f, that XAPA defined for the respective candidate, although, if other XAPA sources overlap with any part of ϵ_f, then events from those pixels are not included when the spectra are produced. The redistribution matrices and area response files necessary for a spectral analysis are then created, using the *XMM* SAS package. These files are ObsID-, camera- and position-dependent and so one needs multiple sets for each candidate.

Every source spectrum generated needs an associated field spectrum for the purposes of background subtraction. The background subtraction in the T_X-pipeline was done using an in-field method, since XCS clusters do not generally have large angular sizes. The background spectra were usually taken from a circular annulus around the source, although in the case of sources very near the edge of the field of view, an ellipse perpendicular to the off-axis direction, with a circular region centred on the cluster excluded,

was used instead. The outer radius of the background annulus is 1.5 times the XAPA -defined major axis of the respective candidate. The inner edge varies depending on the exposure, but is no less than 1.05 times the major axis. Any pixels within the background region that overlap with other XAPA sources are excluded from the background spectrum. The normalization of the background is performed within XSPEC and reflects the ratio of the number of pixels in the source and background extraction regions.

4.2.2 Spectral fitting

The spectral fitting was carried out using XSPEC. The fitting was done using the maximum likelihood Cash statistic (Cash 1979). As mentioned above, there can be multiple spectra per candidate and these were usually all fitted simultaneously. The only exceptions were very low-count spectra, i.e. those with either less than 10 soft-band counts in total or those with less than 10 per cent of the soft-band counts of the spectrum with the most counts. These spectra were excluded from the simultaneous fit because it was found that they degraded the fits.

In XSPEC the photons within each spectrum are grouped into bins before fitting. For the T_X-pipeline we varied the minimum number of counts per bin according to the total number of counts in the spectrum. That way, higher signal-to-noise ratio spectra could be fitted to a higher spectral resolution (and vice versa). For spectra with fewer than 250 counts, the minimum was set at 1 count per bin. For spectra with more than 850 counts, the minimum was 5. In between those limits, the minimum was scaled between 1 and 5 counts using a power-law with an index of 0.75. This particular scaling of the minimum number of counts per bin was chosen after carrying out spectral simulations. It was designed to minimize the bias in the derived parameters while also minimizing the statistical uncertainties.

Four different models are fitted to the data. All of the models include a photoelectric absorption component (WABS; Morrison & McCammon 1983) to simulate the n_H absorption and a hot plasma component (MEKAL; Mewe et al. 1986) to simulate the X-ray emission from the ICM. The different models are

(i) WABS*MEKAL with the hydrogen column, n_H, frozen at the Dickey & Lockman (1990) value and the metallicity, Z, frozen at the canonical, $0.3Z_\odot$, value;

(ii) WABS*MEKAL but with n_H and Z allowed to vary;

(iii) WABS*(MEKAL+POWERLAW), as (ii), but including an extra power-law component to simulate a potential contaminating point source;

(iv) WABS*(MEKAL+MEKAL), as (ii), but with two MEKAL components rather than one, in order to simulate the case where there is a significant cool core in the cluster.

The best-fitting model of these four is usually used to derive the luminosity and temperature of the cluster, but if the best-fitting model does not give sensible parameters, then the next best model will be selected, and so on. The accepted ranges are $0.3\,\mathrm{keV} < T_X < 17.0\,\mathrm{keV}$ and luminosity less than $5 \times 10^{46}\,\mathrm{erg\,s^{-1}}$. It is important to note, however, that these luminosity values are not aperture-corrected and only relate to the luminosity originating from the ϵ_f region defined by XAPA. In general, a cluster will be more extended than this ellipse, and so these aperture luminosities, $L_{X,ap}$, need to be corrected for missing flux using a spatial model. We describe in Section 4.3.1 how such models are fit to XCS candidates.

The 68 per cent uncertainty bounds on the best-fitting $L_{X,ap}$ and T_X values are provided as part of the standard XSPEC fitting process:

the parameter in question is stepped from its best-fitting value until the fit statistic increases by the amount required for the confidence region needed (at each step-point the other free parameters are refit). This stepping is done in both the positive and negative directions to obtain a confidence region.

4.2.3 T_X-pipeline validation

The spectral pipeline is fully automated and so it is important to check the reliability of the results it produces before using them for scientific studies (such as the measurement of cosmological parameters). We have performed these checks using both XSPEC simulations and actual data. The results of the first check are presented in Fig. 16, where temperature uncertainty is plotted against the number of counts in the fitted spectrum. For this test we have simulated cluster spectra (all at $z = 0.5$), using the MEKAL model, with a range of temperatures ($1.5 < T_X < 8$ keV). It can be seen that it is much easier to constrain the temperatures of cool systems (red) than it is for the hottest systems (blue). The constraints also become worse as the number of source counts decreases. It can be seen that below 300 counts the temperature uncertainties exceed 50 per cent for the 8-keV systems, though they are considerably smaller than that for lower temperature systems. Based on the results of this test we have made our decision as to what count limit we should impose on the candidate list when defining a sample for cosmological tests. We have set this limit at 300 because we use T_X values as a mass proxy when measuring cosmological parameters (Sahlén et al. 2009).

The results in Fig. 16 were based on model spectra (albeit with actual *XMM* background contamination) and so should be seen as a best-case scenario: real clusters do not have a perfectly isothermal ICM, nor have zero contamination from point sources. Therefore we have carried out a related test using four real clusters (Table C1), the results of which can be seen in Fig. 17. Here the best-fitting T_X value (and its 1σ uncertainty from XSPEC) are plotted against the total number of counts in the spectra. This was achieved by artificially reducing the exposure time of the respective event file. We note that only one realization of this procedure was performed for each total number of soft-band counts and that the error bars
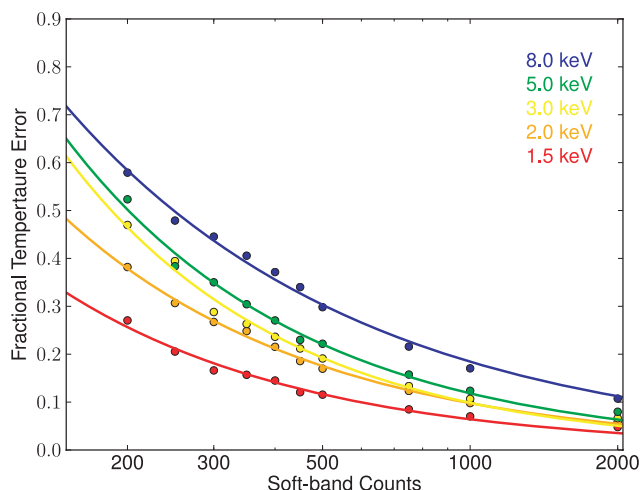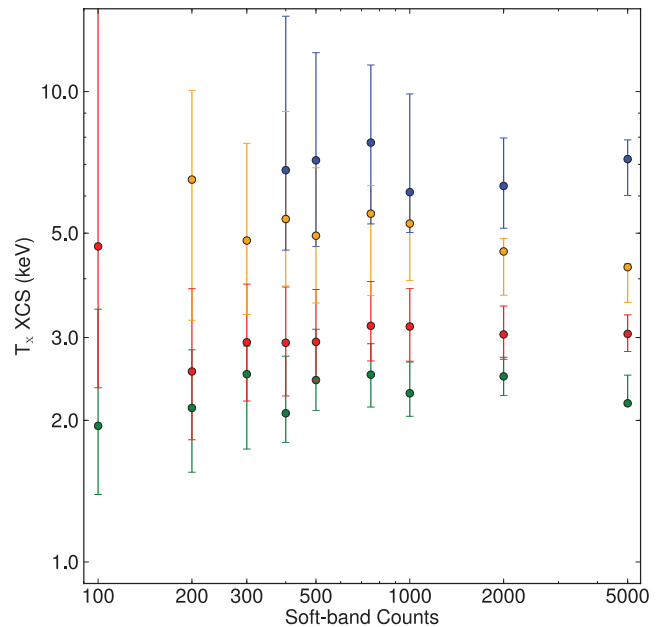

Figure 17. The XCS-determined X-ray temperatures (and uncertainty) as a function of the number of counts in the fitted spectrum. Each colour represents a cluster that was detected with more than 5000 counts. For details of the four clusters used in this plot, see Table C1. The respective exposures were then subdivided to generate lower count spectra. Note that the higher temperature systems do not yield fits at the low-count end. The 1σ error bars come from the XSPEC fitting software (see Section 4.2.2 for details).

are the standard XSPEC -generated values. From Fig. 17 it is clear that there are no systematic biases in the derived values of T_X as the number of counts decreases. The error bars increase, with decreasing counts, in line with the expectation from Fig. 16. The fit failed to converge at low counts for the hotter clusters, but in general Fig. 17 supports our decision to cut the candidate list at 300 counts for cosmological studies. This test also demonstrates that it is still worth fitting candidates with fewer counts, since we can derive reliable T_X values in the galaxy groups regime down to 100 counts.

We have carried out a test to see if the T_X-pipeline works at large off-axis angles, since candidates are located across the *XMM* field of view. There are not many clusters to choose from for this test, but we did identify eight systems that have been observed by *XMM* both on- and off-axis. For this purpose we define off-axis (on-axis) as a source centroid more than 10 (less than 3) arcmin from the ObsID aim point. The standard XCS spectral reduction was undertaken and the results can be seen in Fig. 18. It can be seen that the fits to spectra taken off-axis, while in general having larger uncertainties due to having a lower signal-to-noise ratio, are consistent with the corresponding on-axis results. We can therefore be confident that the pipeline produced XCS T_X values that are not biased in cases where the objects are located on the outskirts of the field of view.

The test in Fig. 18 shows that the XCS pipeline is internally consistent, but it is also important to compare XCS parameters to those derived externally, i.e. by other authors, since they will use different approaches. In particular, most cluster spectral fitting is done on an object-by-object basis, with the background regions and the light-curve cleaning being adjusted by hand. In contrast the XCS pipeline is completely automated because we do not have the resources to fit hundreds of candidates individually. We have

Figure 16. Fractional temperature uncertainty as a function of the number of soft-band counts as a result of fitting simulated $z = 0.5$ MEKAL spectra with different temperatures, going from cool to hot clusters. The red, orange, yellow, green and blue points represent spectra of 1.5, 2, 3, 5 and 8 keV, respectively.
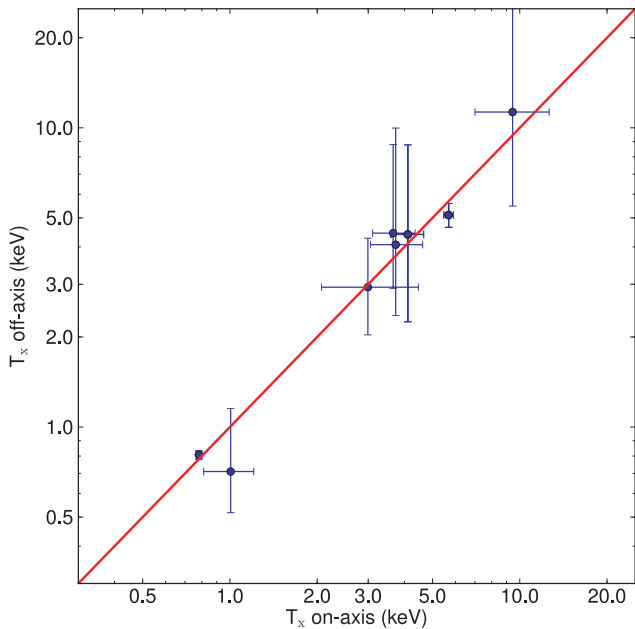
Figure 18. Comparison of XCS-determined X-ray temperatures when the cluster is observed off-axis (y-axis) or on-axis (x-axis). For details of the eight clusters used in this plot, see Table C2. The solid line shows the one-to-one relationship. The error bars are 1σ. Both x- and y-errors come from XSPEC (see Section 4.2.2 for details).

therefore tested the quality of the results from our pipeline using previously published results. We have constructed a sample of 11 XCS clusters which have previously published temperatures measured with *XMM* (Gastaldello et al. 2007; Pacaud et al. 2007; Hoeft et al. 2008). The results can be seen in Fig. 19, where the temperatures derived from the XCS pipeline are plotted against those
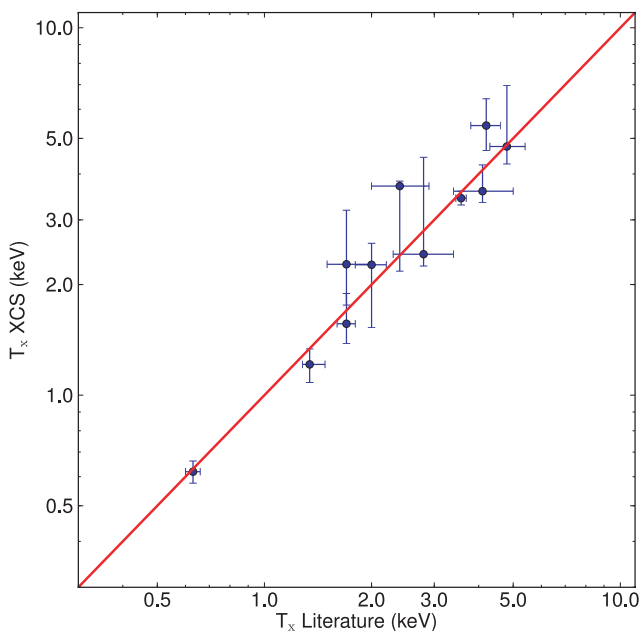
Figure 19. Comparison of XCS determined X-ray temperatures with values determined by other authors. For details of the 11 clusters used in this plot, see Table C3. The solid line shows the one-to-one relationship. The error bars are 1σ. The x-errors are as quoted in the literature. The y-errors come from XSPEC (Section 4.2.2 for details).

measured by other authors. It can be seen that there does not appear to be any systematic offset and the XCS temperatures are consistent with the literature values. This final test demonstrates that the XCS T_X values are reliable and hence suitable for science applications without the need for a further 'hands on' analysis stage.

4.3 Spatial fitting: X-ray luminosities

As mentioned above (Section 4.2.2), the spectral pipeline produces both luminosity and temperature fits, but the luminosities, $L_{X,ap}$, are within an aperture and are not corrected for missing flux. In order to extrapolate the cluster emission to R_{500}, so that the total cluster luminosity can be calculated, it is necessary to measure the surface brightness profile. This is achieved in the XCS spatial pipeline, the L_X-pipeline hereafter, by fitting an analytical function to the cluster and then using this to extrapolate to R_{500}. We decided against using the alternative, non-parametric approach that produces deprojected gas densities, e.g. Croston et al. (2008), because it is complex and, importantly, does not allow us to extrapolate fluxes to R_{500}.

4.3.1 Spatial models

Surface brightness fits are performed for every candidate that passes the spectroscopic pipeline. The main function used to characterize the shape of the clusters was a simple one-dimensional, spherically symmetric, β-profile model (Cavaliere & Fusco-Femiano 1976):

$$S(r) = S(0) \left[1 + \left(\frac{r}{r_c} \right)^2 \right]^{-3\beta+1/2}, \qquad (2)$$

where r_c is the core radius and β is the density index parameter, which encodes the power-law decline. Three different types of β-profile models were fitted to the data as follows.

(i) One with β frozen at the canonical value of 2/3. The free parameters are the normalization and the RA and Dec. of the centroid.

(ii) One as (i), but with β allowed to vary.

(iii) One with an inner power-law cusp inside a certain parameterized radius (usually of the order of the core radius). This gives us a crude description of clusters with cool cores or AGN contamination. The free parameters are as (ii), plus both an extra normalization and an extra power law index.

The same background regions were used for the surface brightness fitting as were used for the spectral fitting. However, in addition to knowing the total number of counts in the background region, it is also necessary to know how those counts are distributed. The total *XMM* background varies considerably across the field of view and so for extended sources, such as clusters, one cannot assume that the background counts are divided equally between all the pixels. The background can be considered as having two components, an 'X-ray component' that is focused (and so vignetted) by the telescope mirrors and a 'particle component' that is not. In reality, these terms are not particularly accurate, since the X-ray component includes soft protons that are focused by the mirrors and the particle component includes high-energy photons that are created as a result of particle collisions with the telescope structure.

The X-ray and particle components need to be treated separately during spatial fitting because their spatial variation is different: the

X-ray component is assumed to vary in the same way as the exposure map, because it is vignetted, whereas the particle component should show no positional dependence. To determine the particle background count-rate per pixel, we have selected, from the respective ObsID, two or more source regions that are at significantly different off-axis angles. We then compare the ratio of the normalized counts within those regions to the ratio of the same regions in the exposure map. The difference between those ratios tells us at what level the counts are contaminated by the particle background. This process is illustrated in Fig. B8.

4.3.2 Spatial fitting

As with the spectroscopic pipeline, the fit is performed simultaneously on all ObsIDs, and subexposures, in which the candidate was observed (barring those with very few counts); see Section 4.2.1. For the spatial fitting we generated new, 4.35 arcsec pixel, image files in the same 0.3–7.9 keV band as was used for the T_X-pipeline. The three spatial models (see above) were convolved with the 1D EAM (Section 2.1.1) *XMM* PSF model before the fitting took place. They were then multiplied by the exposure map at the respective location, in order to add observational effects such as vignetting and chip gaps. The background was accounted for as described above.

The maximum-likelihood Cash statistic was used for the comparison between the model and the data. The MINUIT package (James & Roos 1975) of minimization algorithms was used to find the best-fitting of the three models. The best-fitting model was then used to calculate the scaling of the luminosity from the spectral extraction region to R_{500}. This was achieved by calculating the ratio of the summed emission from the spectral extraction region (i.e. ϵ_f) to that from a circular region, radius R_{500}. The ratio was then used to scale $L_{X,ap}$ to $L_{X,500}^{fit}$. This luminosity scaling value is typically in the range of 0.9–3.0, depending on the complex interplay between the cluster size and redshift, the location on the field of view, and the depth of the exposure.

The 1σ uncertainty bounds on the free parameters in the spatial fits were generated in a similar way to that used in the T_X-pipeline (Section 4.2.2), i.e. by stepping, and fixing, the parameter of interest. This was done separately for each of the three models used in the spatial fitting. The uncertainty bounds on the $L_{X,500}^{fit}$ were not so straightforward to calculate, however. This is because the $L_{X,500}^{fit}$ calculation involves both the T_X-pipeline and the L_X-pipeline and, since they are performed separately, there is no information on the correlations between them. Ideally one would carry out a simultaneous fitting of a spectral and spatial model to a data cube (X,Y and energy), as was demonstrated by Lloyd-Davies, Ponman & Cannon (2000), but this process would be too complex and CPU-intensive for the batch processing required by XCS. Therefore, we adopt the conservative approach of taking the uncertainty bounds for the two quantities (i.e. on the luminosity scaling value and on $L_{X,ap}$) and calculating the luminosities for the four most extreme combinations (Fig. 20). The highest and lowest luminosities are then used as the uncertainty bounds of the $L_{X,500}^{fit}$ measurement, although this will almost certainly be an overestimate.

We note that not all of the candidates that are passed to the spatial fitting generate acceptable fits. When the spatial fitting fails for a candidate, we then estimate the luminosity, $L_{X,500}^{est}$, by extrapolating the $L_{X,ap}$ value assuming a standard cluster profile. For this, we fix the power-law slope to be $\beta = 2/3$ and use a core radius appropriate to the T_X. Thus all candidates that pass the spectroscopic
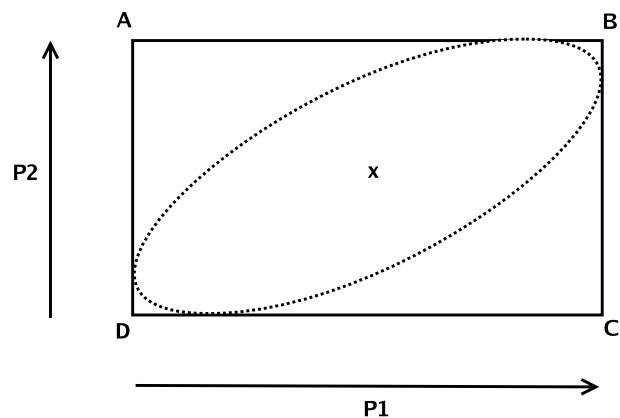

Figure 20. Schematic of how the luminosity uncertainties are calculated by combining the uncertainties on the two quantities (P1, P2). The cross represents the best-fitting point and the dotted line represents the (unknown) 1σ confidence contour. The model luminosity is evaluated at A, B, C and D, and the maximum and minimum values are used as upper and lower uncertainty bounds.

pipeline will have either an associated $L_{X,500}^{fit}$ measurement or $L_{X,500}^{est}$ estimate.

4.3.3 L_X-pipeline validation

The L_X-pipeline is fully automated and so it is important to check the reliability of the results it produces before using them for scientific studies (such as the study of the evolution of the L_X–T_X relation). We have done this by comparing the XCS-derived results for β and $L_{X,500}^{fit}$ with those derived by other authors. First we have examined clusters in common with the sample of Alshino et al. (2010). This should be a fair comparison since this sample is a subset of the *XMM*-LSS (Pacaud et al. 2007) and is therefore, like XCS, drawn from an *XMM* survey. Fig. 21 shows XCS β for four clusters plotted against the β values taken from Alshino et al. (2010). It can be seen that there is no systematic bias between the values. In addition, the scatter about the one-to-one relation (solid line) is consistent with the measurement uncertainties. We can therefore be reassured that the spatial parameters we obtain from the spatial pipeline are reliable.

We have also compared the XCS $L_{X,500}^{fit}$ values against published values obtained from clusters with 300 or more counts in the *XMM*-LSS sample (Pacaud et al. 2007, note that these are 300 *XMM*-LSS, not XAPA, counts). The XCS $L_{X,500}^{fit}$ values are plotted against those of Pacaud et al. (2007) in Fig. 22. It can be seen that they closely follow the one-to-one relation (solid line). This test demonstrates that the XCS $L_{X,500}^{fit}$ values are reliable and hence suitable for science applications without the need for a further hands-on analysis stage.

4.4 Spectral fitting: X-ray redshifts

We cannot exploit the thousands of candidates that XAPA has produced without first determining their redshifts. As mentioned above (Section 4.1), only a small fraction have redshifts available from the literature, so we have carried out both an intensive optical follow-up campaign, and exploited the SDSS archive, to gather more redshifts. This effort has yielded redshift information for $\simeq900$ additional candidates to date (Mehrtens et al., in preparation), but the redshift follow-up of the XCS is still far from complete. We therefore
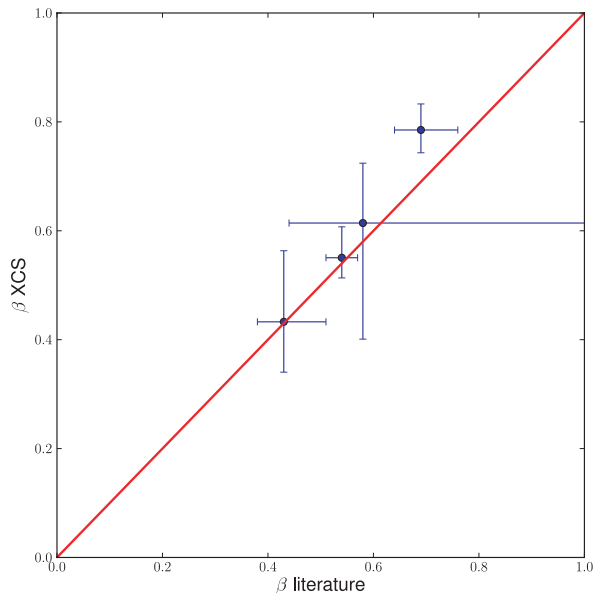
Figure 21. Comparison of the XCS-determined outer slope (or β) with that derived by other authors. In both cases, the surface brightness of the clusters was fit using a circularly symmetric King profile (without a central cusp, see Equation 2). For details of the four clusters used in this plot, see Table C4. The solid line shows the one-to-one relationship. The error bars are 1σ. The x-errors are as quoted in the literature. The y-errors come from the XCS fitting software (see Section 4.3.2 for details).
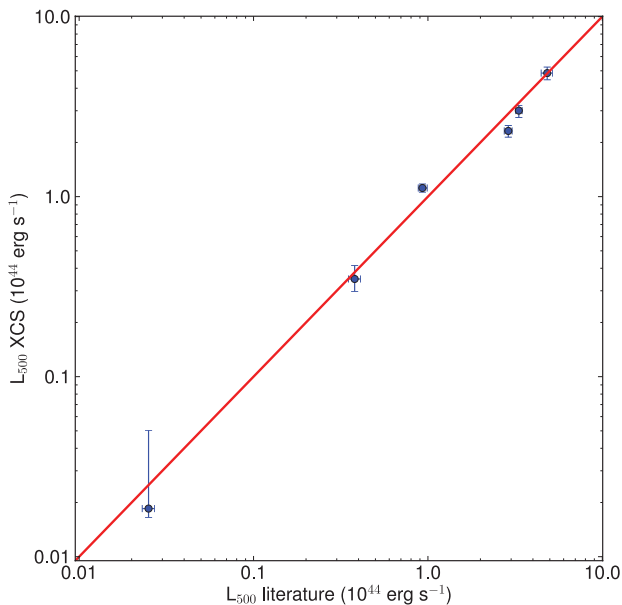
Figure 22. Comparison of XCS-determined bolometric luminosities within R_{500} with values determined by other authors. For details of the six clusters used in this plot, see Table C5. The solid line shows the one-to-one relationship. The error bars are 1σ. The x-errors are as quoted in the literature. The y-errors come from the XCS fitting software (see Section 4.3.2 for details).

decided to use the *XMM* data itself to constrain candidate redshifts. This process, of measuring 'X-ray redshifts' or z_X, has been demonstrated by several authors for individual clusters (Hashimoto et al. 2004; Werner et al. 2007; Lamer et al. 2008; Rosati et al. 2009) and recently on a sample of Chandra clusters by Yu et al. (2011), and has even been used to study bulk motions of the gas within the

bright, nearby clusters (Dupke & Bregman 2001b,a), but has never been used on the industrial scale we need for XCS. In the following we describe the X-ray redshift pipeline, z_X-pipeline hereafter, and its verification using XCS clusters with known redshifts.

4.4.1 Generating and fitting the spectra

Similar to the T_X-pipeline (Section 4.2), all exposures that overlapped with a particular candidate were used and a simultaneous fit was carried across all the respective spectra. Because this pipeline will be run on the many thousands of candidates that XAPA produces, we needed to keep the processing time per candidate to a minimum. We therefore chose a single-temperature MEKAL model, convolved with a photoelectric absorption model. Moreover, during the fitting, only the spectral normalization was left free. By design, we do not want to assume the redshift, so we ran a series of fits stepping from $z = 0.01$ to 2, in steps of 0.01. At all of these steps, the metallicity was fixed at $Z = 0.3 \times$ solar and n_H at the Dickey & Lockman (1990) value. The T_X was not free either, but rather calculated (via the Arnaud & Evrard 1999 L_X–T_X relation) from the best-fitting normalization at that redshift step (assuming no scatter in the L_X–T_X relation).

At each redshift step, the Cash statistic was recorded, as demonstrated in Fig. A5. The z_X for the candidates is then chosen by searching for minima in the distribution of Cash statistic values. Usually the redshift corresponding to the lowest Cash statistic was used, but if the corresponding temperature was $T_X > 8\,\text{keV}$ then the next deepest minimum was chosen, and so on. This limit was placed on the allowed temperature because very few $T_X > 8\,\text{keV}$ are expected to be detected by XCS (Sahlén et al. 2009). The 1σ uncertainty on z_X, σ_{z_X}, was also determined from the Cash statistic distribution. We note that in the following we refer to statistical uncertainties expressed as a percentage, and by this we mean $100 \times \sigma_{z_X}/z_X$.

4.4.2 z_X-pipeline validation

X-ray redshift measurements have not been attempted with this level of automation before. We checked our results to see if our z_X values are suitable for science applications. We did this using clusters in XCS-DR1 that had optically determined redshifts. Under the assumption that the optical redshift is correct, we have compared the z_X to the z_{true} for 42 XCS-DR1 clusters where the z_X fit yields a statistical uncertainty of <20 per cent and $\sigma_{z_X} < 0.05$ (Fig. 23). We chose $\sigma_{z_X} = 0.05$ as the upper limit for this comparison because this is the typical error on the single colour photometric redshifts presented in Mehrtens et al. (in preparation). As shown in Fig. 23, the z_X fits are usually good (to within the errors), but the level of catastrophic failures, i.e. where $z_{\text{true}} - z_X = \Delta_z \gg \sigma_{z_X}$, is high. The failure rate for X-ray redshifts is 24 per cent, compared to ≤ 7 per cent for the photometric redshifts presented in Mehrtens et al. (in preparation).

These catastrophic failures are not unexpected, even for high signal-to-noise ratio spectra, when the single-temperature spectral model is too simple, e.g. if there is AGN contamination or a cool core. Similarly, if the n_H and/or abundance was fixed at the wrong value, or the cluster is not close enough to the Arnaud & Evrard (1999) L_X–T_X relation, then one might obtain z_X values with small errors, but that are not physically realistic. We have investigated the possibility of making cuts on the sample to objectively weed out the catastrophic failures. However, it does not seem to be possible to
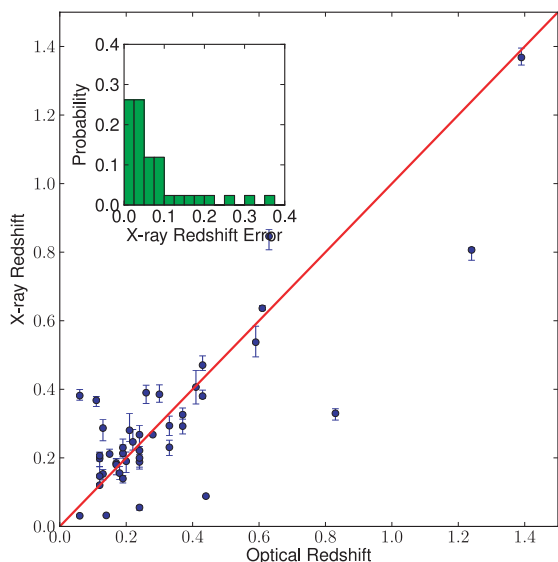
Figure 23. Measured X-ray redshifts plotted against optically determined redshifts for clusters in XCS-DR1 (Mehrtens et al., in preparation). The solid line shows the one-to-one relationship. Only X-ray redshifts with statistical uncertainties of 20 per cent or less are shown. The insert shows a histogram of the difference between the X-ray redshifts and optically determined redshifts.

predict, a priori, that a given z_X estimate would be unreliable from any combination of number of counts, cluster temperature or n_H. We discuss in Section 5 how we plan to use the z_X values, despite their high failure rate.

5 DISCUSSION

In Fig. 1 we introduced the complex methodology associated with the generation of a cluster catalogue based on serendipitous detections in the *XMM* archive. We went on to describe, and verify, all the steps in the methodology that involve *XMM* data [other steps are described in our companion paper, Mehrtens et al. (in preparation)]. In this section we will discuss each step again, making reference where appropriate to predictions made in our pre-launch paper Romer et al. (2001, R01 hereafter), and in our cosmology-forecasting paper Sahlén et al. (2009). We also highlight areas of improvement.

In Section 2.2 we described the download of data from the archive and showed how the area covered by the archive has grown over the last 10 years. By now there are over 600 deg^2 of the sky covered by *XMM*, and 51, 276 and 410 deg^2 (at >40, >10 and >0 ks depths, respectively) are in regions suitable for cluster searching. We note that the exposure times used in Fig. 2 are after flare correction (Section 2.3.3), and that flares typically affect one-third of the exposure. The rate of addition of new area is slowing over time, reflecting the trend towards repeated observations and fewer, but longer, exposures. It is, therefore, almost certain that XCS will not reach the 800-deg^2 target set in R01. The revised target of 500 deg^2 target set in Sahlén et al. (2009) does remain achievable though (as long as no minimum exposure time cut is applied).

The distribution and average (requested) exposure times of ObsIDs in the public archive is close to what was anticipated in R01, but due to the unanticipated need for flare correction, the average usable exposure is only 13 ks (compared to a requested average of 20 ks and a predicted, in R01, average of 22 ks). These decreases,

in exposure time and areal coverage, will certainly impact the size of the final XCS cluster catalogue. However, we have been able to use a lower minimum acceptable source significance (4σ rather than the 8σ used in R01) – because the XAPA extent determination is more effective at low signal-to-noise ratio than expected of previous experience with ROSAT – and this will help one to keep the cluster numbers up.

In Sections 2.3 and 2.4 we described the reduction of the downloaded data, including mitigation of time periods affected by background flares, and the production of images. This was done in a fairly standard way, albeit on a much larger, and more automated, scale than is typical. In R01 we expected that XCS source detection would be carried out only in EPIC-pn images, because we assumed that it would be too complicated to carry out selection function tests on merged images. However, in practice we have been able to run source searching and selection functions on merged images without any difficulty (Section 3.4.3). This has helped compensate for the decreased sensitivity, and increased background levels, of the EPIC-pn CCDs compared to pre-launch predictions. One thing that was not anticipated in R01 was the need to create mask files by hand for about a third of the ObsIDs (Section 2.4.1). This tedious process has been carried out by student volunteers and has not actually held up the processing of the archive significantly.

Overall, we are satisfied with the performance of the procedures described in Section 2 and do not plan any major modifications in future. That said, we did uncover during the quality control stage that some of the masks were too small and also that a small fraction of the reduced image had an atypically high background (see Mehrtens et al., in preparation). These two factors have resulted in contamination of the candidate list at the $\simeq 7$ per cent level. To avoid such contamination in future, we have improved the way that eyeball checks of reduced images will be carried out.

In Section 3 we described the generation of the XCS source catalogue using the XCS Automated Pipeline Algorithm (XAPA), and the tests we have carried out to demonstrate its efficacy. In Section 3.1, we explained how XAPA applies multiscale wavelets to generate a source list per ObsID, and discussed some of the successes of XAPA, including the ability to detect sources over a wide range of sizes and signal-to-noise ratios: only very rarely does one look at an image of an ObsID, with XAPA ellipses overlaid, and see the real sources that have been missed or artefacts (e.g. chip edges, where there can be discontinuities in the background level) misidentified as sources. We are especially pleased with the XAPA vision model (Section 3.1.2) because of its ability to detect sources within sources and to fit source ellipses. During the development of XAPA, it was found that vagaries of the *XMM* optics could result in false source detections (e.g. when point sources had extended lobes, due to the complex off-axis PSF), or incorrect size measurements (e.g. when an extended source had a cuspy core). These issues were addressed with additional subalgorithms.

As shown in Section 3.3, the parametrization of point sources (fluxes, positions etc.) is very good. Not only does this give us confidence that the extended source centroids are suitable for cluster searching, it also demonstrates that the point source catalogue itself can be used for science applications. XCS members and collaborators are using the data products in the point source catalogue in a variety of ways, including a study of the evolution of quasar X-ray spectra and a search for X-ray Dim Isolated Neutron Stars (or XDINS, see Haberl 2007 for a review of XDINS).

We do have some concerns, however, about the parametrization of extent by XAPA (Section 3.2.3) because the available PSF models are known to have deficiencies, especially off-axis. Occasionally,

during the quality control stage (Mehrtens et al., in preparation), we see sources that are obviously (from the X-ray data themselves and/or from the related optical image) point-like but that have been classified as extended and erroneously entered into the candidate list. Likewise there are likely to be incidences of extended sources that are detected but falsely classified as point-like (or flagged as PSF-like). The latter effect was indicated by the selection function test using numerically generated synthetic clusters (Section 3.4.5; see below). For these reasons we plan to adapt XAPA to use a new 2D PSF model that is currently under development by Read et al. (2010). The implementation of this new 2D PSF model would be a major undertaking because it would require the extent determination subalgorithm of XAPA to be rewritten and also necessitate the recalculation of the survey selection functions. It is also worth pointing out that even a perfect PSF model cannot prevent very nearby (on the sky) sources from becoming blended into a single source, especially when the signal-to-noise ratio is low. These blends will always affect our candidate list at some level (as they will any cluster searching project based on *XMM* detections); some will be obvious from the optical follow-up [see Mehrtens et al. (in preparation) for examples], but some might well require higher resolution imaging, e.g. from Chandra, to be identified.

The collation of a Master Detection List (MDL) for the survey (Section 3.2.7) has been fairly straightforward, despite the fact that so many ObsIDs overlap (and so many sources are detected multiple times in the archive). However, we have found that the process by which duplicate extended sources are identified (via a fixed matching radius of 30 arcsec) does not always work at low redshifts ($z \lesssim 0.2$), so we are in the process of improving this aspect of XAPA. As of 2010 May 1, the MDL contained 114 711 point sources and 12 582 extended sources, although, as just noted, a small number of the extended sources will be duplicate entries. We have selected 3675 of the extended sources as cluster candidates, after making a series of cuts to the extended source list, and these have then been passed on to the post processing and optical follow-up steps described in Section 4 and in Mehrtens et al. (in preparation).

We stress that the MDL, and hence the candidate list, was derived from the analysis of individual ObsIDs, even in regions where different ObsIDs overlap. In fact, approximately 40 per cent of ObsIDs in the XMM public archive have significant overlap with other ObsIDs, with a median additional exposure time of 70 per cent. Therefore, it would be possible to increase the number of sources, and hence candidates, detected by XAPA using co-adding ObsIDs. However, this would require a major overhaul of both XAPA and the selection function methodology (in the latter case because the PSF would be significantly more complicated), and we have no plans to use co-added ObsIDs in XCS. That said, we do take advantage of multiple exposure when running the T_X- and L_X-pipelines.

We note that the XCS is not the largest compilation of *XMM* detections; the *XMM*–Newton Serendipitous Survey 2XMM catalogue (Watson et al. 2009) contains 191 870 sources discovered in 3491 *XMM* ObsIDs. We made comparisons between XCS and 2XMM point and extended sources in Sections 3.3 and 3.4.1, finding them to be in good agreement. We compared XCS to another sample of *XMM*-selected clusters (the *XMM*-LSS, Section 3.4.2), and found them too to be in a good agreement: only four of the 33 *XMM*-LSS 'class 1' extended sources did not make it into the candidate list (because they did not meet the extent and/or signal-to-noise ratio criteria).

Selection functions are very important to any cluster survey that plans to carry out statistical studies, such as the measurement of

scaling relations or cosmological parameters. The XCS selection functions need to describe survey completeness as a function of a wide number of parameters, and are thus very CPU intensive. Examples of our selection function work so far are given in Sections 3.4.4 and 3.4.5. We have demonstrated, using simple analytical models for the ICM distribution, that XCS can detect typical (for the local L_X–T_X relation), 3- and 6-keV clusters to high redshifts, but that the percentage recovery of the cooler (i.e. fainter) clusters drops off rapidly, e.g. from roughly 60 per cent at $z = 0.6$–10 per cent at $z = 0.9$ for 3-keV clusters (Fig. A3). These predictions of the selection function redshift dependence were based on the assumption that all clusters have core radii of $r_c = 160$ kpc, so we have also investigated our sensitivity to smaller and larger clusters (Fig. 14). We found that for clusters with more than 300 counts, the cluster recovery rate is good (≥ 70 per cent) when the extent is in the range of $\simeq 10$–20 arcsec. These limits roughly translate to $0.1 < z < 0.6$ for $r_c = 50$ kpc and $z > 0.3$ for, more typical, $r_c = 160$ kpc. XCS is not as sensitive to clusters with core radii at the top end of the Jones & Forman (1984) range; roughly only 20 per cent (40 per cent) of 300-count clusters with $r_c = 400$ kpc are recovered at $z > 0.3$ ($z > 1$), although this rises to 60 per cent (75 per cent) for 1000-count clusters. This insensitivity was not anticipated in R01 (i.e. before we had access to realistic selection functions); we claimed therein that all clusters with core radii larger than 20 arcsec would be flagged as extended sources. It may seem counterintuitive that clusters of larger angular extent are harder to recover, but this is due to two factors; first, more extended clusters have lower surface brightnesses and correspondingly lower contrast against the background, making it harder to detect them, and second, our wavelet scales were chosen with more compact clusters (<250 kpc, R01) in mind.

In Section 3.4.5 we have used numerically generated 'clusters', from the CLEF hydrodynamical simulation of Kay et al. (2007), to investigate whether factors such as cool cores (which result in luminosity enhancements at small radii), or recent merging activity, might impact the ability of XCS to detect clusters. We found that the numerical 'clusters' are easier to detect than the analytical β-profile 'clusters', but that they are more likely to be misclassified (as point-like) when they contain cool cores. This effect is reduced as the number of 'source' counts increases, and above 500 counts no longer occurs (Fig. 15). This test justifies the use of selection functions based on simple analytical cluster profiles in the XCS cosmology forecasting paper (Sahlén et al. 2009, which was based on a minimum count threshold of 500). This is important because CLEF, and most other hydrodynamical samples, are only available for a single assumed underlying cosmology and, in order to use XCS to measure the underlying cosmology, we need to know the selection function across a range of cosmological parameters. However, this test further suggests that it may not be appropriate to use only simple analytical profiles when establishing selection functions for a minimum count threshold of 300 (which we have determined to be the limit to which we can expect to recover reliable T_X measurements; see below), and this is something we plan to investigate further.

In Section 4.1, we described an automated search for redshifts available in the literature. When this search was run in 2011 May, a total of 493 candidates were associated with published redshifts using NED. We have found, however, during the preparation of the first XCS data release [XCS-DR1, Mehrtens et al. (in preparation)] that the NED redshifts are not always appropriate for their respective candidate: the match to a given NED cluster might be erroneous, especially at low redshifts (where the allowed matching
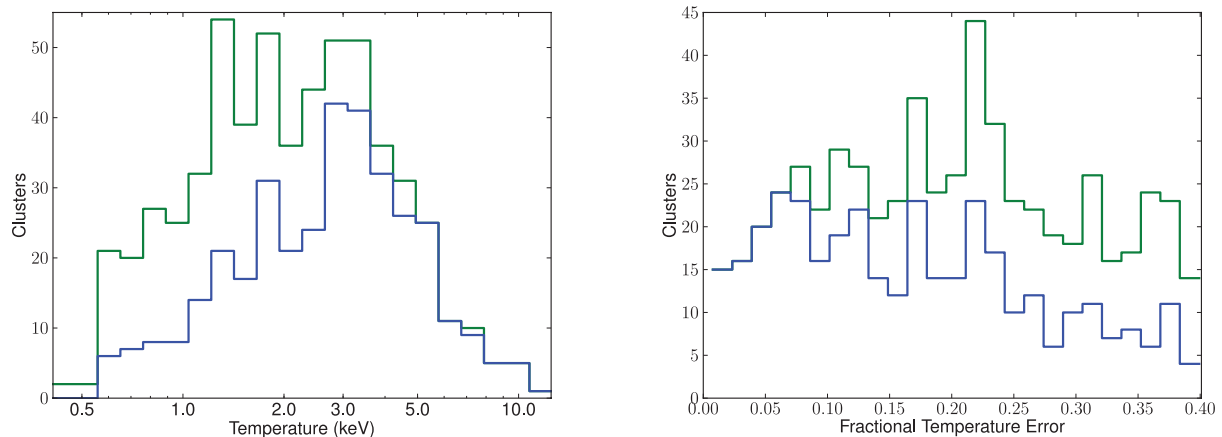
Figure 24. Number of clusters with less than 40 per cent temperature errors and with more than 300 soft-band counts (blue) and with no count cut (green) plotted against measured temperature (left panel) and fractional temperature error (right panel).

radius is large). We have therefore imposed a default[4] minimum redshift limit of $z = 0.08$ when using literature redshifts. Moreover, even if the match to the NED cluster is correct, the default NED redshift for that cluster might not be the best one available in the literature. Therefore, of the 493 redshifts automatically selected from NED, only 121 were included in XCS-DR1. That said, NED still contributed more redshifts to the catalogue than any of the other optical follow-up methods used in Mehrtens et al. (in preparation).

In Section 4.2 we have described and verified an automated method to derive X-ray temperatures from the *XMM* archive. We have shown that reliable T_X values can be obtained for most clusters if more than 300 soft-band counts are available in the background-subtracted spectrum (Figs 16 and 17). We have further shown that our technique works well even at large off-axis angles (Fig. 18) and that our automatically generated results are consistent with those derived by other authors using more traditional spectral fitting methods (Fig. 19). We note that being able to fit T_X down to 300 counts was not anticipated in R01, where we assumed the minimum counts threshold for T_X measurement would be 1000. This decrease can be attributed to our adaptive spectral binning technique and our use of Cash (rather than Gaussian) statistics in the fitting.

In R01 we predicted that up to 1800 *XMM* clusters might yield temperatures (with <20 per cent errors). By comparison (by 2011 May 01), we had made only 292 T_X measurements (with <20 per cent errors) for candidates with optically determined redshifts, although, when the error threshold is relaxed to <40 per cent, the number rises to 587. Of these 587, 357 were determined from candidates detected with 300 or more background-subtracted counts (Fig. 24). Even when we set the error threshold at 10 per cent (the calibration uncertainty for the satellite), we still have 122 clusters (112 with over 300 counts) remaining. For these 122, it would not be worth doing further *XMM* follow-up, although some high-resolution Chandra imaging would be worthwhile to elucidate the impact of point-source contamination on derived temperatures (as we have done successfully for *XMM* XCS J2215.9−1738, see Hilton et al. 2010). Errors on T_X of 40 per cent are too large for some of the science applications we have in mind for XCS, e.g. studies of

the evolution in the scatter on the L_X–T_X relation (since the intrinsic scatter is <40 per cent). Therefore, we have made requests for additional *XMM* follow-up of certain XCS clusters. We note that only 278 of the 587 candidates with T_X measurements (<40 per cent errors) are included in XCS-DR1.[5] That is not to say that the remainder are not clusters, but rather that they cannot be confirmed as being so using the currently available optical and X-ray data (see Mehrtens et al., in preparation). Even so, the size of the XCS-DR1 T_X sample is still much larger than any previous published compilations of cluster T_X measurements (with either heterogeneous or homogeneous selection).

In Section 4.3 we described and verified an automated method to derive X-ray luminosities from the *XMM* archive, the L_X-pipeline. This pipeline is run on any candidate for which a T_X measurement has been made via the T_X-pipeline. We demonstrated that the parameters that come out of the spatial fitting are robust, as compared to previously published work (Figs 21 and 22). Limitations of the current method include reliance on circularly symmetric models and a lack of covariance information between the T_X-pipeline and L_X-pipelines. Addressing these two issues is possible, but given that we are often fitting to only a few hundred counts, and currently using a circularly symmetric PSF, we have no plans to adjust the L_X-pipeline accordingly (because it would increase the computational complexity significantly). A further limitation of the current L_X-pipeline is that the error on the input redshift is assumed to be zero. This simplification is justified for spectroscopically determined redshifts, but not for photometric, or X-ray (see below), redshifts. This issue will be addressed before the L_X values are used in a future study of the evolution of the L_X–T_X relation.

The impact of redshift errors notwithstanding, the uncertainty on the L_X value for a given candidate will be much smaller than the associated error on T_X (because the ICM emission is only a weak function of T_X). Therefore, we do not think it is necessary to carry out a large-scale *XMM* follow-up campaign in order to improve the L_X measurements. That said, we are planning to request *XMM* snapshots of clusters that were discovered so close to the edge of the field of view that a large fraction of their flux was not captured in their respective EPIC images. These snapshots will allow us to get a better estimate of their total flux. We also plan to make *Chandra*

---

[4] In principle we would have been prepared to assign $z_{lit} < 0.08$ values to XCS-DR1 clusters if they had a measured photometric redshift of $z_{phot} \leq 0.1$. However, in practice there were no such cases, see Mehrtens et al. (in preparation).

[5] An additional 124 more T_X measurements with larger errors are also included in XCS-DR1.
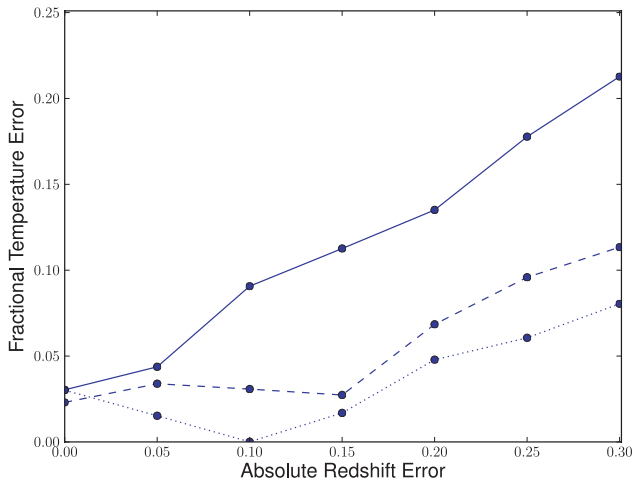
Figure 25. The results of fitting simulated MEKAL spectra with a temperature of 3.0 keV and redshifts of 0.2 (solid), 0.6 (dashed) and 1.0 (dotted) showing fractional temperature uncertainty as a function of absolute redshift uncertainty.

snapshot requests of a representative subsample of XCS clusters, in order to gauge, in a statistical sense, the impact of point source contamination on XCS L_X values (although this test may be possible using the existing Chandra archive and we will explore that avenue first).

In Section 4.4 we described a method to extract X-ray redshifts directly from the discovery data to supplement the XCS optical follow-up efforts. As shown in Fig. 23, acceptable ($\Delta_z < 0.1$) redshift measurements are made in $\simeq 75$ per cent of the cases when thresholds on the z_X-pipeline errors are set at <20 per cent and $\sigma_{z_X} < 0.05$. To date, 125 candidates have yielded z_X measurements that meet these criteria. We have used z_X estimates to pre-select candidates for optical follow-up and this approach has been successful, e.g. one cluster with $z_X = 0.84$ was demonstrated to have a true redshift of $z = 0.83$ based on subsequent Gemini GMOS spectroscopic observations (Mehrtens et al., in preparation). The level of catastrophic redshift errors is much higher for X-ray redshifts than for optical photometric techniques, so all clusters with only z_X values will ultimately have to be followed up with optical photometry or spectroscopy.

The impact of z_X errors on L_X measurements is significant, and so L_X values that rely on z_X will not be used for science applications. However, we have determined that the impact of z_X errors on T_X measurements is not significant: as shown in Fig. 25, an absolute redshift uncertainty of $\Delta_z = 0.3$ induces a less than 30 per cent T_X uncertainty. For this reason, it will be possible to use z_X-determined T_X values to select XCS clusters for SZ follow-up, because the SZ effect is (roughly speaking) redshift independent. Most current SZ instruments are only sensitive to $T_X > 5$ keV clusters, but few of those have been catalogued yet, especially at $z > 0.5$: in the BAX[6] data base there are only 39 such clusters listed. In XCS-DR1 there are 31 such clusters (of which 25 are in addition to the BAX sample). By comparison, using the X-ray redshift technique we have identified 17 more candidates (without other redshift information) meeting those criteria. In summary, X-ray redshifts are not a 'magic bullet' and optical follow-up is still required in order

to secure redshift measurements; however, they do provide some useful information, as long as they are used judiciously.

In summary, we have demonstrated that the X-ray algorithms developed for XCS are suitable for the compilation and analysis of large samples of clusters detected serendipitously by *XMM*. In our companion paper (Mehrtens et al., in preparation) we discuss the optical follow-up of candidates and present the first XCS data release (XCS-DR1), including 402 T_X measurements. On going science exploitation of XCS-DR1 includes projects related to cluster scaling relations, fossil groups, the SZ effect and the derivation of cosmological parameters. We also plan to apply our post-processing pipelines, that were designed with serendipitous clusters in mind, to the many hundreds of target clusters in the *XMM* archive, so that they too can benefit from a uniform set of T_X and L_X measurements. Even though these target clusters cannot be used for XCS statistical studies, we think this will be a valuable resource for the community, especially now that Planck is in full operation.

6 CONCLUSIONS

(i) We have demonstrated that the *XMM* archive is a rich resource for serendipitous cluster detection out to redshifts of at least $z = 1.5$.

(ii) The archive now covers over 600 deg² that can be used for serendipitous source detection and, of this, 51, 276 and 410 deg² (at >40, >10 and >0 ks depths, respectively) are available for cluster detection.

(iii) We have shown that typically one-third of a given *XMM* exposure is rendered unusable due to background flares.

(iv) We have shown that it is possible to exploit the whole *XMM* archive in a uniform and reproducible way.

(v) We have developed a source detection pipeline that operates across the entire *XMM* field of view, and is effective over a wide range of angular scales and signal-to-noise ratios. It has many features, including the ability to determine which sources are extended beyond the PSF model and to detect point-like sources that lie along the line of sight to extended sources.

(vi) We have developed a pipeline that can measure reliable X-ray cluster temperatures. This pipeline has been shown to work well even when the cluster is discovered on the outskirts of the field of view. We have demonstrated that with 300 or more background-subtracted counts, one can measure robust, unbiased temperatures for most clusters.

(vii) We have developed a pipeline that can measure reliable X-ray luminosities by making spatial fits to *XMM* images. The derived luminosity values have been shown to be robust, as have the fitted spatial parameters.

(viii) We have developed a pipeline that can measure 'X-ray redshifts' for clusters using *XMM* spectra. These redshifts can help increase the number of clusters with X-ray temperature measurements; acceptable ($\Delta_z < 0.1$) redshift measurements are made in $\simeq 75$ per cent of the cases (once errors thresholds have been imposed).

(ix) To date (2011 May 01), some key statistics for XCS are as follows: 5776 ObsIDs have been downloaded from the *XMM* archive; 5642 ObsIDs have run through the event list cleaning pipeline; 4029 ObsIDs have been processed by the source detection pipeline; 114 711 point sources and 12 582 extended sources have been catalogued; 3675 cluster candidates have been selected, of which 993 were detected with more than 300 background-subtracted counts; 587 (122) X-ray temperatures have been measured with <40 (<10) per cent errors.

REFERENCES

Abazajian K. N. et al., 2009, ApJS, 182, 543
Abbey T., Carpenter J., Read A., Wells A., XMM Science Centre Swift Mission Operations Center 2006, in Wilson A., ed., ESA Special Publ. Vol. 604, The X-ray Universe 2005. ESA, Noordwijk, p. 943
Abell A. O., 1958, ApJS, 3, 211
Adami C., Ulmer M. P., Romer A. K., Nichol R. C., Holden B. P., Pildis R. A., 2000, ApJS, 131, 391
Adami C. et al., 2011, A&A, 526, A18
Ade P. A. R. et al., 2011, preprint (arXiv:1101.2024)
Allen S. W., Evrard A. E., Mantz A. B., 2011, ARA&A, 49, 409
Alshino A., Ponman T., Pacaud F., Pierre M., 2010, MNRAS, 407, 2543
Altieri B., Chen B., Gabriel C., Gondoin P., Lammers U., Lumb D., Kirsch., 2004, Technical Report XMM-PS-GM-20 Issue 3.1, XMM-Newton Calibration Access and Data Handbook. ESA, Noordwijk
Anders E., Grevesse N., 1989, Geochimica et Cosmochimica Acta, 53, 197
Arabadjis J. S., Bregman J. N., 1999, ApJ, 510, 806
Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, ASP Conf. Ser. Vol. 101, Astronomical Data Analysis Software and Systems V. Astron. Soc. Pac., San Franciso, p. 17
Arnaud M., Evrard A. E., 1999, MNRAS, 305, 631
Arnaud M., Pointecouteau E., Pratt G. W., 2005, A&A, 441, 893
Arnaud M., Pratt G. W., Piffaretti R., Böhringer H., Croston J. H., Pointecouteau E., 2010, A&A, 517, A92
Arviset C., Guainazzi M., Hernandez J., Dowson J., Osuna P., Venet A., 2002, ArXiv Astrophysics e-prints
Arviset C., Osuna P., Salgado J., 2004, in Ochsenbein F., Allen M. G., Egret D., eds, ASP Conf. Ser. Vol. 314, Astronomical Data Analysis Software and Systems (ADASS) XIII. Astron. Soc. Pac., San Francisco, p. 574

Barkhouse W. A. et al., 2006, ApJ, 645, 955
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Bird C., Mushotzky R. F., Metzler C. A., 1995, ApJ, 453, 40
Blackburn J. K., 1995, in Shaw R. A., Payne H. E., Hayes J. J. E., eds, ASP Conf. Ser. Vol. 77, Astronomical Data Analysis Software and Systems IV. Astron. Soc. Pac., San Francisco, p. 367
Blanton E. L., Gregg M. D., Helfand D. J., Becker R. H., White R. L., 2003, AJ, 125, 1635
Böhringer H. et al., 2000, ApJS, 129, 435
Böhringer H. et al., 2004, A&A, 425, 367
Boyle W. S., Smith G. E., 1970, Bell Sys. Tech. J., 49, 587
Bremer M. N. et al., 2006, MNRAS, 371, 1427
Burenin R. A., Vikhlinin A., Hornstrup A., Ebeling H., Quintana H., Mescheryakov A., 2007, ApJS, 172, 561
Burke D. J., Collins C. A., Sharples R. M., Romer A. K., Nichol R. C., 2003, MNRAS, 341, 1093
Cash W., 1979, ApJ, 228, 939
Cavaliere A., Fusco-Femiano R., 1976, A&A, 49, 137
Clavel J., 1998, in Science with XMM, XMM Science Archive
Collins C. A. et al., 2009, Nat, 458, 603
Croston J. H. et al., 2008, A&A, 487, 431
Cruddace R. et al., 2002, ApJS, 140, 239
Cunha C., Huterer D., Frieman J. A., 2009, Phys. Rev. D, 80, 063532
Davidson M., 2006, PhD thesis, Univ. Edinburgh
Deponte Evans J., Cresitello-Dittmar M., Doe S., Evans I. N., Germain G., Glotfelty K., Overly J., 2008, in Argyle R. W., Bunclark P. S., Lewis J. R., eds, ASP Conf. Ser. Vol. 394, Astronomical Data Analysis Software and Systems XVII. Astron. Soc. Pac., San Francisco, p. 627
Dickey J. M., Lockman F. J., 1990, ARA&A, 28, 215
Doe S., Noble M., Smith R., 2001, in Harnden F. R., Jr, Primini F. A., Payne H. E., eds, ASP Conf. Ser. Vol. 238, Astronomical Data Analysis Software and Systems X. Astron. Soc. Pac., San Francisco, p. 310
Dunkley J. et al., 2011, ApJ, 739, 52
Dupke R. A., Bregman J. N., 2001a, ApJ, 562, 266
Dupke R. A., Bregman J. N., 2001b, ApJ, 547, 705
Ebeling H., Edge A. C., Böhringer H., Allen S. W., Crawford C. S., Fabian A. C., Voges W., Huchra J. P., 1998, MNRAS, 301, 881
Ebeling H., Edge A., Allen S., Crawford C., Fabian A., Huchra J., 2000, MNRAS, 318, 333
Ebeling H., Edge A. C., Henry J. P., 2001, ApJ, 553, 668
Ebeling H., Mullis C. R., Tully R. B., 2002, ApJ, 580, 774
Fassbender R., Böhringer H., Lamer G., Mullis C. R., Rosati P., Schwope A., Kohnert J., Santos J. S., 2008, A&A, 481, L73
Fassbender R. et al., 2011, A&A, 527, 78
Finoguenov A. et al., 2007, ApJS, 172, 182
Finoguenov A. et al., 2010, MNRAS, 403, 2063
Foley R. J. et al., 2011, ApJ, 731, 86
Freeman P. E., Kashyap V., Rosner R., Lamb D. Q., 2002, ApJS, 138, 185 (F02)
Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1990, ApJ, 351, 10
Fruscione A. et al., 2006, in Silva D. R., Doxsey R. E., eds, SPIE Conf. Ser. Vol. 6270, Observatory Operations: Strategies, Processes, and Systems. SPIE, Bellingham, 62701V
Gabriel C. et al., 2004, in Ochsenbein F., Allen M. G., Egret D., eds, ASP Conf. Ser. Vol. 314, Astronomical Data Analysis Software and Systems (ADASS) XIII. Astron. Soc. Pac., San Francisco, p. 759
Gastaldello F., Buote D. A., Humphrey P. J., Zappacosta L., Seigar M. S., Barth A. J., Brighenti F., Mathews W. G., 2007, ApJ, 662, 923
Ghizzardi S., 2001, Technical Report EPIC-MCT-TN-011, In Flight Calibration of the PSF for the MOS1 and MOS2 Cameras. ESA, Noordwijk
Ghizzardi S., 2002, Technical Report EPIC-MCT-TN-012, In Flight Calibration of the PSF for the PN Camera. ESA, Noordwijk
Gioia I. M., Maccacaro T., Schild R. E., Wolter A., Stocke J. T., Morris S. L., Henry J. P., 1990, ApJS, 72, 567
Gioia I. M., Henry J. P., Mullis C. R., Böhringer H., Briel U. G., Voges W., Huchra J. P., 2003, ApJS, 149, 29
Gladders M. D., Yee H. K. C., 2000, AJ, 120, 2148
Gobat R. et al., 2011, A&A, 526, A133

Gondoin P., Beijersbergen M., Willingale R., de Chambure D., van Katwijk K., Lumb D. H., Peacock A. J., 1996, in Siegmund O. H., Gummin M. A., eds, SPIE Conf. Ser. Vol. 2808, Simulation of the XMM Mirror Performance Based on Metrology Data. SPIE, Bellingham, p. 376

Gondoin P., Aschenbach B. R., Beijersbergen M. W., Egger R., Jansen F. A., Stockman Y., Tock J., 1998, in Hoover R. B., Walker A. B., eds, SPIE Conf. Ser Vol. 3444, Calibration of the First XMM Flight Mirror Module: I. Image Quality. SPIE, Bellingham, p. 278

Gondoin P., Aschenbach B., Erd C., Lumb D. H., Majerowicz S., Neumann D., Sauvageot J. L., 2000, in Flanagan K. A., Siegmund O. H., eds, SPIE Conf. Ser. Vol. 4140, In-orbit Calibration of the *XMM–Newton* Telescopes. SPIE, Bellingham, p. 1

Gonzalez A. H., Zaritsky D., Zabludoff A. I., 2007, ApJ, 666, 147

Haberl F., 2007, Ap&SS, 308, 181

Hashimoto Y., Barcons X., Böhringer H., Fabian A. C., Hasinger G., Mainieri V., Brunner H., 2004, A&A, 417, 819

Henry J. P., 2004, ApJ, 609, 603

Henry J. P., Mullis C. R., Voges W., Böhringer H., Briel U. G., Gioia I. M., Huchra J. P., 2006, ApJS, 162, 304

Henry J. P. et al., 2010, ApJ, 725, 615

Hilton M. et al., 2007, ApJ, 670, 1000

Hilton M. et al., 2009, ApJ, 697, 436

Hilton M. et al., 2010, ApJ, 718, 133

Hoeft M., Lamer G., Kohnert J., Schwope A., 2008, preprint (arXiv:0809.2265)

Horner D. J., 2001, PhD thesis, Univ. Maryland, College Park

Horner D. J., Perlman E. S., Ebeling H., Jones L. R., Scharf C. A., Wegner G., Malkan M., Maughan B., 2008, ApJS, 176, 374

Hosmer M., 2010, PhD thesis, Astronomy Centre, Univ. Sussex

Hoyle B., Jimenez R., Verde L., 2010, preprint (arXiv:1009.3884)

James F., Roos M., 1975, Comput. Phys. Communications, 10, 343

Jansen F. A., Laine R., 1997, BAAS, 29, 1365

Jones C., Forman W., 1984, ApJ, 276, 38

Kay S. T., da Silva A. C., Aghanim N., Blanchard A., Liddle A. R., Puget J.-L., Sadat R., Thomas P. A., 2007, MNRAS, 377, 317

Kessler R. et al., 2009, ApJS, 185, 32

Koester B. P. et al., 2007, ApJ, 660, 239

Kolokotronis V., Georgakakis A., Basilakos S., Kitsionas S., Plionis M., Georgantopoulos I., Gaga T., 2006, MNRAS, 366, 163

Lamer G., Hoeft M., Kohnert J., Schwope A., Storm J., 2008, A&A, 487, L33

Larson D. et al., 2011, ApJS, 192, 16

Lloyd-Davies E. J., Ponman T. J., Cannon D. B., 2000, MNRAS, 315, 689

Lumb D. H., Warwick R. S., Page M., De Luca A., 2002, A&A, 389, 93

Majumdar S., Mohr J. J., 2004, ApJ, 613, 41

Mantz A., Allen S. W., Ebeling H., Rapetti D., Drlica-Wagner A., 2010, MNRAS, 406, 1773

Mantz A., Allen S. W., Rapetti D., 2010, MNRAS, 406, 1805

Marriage T. A. et al., 2011, ApJ, 737, 61

Maughan B. J., 2007, ApJ, 668, 772

Maughan B. J., Jones C., Forman W., Van Speybroeck L., 2008, ApJS, 174, 117

Mehrtens N. et al., 2011, preprint (arXiv:1106.3056)

Melin J., Bartlett J. G., Delabrouille J., 2005, A&A, 429, 417

Menanteau F. et al., 2010, ApJ, 123, 1523

Mewe R., Lemen J. R., van den Oord G., 1986, A&A, 65, 511

Miller C. J. et al., 2005, AJ, 130, 968

Morrison R., McCammon D., 1983, ApJ, 270, 119

Mullis C. R. et al., 2003, ApJ, 594, 154

Mullis C. R., Rosati P., Lamer G., Böhringer H., Schwope A., Schuecker P., Fassbender R., 2005, ApJ, 623, L85

Mushotzky R. F., Done C., Pounds K. A., 1993, ARA&A, 31, 717

Olsen L. F. et al., 2008, A&A, 478, 93

Onuora L. I., Kay S. T., Thomas P. A., 2003, MNRAS, 341, 1246

Oukbir J., Blanchard A., 1992, A&A, 262, L21

Pacaud F. et al., 2007, MNRAS, 382, 1289

Papovich C. et al., 2010, ApJ, 716, 1503

Perlman E. S., Horner D. J., Jones L. R., Scharf C. A., Ebeling H., Wegner G., Malkan M., 2002, ApJS, 140, 265

Piccinotti G., Mushotzky R. F., Boldt E. A., Holt S. S., Marshall F. E., Serlemitsos P. J., Shafer R. A., 1982, ApJ, 253, 485

Pierre M. et al., 2006, MNRAS, 372, 591

Ponman T. J., Cannon D. B., Navarro J. F., 1999, Nat, 397, 135

Pradas J., Kerp J., 2005, A&A, 443, 721

Pratt G. W., Croston J. H., Arnaud M., Böhringer H., 2009, A&A, 498, 361

Predehl P. et al., 2006, in Turner M. J. L., Hasinger G., eds, Proc. SPIE Vol. 6266, Space Telescopes and Instrumentation II: Ultraviolet to Gamma Ray. SPIE, Bellingham, 62660P

Rapetti D., Allen S. W., Mantz A., Ebeling H., 2010, MNRAS, 406, 1796

Read A. M., 2004, Technical Report XMM-CCF-REl-167, PSF of the X-ray telescopes. ESA, Noordwijk

Read A. M., Ponman T. J., 2003, A&A, 409, 395

Read A. M., Saxton R., Guainazzi M., Rosen S., Stuhlinger M., 2010, Technical Report XMM-CCF-REl-263, 2-D PSF Parametrization. ESA, Noordwijk

Reiprich T. H., Böhringer H., 2002, ApJ, 567, 716

Romer A. K. et al., 2000, ApJS, 126, 209

Romer A. K., Viana P. T. P., Liddle A. R., Mann R. G., 2001, ApJ, 547, 594 (R01)

Rosati P., della Ceca R., Norman C., Giacconi R., 1998, ApJ, 492, L21

Rosati P. et al., 2009, A&A, 508, 583

Rozo E. et al., 2010, ApJ, 708, 645

Sahlén M. et al., 2009, MNRAS, 397, 577

Santos J. S. et al., 2009, A&A, 501, 49

Schneider D. P. et al., 2007, AJ, 134, 102

Schwope A. D., Lamer G., Burke D., Elvis M., Watson M. G., Schulze M. P., Szokoly G., Urrutia T., 2004, Advances Space Res., 34, 2604

Schwope A. D. et al., 2010, A&A, 513, L10

Sehgal N. et al., 2011, ApJ, 732, 44

Slezak E., Bijaoui A., Mars G., 1990, A&A, 227, 301

Stanford S. A. et al., 2006, ApJ, 646, L13

Staniszewski Z. et al., 2009, ApJ, 701, 32

Stark A. A., Gammie C. F., Wilson R. W., Bally J., Linke R. A., Heiles C., Hurwitz M., 1992, ApJS, 79, 77

Stockman Y., Mazy E., Tock J. P., de Chambure D., Gondoin P., 1997, in Guyenne T.-D., ed., ESA Special Publ. Vol. 408, Environmental Testing for Space Programms. ESA, Noordwijk, p. 169

Stott J. P. et al., 2010, ApJ, 718, 23

Suhada R. et al., 2010, A&A, 514, L3

Suhada R. et al., 2011, A&A, 530, 110

Sunyaev R. A., Zeldovich Y. B., 1972, Comments Astrophys. Space Phys., 4, 173

Tanaka M., Finoguenov A., Ueda Y., 2010, ApJ, 716, L152

Vanderlinde K. et al., 2010, ApJ, 722, 1180

Véron-Cetty M., Véron P., 2006, A&A, 455, 773 (VeronCat)

Vikhlinin A. et al., 2009a, ApJ, 692, 1033

Vikhlinin A. et al., 2009b, ApJ, 692, 1060

Villa G. E. et al., 1996, in Siegmund O. H., Gummin M. A., eds, SPIE Conf. Ser. Vol. 2808, EPIC System Onboard the ESA XMM. SPIE, Bellingham, p. 402

Voit G. M., 2005, Rev. Modern Phys., 77, 207

Watson M. G. et al., 2009, A&A, 493, 339

Weisskopf M. C., 1999, preprint (arXiv:astro-ph/9912097)

Werner N. et al., 2007, A&A, 474, 707

Williamson R. et al., 2011, ApJ, 738, 139

Wilms J., Allen A., McCray R., 2000, ApJ, 542, 914

Wilson G. et al., 2009, ApJ, 698, 1943

Wittman D., Margoniner V. E., Tyson J. A., Cohen J. G., Becker A. C., Dell'Antonio I. P., 2003, ApJ, 597, 218

Wolter H., 1952a, Ann. Phys., 10, 286

Wolter H., 1952b, Ann. Phys., 10, 94

Wu H., Rozo E., Wechsler R. H., 2010, ApJ, 713, 1207

Yu H., Tozzi P., Borgani S., Rosati P., Zhu Z.-H., 2011, A&A, 529, A65

Zwicky F., Herzog E., Wild P., 1968, Catalogue of Galaxies and of Clusters of Galaxies. California Institute of Technology, Pasadena, CA
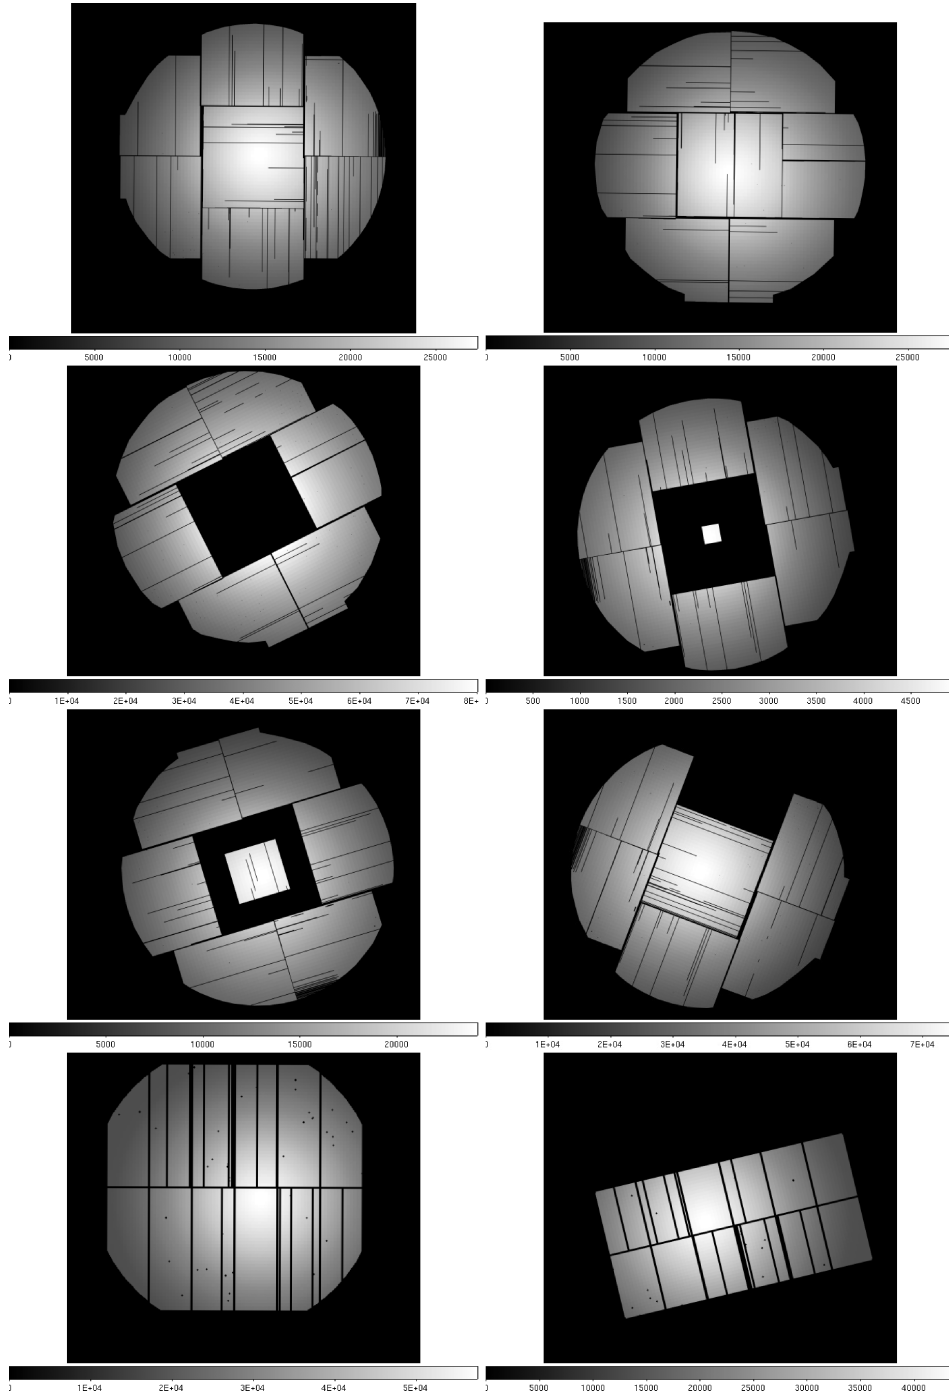
APPENDIX A: ADDITIONAL SUPPORTING FIGURES

Figure A1. Exposure maps relating to differing EPIC observing modes. In order from top left: MOS1 full window mode, MOS2 full window mode, MOS fast uncompressed, MOS partial window W2 or W4 mode, MOS partial window W3 or W5 mode, MOS1 full window mode with CCD6 switched off, pn full window mode, pn large window mode.
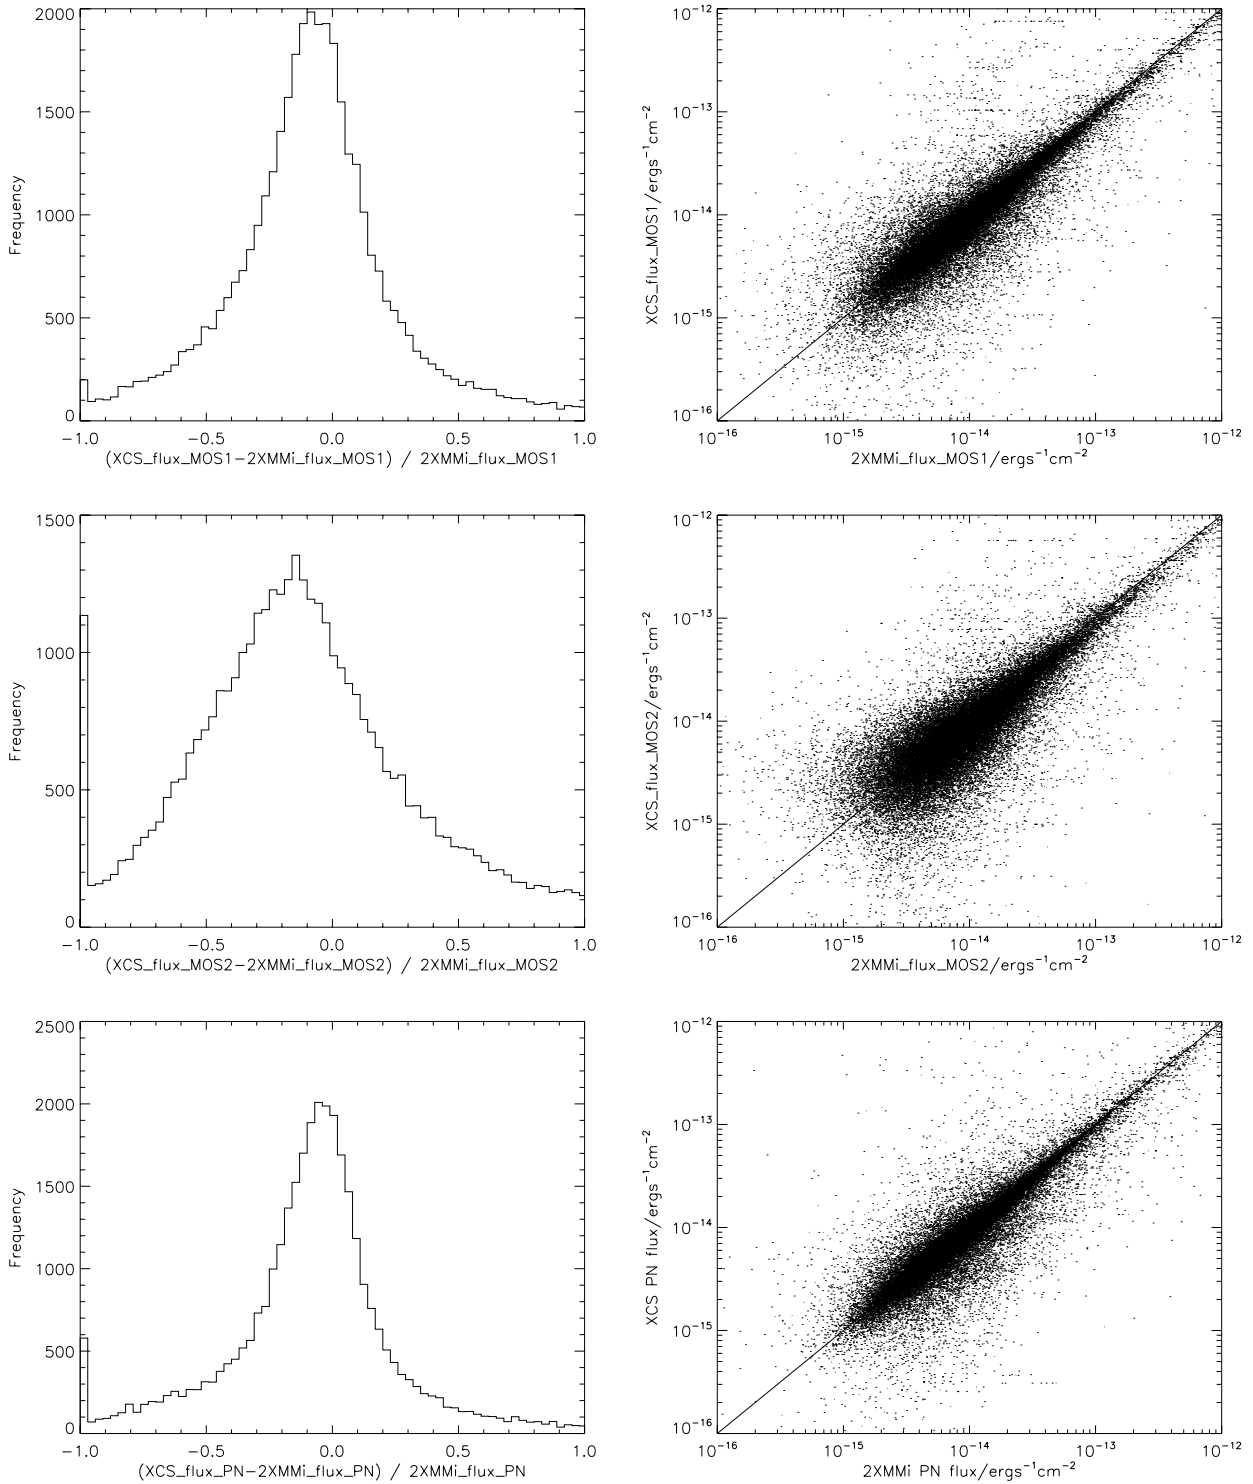
Figure A2. A comparison of the individual camera fluxes of XCS sources with their matches in the 2XMM catalogue. Top: EPIC-mos1; Middle: EPC-mos2; Bottom: EPIC-pn.
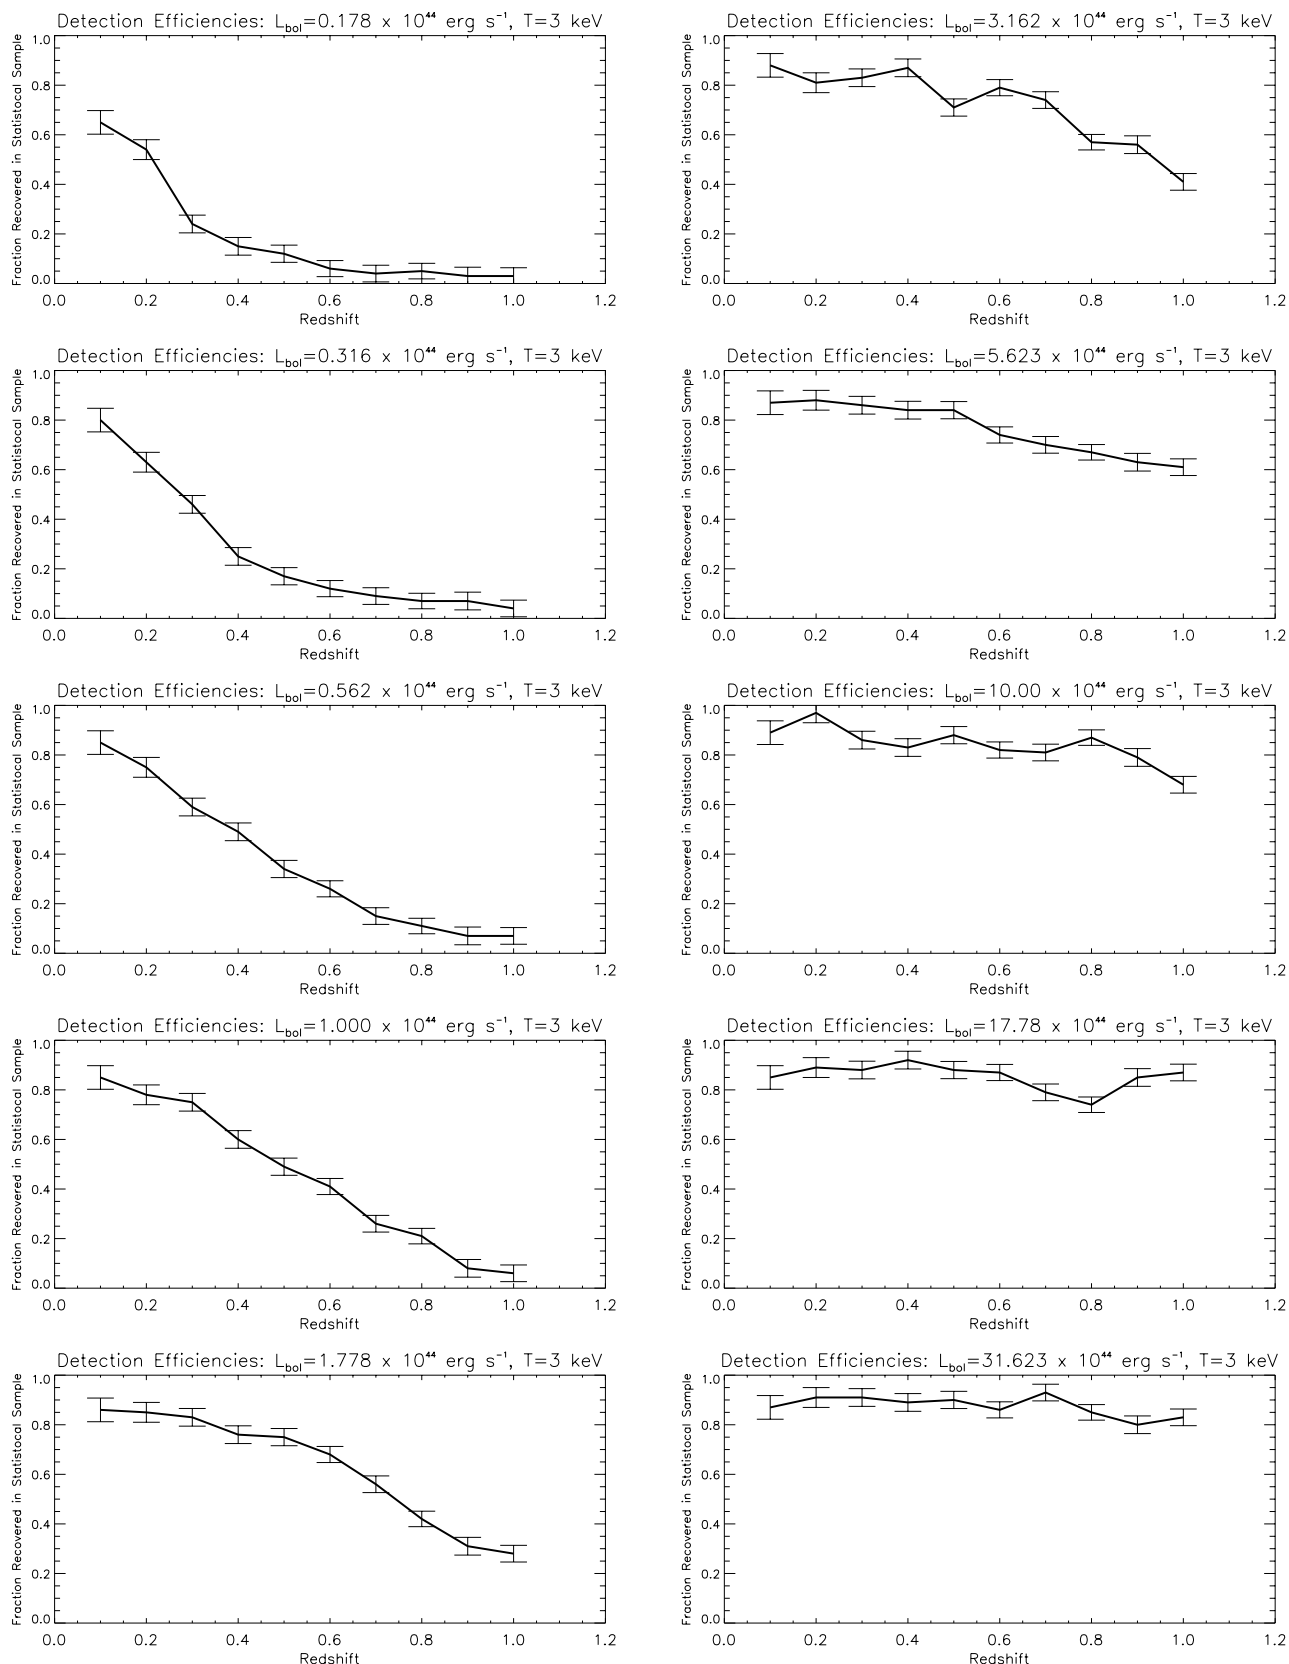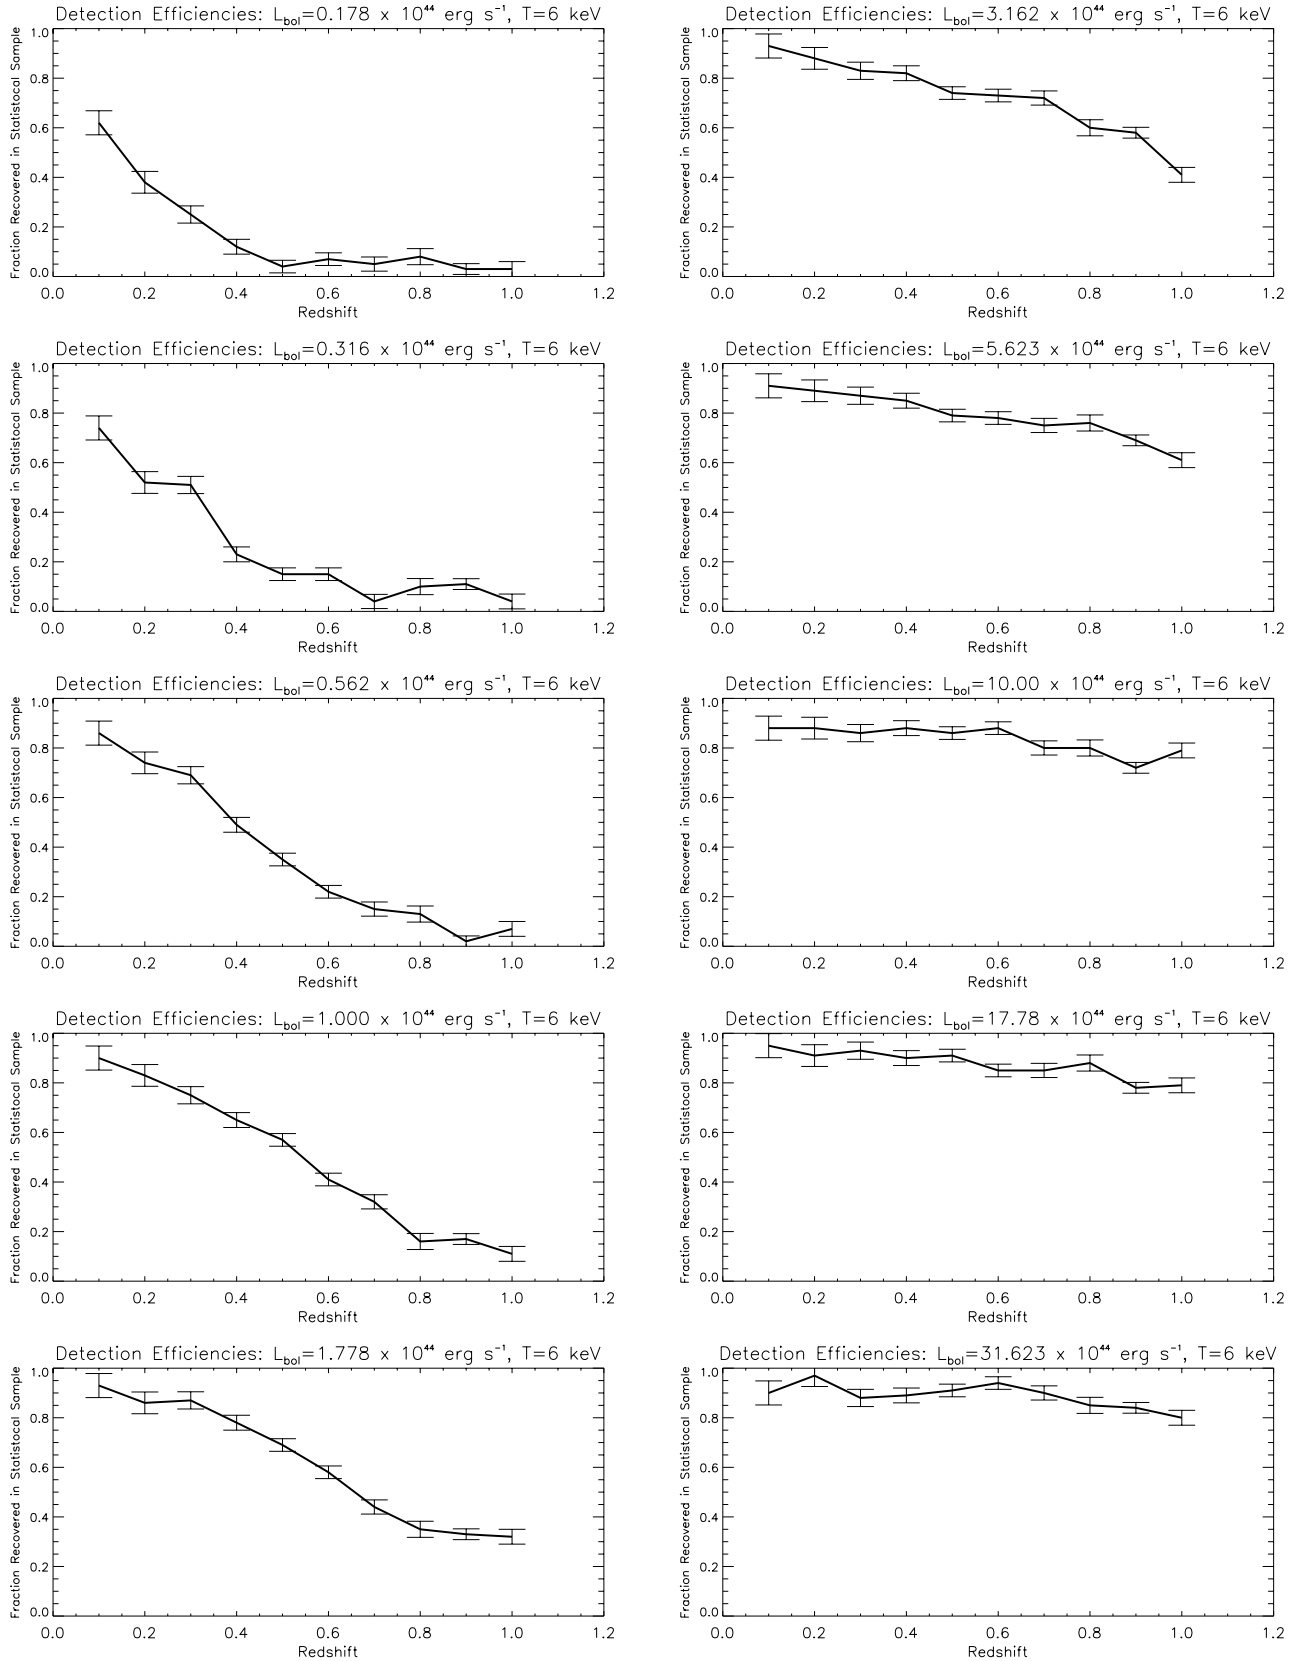
Figure A3. Predicted recovery efficiency as a function of redshift for 3-keV clusters with a range of X-ray luminosities (bolometric). The synthetic clusters used for this test had circularly symmetric β-profiles ($\beta = 2/3$) with core radii of 160 kpc. The typical luminosity of a 3-keV cluster based on the local L_X–T_X relation is 1 to 2×10^{44} erg s^{-1}.
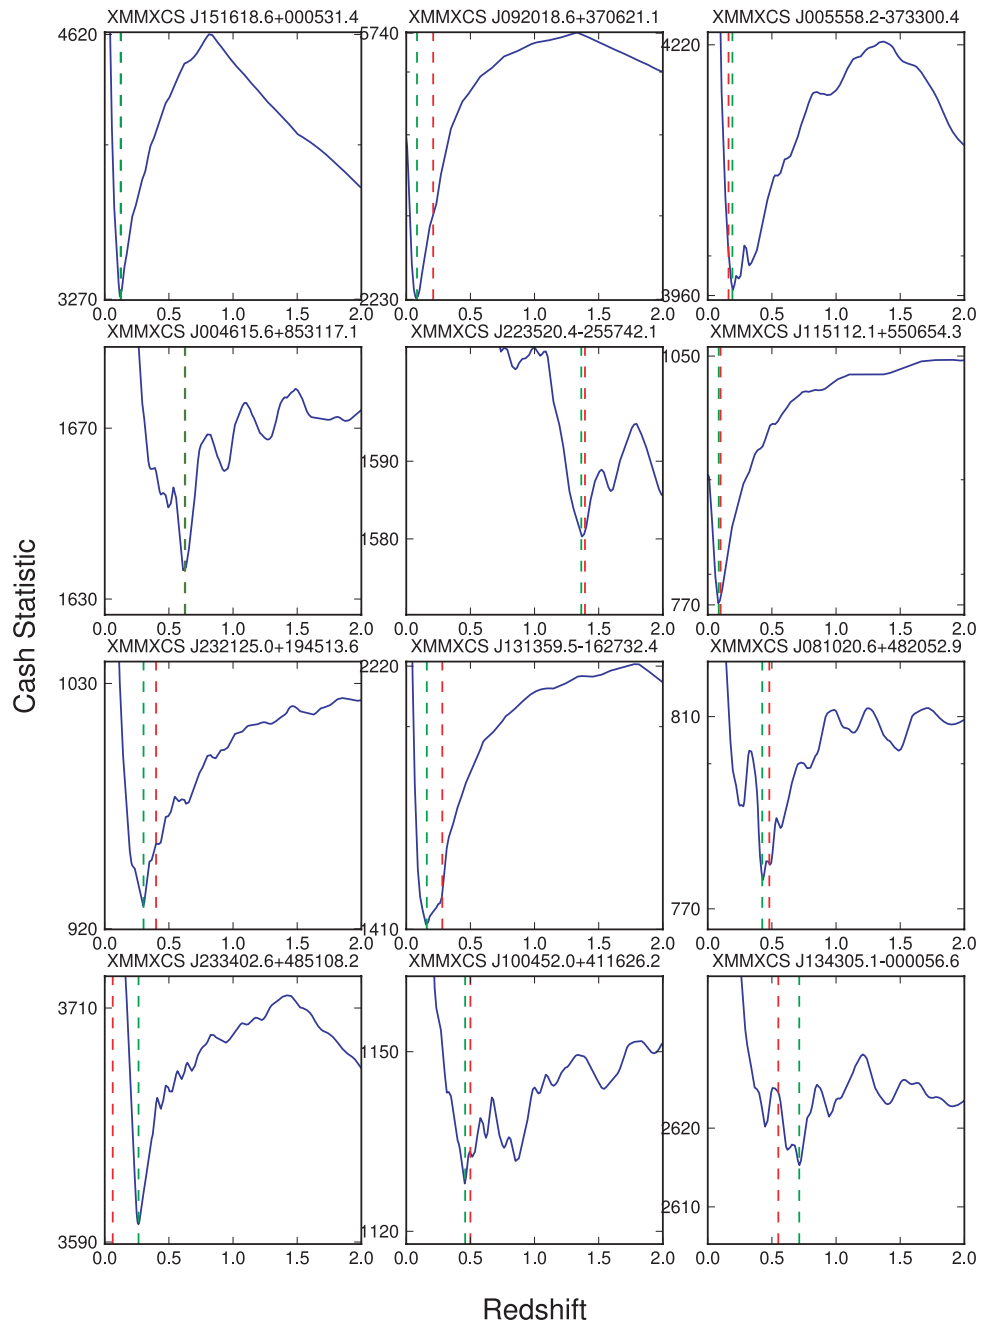
Figure A4. Predicted recovery efficiency as a function of redshift for 6-keV clusters with a range of X-ray luminosities (bolometric). The synthetic clusters used for this test had circularly symmetry β-profiles ($\beta = 2/3$) with core radii of 160 kpc. The typical luminosity of a 6-keV cluster based on the local L_X–T_X relation is 8 to 15×10^{44} erg s^{-1}.

Figure A5. Cash statistic output from the X-ray redshift fitting code, plotted against redshift. We show 12 XCS clusters that have both well measured (<2.5 per cent statistical uncertainty) X-ray redshifts and spectroscopically determined optical redshifts. The optical and X-ray redshifts are indicated with red and green dotted lines, respectively.
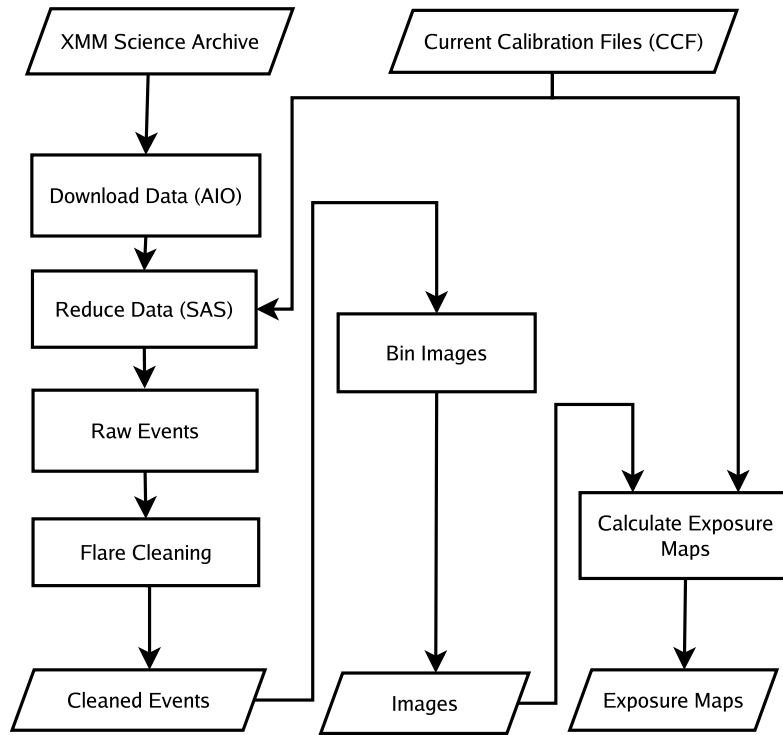
APPENDIX B: XAPA FLOW CHARTS

Figure B1. Flow chart of the process of *XMM* data reduction to produce cleaned event files, images and exposure maps. This overviews the process by which the *XMM* data are acquired, reduced and cleaned, and the products generated.
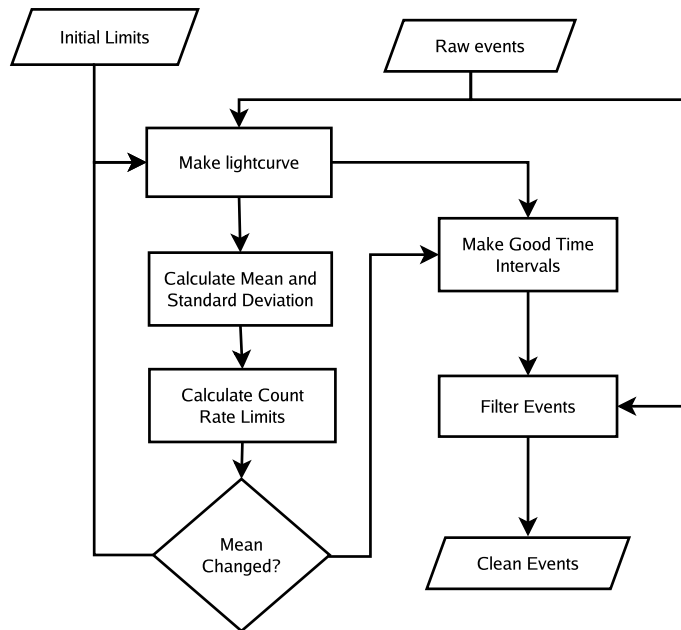
Figure B2. Flow chart of the process for removing periods of high background due to variations in the particle flux to which the instruments are exposed. This illustrates the process by which light curves of the raw event files have an iterative 3σ clipping applied to them until the mean count-rate stops changing. Cleaned event files are then produced for the good time intervals identified.
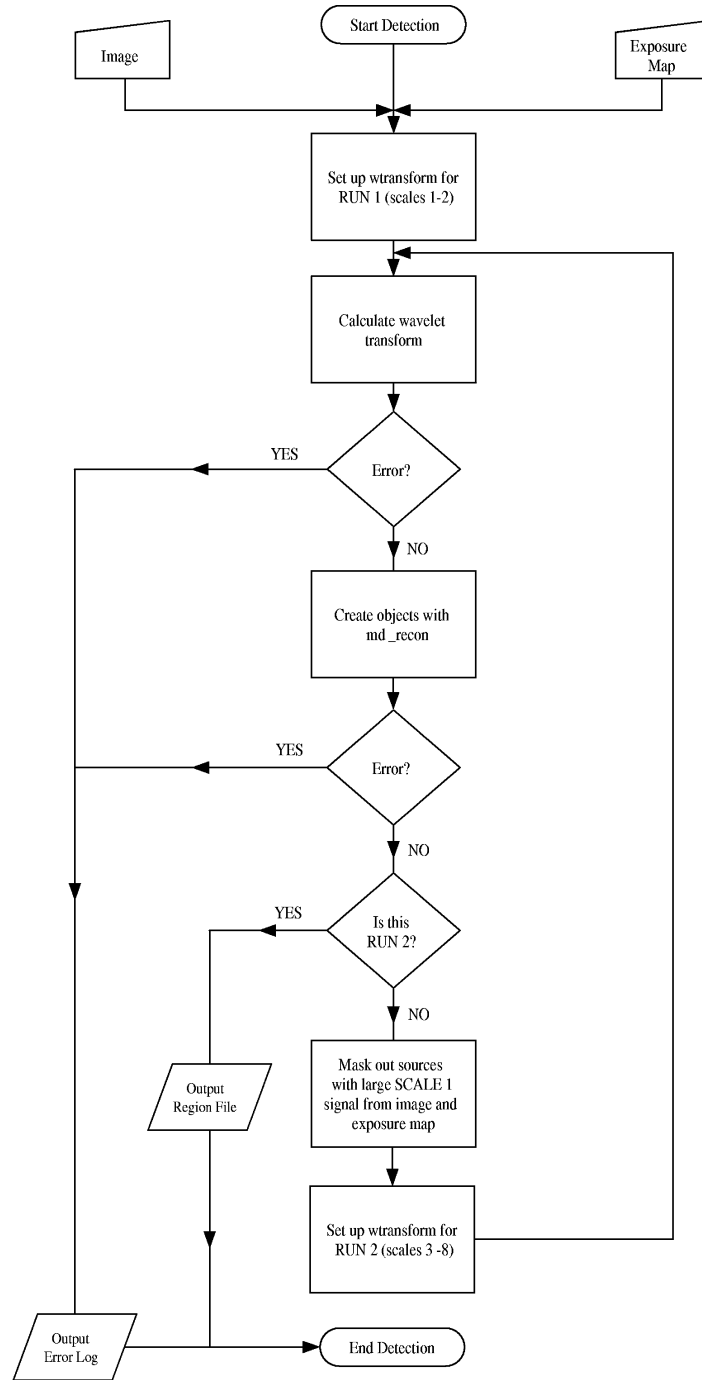
Figure B3. Flow chart for the `md_detect` routine, showing the two-stage (wavelet transform and reconstruction), two-pass (to remove pollution of the wavelet signal by bright, compact sources) process.
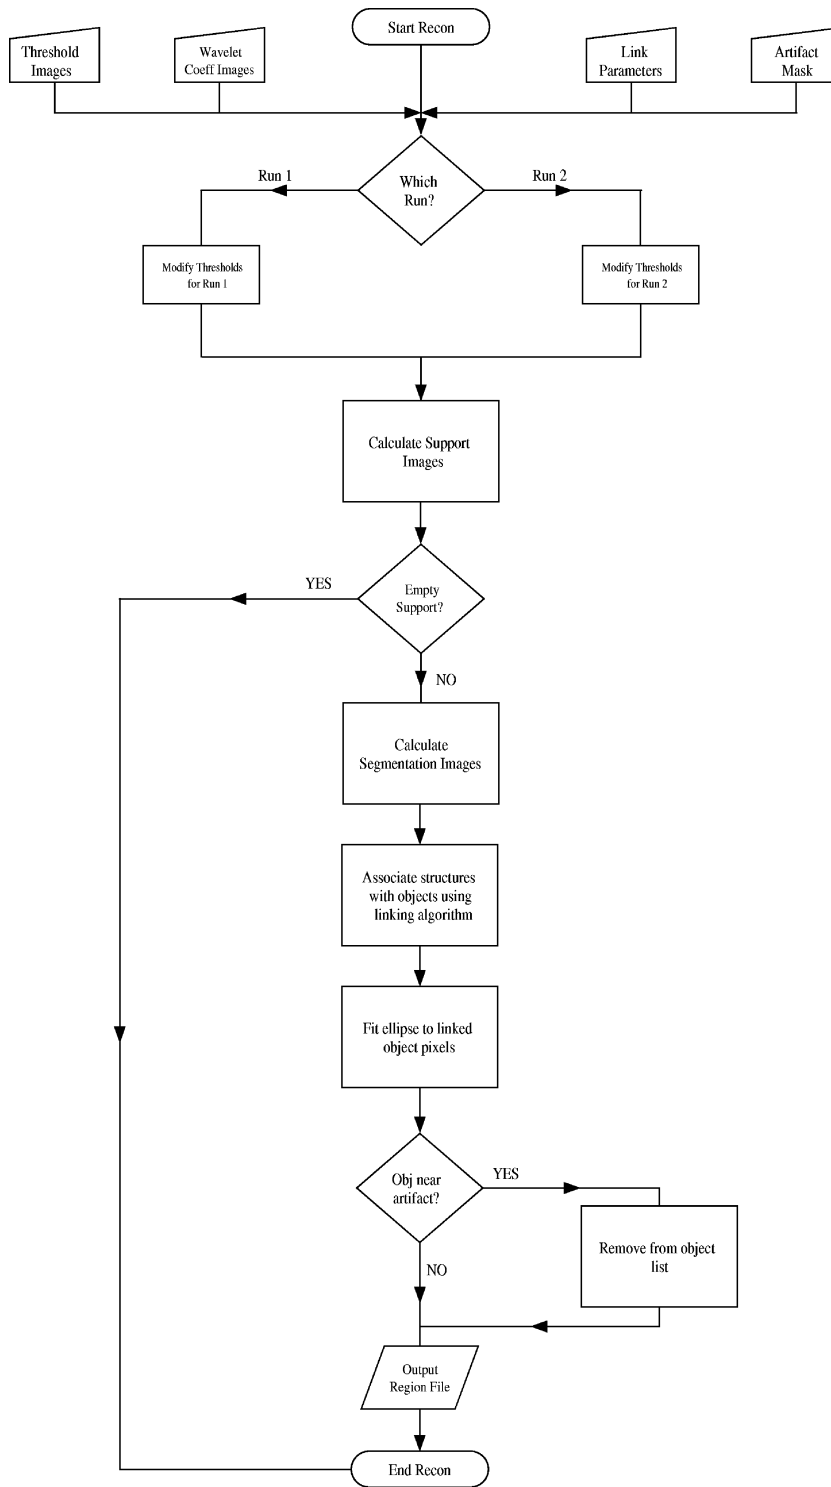
Figure B4. Flow chart for the `md_recon` routine, showing the stages of the process to reconstruct a source list from the outputs of `wtransform` on different scales.
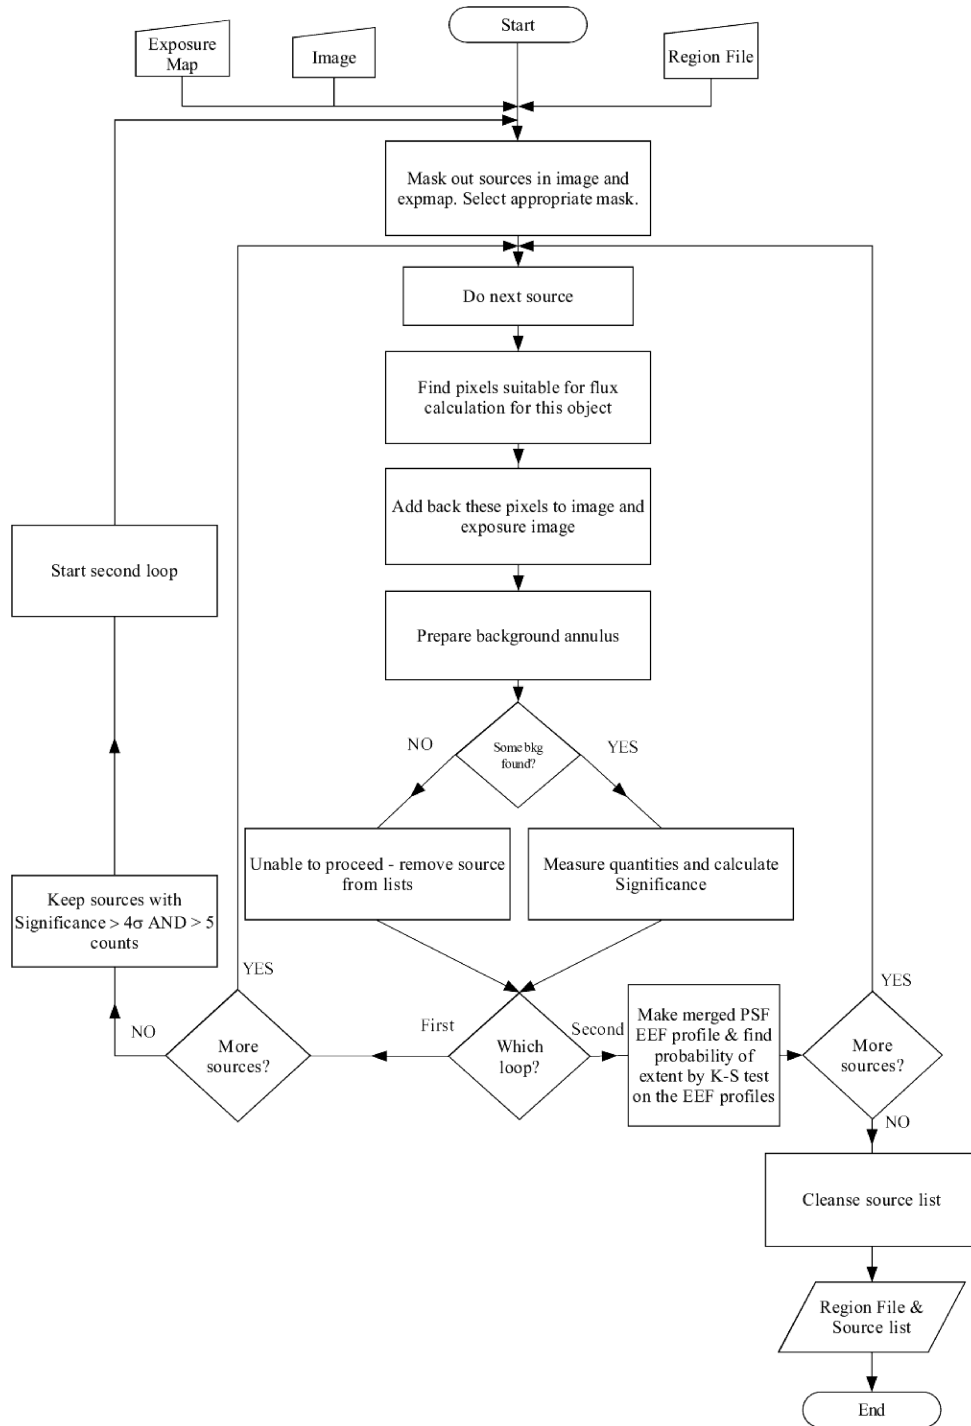
Figure B5. Flow chart for the `find_srcprop` routine, showing the stages of the process to derive the properties of each source. The routine is run in two stages: the first filters out sources with $<4\sigma$ significance, and the second determines the extent probability.
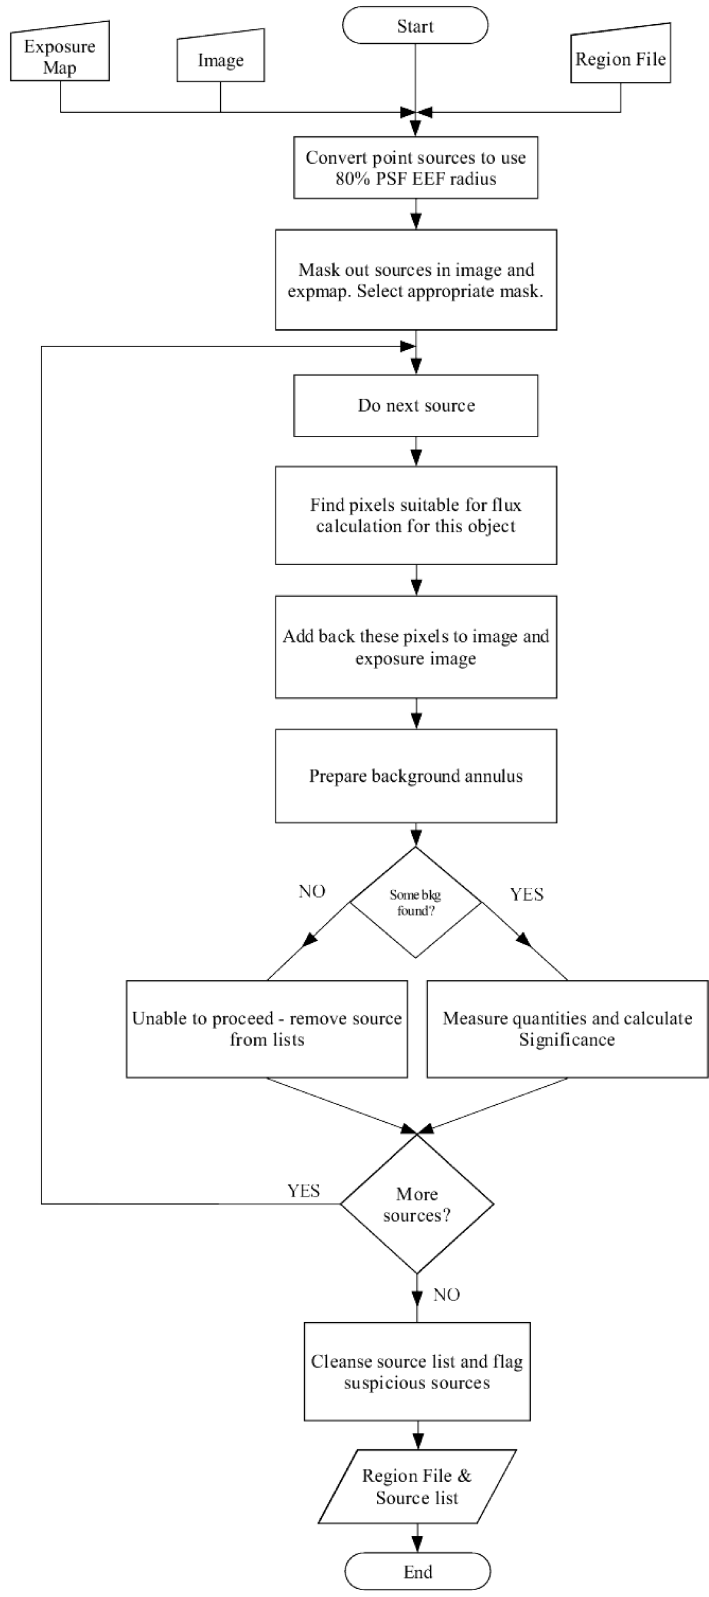
Figure B6. Flow chart for the `find_srcprop_final` routine. This routine measures properties for every source that is output from `find_srcprop`.
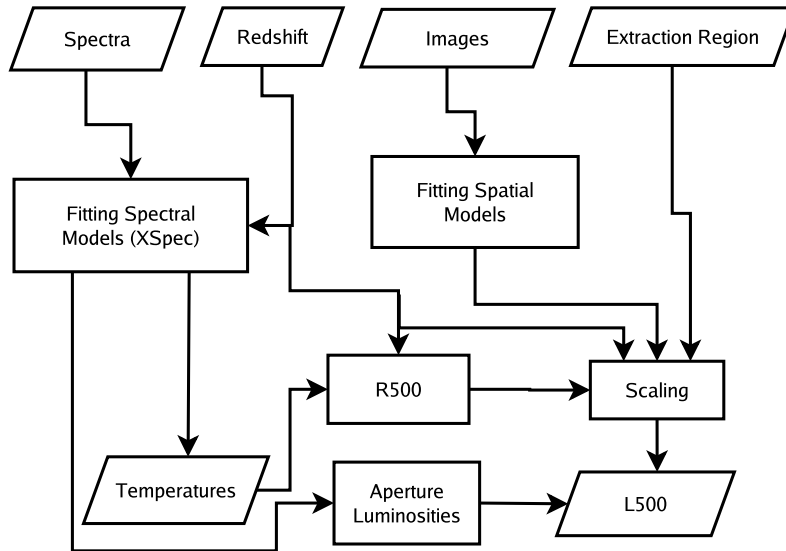
Figure B7. Flow chart of process for deriving luminosities and temperature from cluster spectra and images. This illustrates the process by which models are fitted to the X-ray spectra and images to produce temperatures and aperture luminosities, which are corrected using the fitted surface-brightness model to produce luminosities within R_{500}.
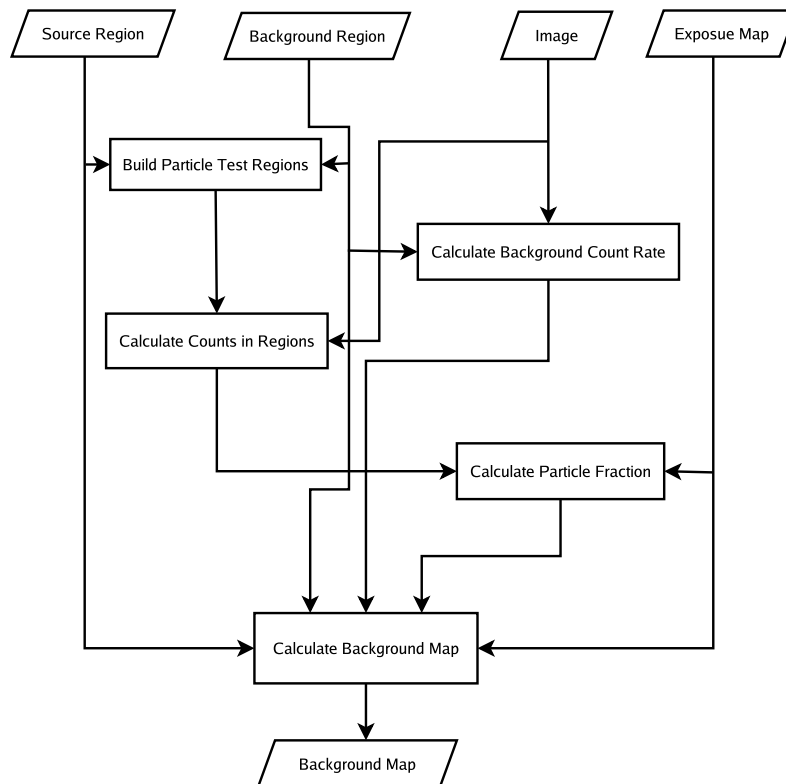

Figure B8. Flow chart of process for generating background maps for use in the XCS surface brightness fitting. This illustrates the process by which the background measured in an annulus around the source is extrapolated to all positions in the image, taking into account the exposure map and the fraction of particles in the background that are not vignetted by the telescope.

APPENDIX C: CLUSTERS USED FOR T_X AND L_X PIPELINE VALIDATION

Table C1. Clusters used for the comparison of temperature measurements with different numbers of soft-band source counts per spectrum.

Cluster Name	z	n_H (10^{20} cm^{-2})	T_X^{200} (keV)	T_X^{300} (keV)	T_X^{400} (keV)	T_X^{500} (keV)	T_X^{1000} (keV)	T_X^{2000} (keV)
XMMXCS J001737.4−005235.4	0.20	2.72	$6.50^{+3.55}_{-3.23}$	$4.82^{+2.94}_{-1.46}$	$5.36^{+3.71}_{-1.50}$	$4.94^{+1.95}_{-1.39}$	$5.24^{+0.85}_{-1.28}$	$4.57^{+0.30}_{-0.88}$
XMMXCS J092018.9+370617.7	0.19	1.57	$2.13^{+0.70}_{-0.57}$	$2.51^{+0.37}_{-0.77}$	$2.07^{+0.67}_{-0.28}$	$2.44^{+0.69}_{-0.34}$	$2.28^{+0.38}_{-0.24}$	$2.48^{+0.24}_{-0.22}$
XMMXCS J130749.6+292549.2	0.25	1.01	$2.54^{+1.27}_{-0.72}$	$2.93^{+0.97}_{-0.73}$	$2.92^{+0.92}_{-0.67}$	$2.94^{+0.86}_{-0.50}$	$3.17^{+0.65}_{-0.49}$	$3.05^{+0.45}_{-0.35}$
XMMXCS J141832.3+251104.9	0.30	1.84	–	–	$6.81^{+7.66}_{-2.21}$	$7.14^{+4.95}_{-2.45}$	$7.79^{+3.60}_{-2.56}$	$6.30^{+1.67}_{-1.18}$

Table C2. Clusters used for the on-/off-axis comparison of temperature measurements.

Cluster name	z	n_H (10^{20} cm^{-2})	On-axis T_X (keV)	Off-axis T_X (keV)
XMMXCS J100304.6+325337.9	0.42	1.55	$4.11^{+0.55}_{-0.52}$	$4.40^{+4.37}_{-2.16}$
XMMXCS J151618.6+000531.3	0.13	4.66	$5.68^{+0.22}_{-0.21}$	$5.11^{+0.47}_{-0.46}$
XMMXCS J184718.3−631959.3	0.02	6.87	$0.78^{+0.01}_{-0.02}$	$0.81^{+0.02}_{-0.03}$
XMMXCS J130832.6+534213.8	0.33	1.58	$3.66^{+0.70}_{-0.56}$	$4.45^{+4.36}_{-1.54}$
XMMXCS J072054.3+710900.5	0.23	3.88	$2.99^{+1.48}_{-0.92}$	$2.93^{+1.34}_{-0.90}$
XMMXCS J132508.7+655027.9	0.18	2.00	$1.01^{+0.20}_{-0.19}$	$0.71^{+0.44}_{-0.19}$
XMMXCS J022403.8−041333.4	1.05	2.51	$3.73^{+0.89}_{-0.67}$	$4.07^{+5.92}_{-1.71}$
XMMXCS J223520.4−255742.1	1.39	1.47	$9.45^{+3.19}_{-2.44}$	$11.29^{+14.03}_{-5.82}$

Table C3. Clusters used for the literature comparison of temperature measurements. Redshifts and T_X values for the XLSS clusters come from Pacaud et al. (2007), those for the XMMU and 2XMM clusters from Gastaldello et al. (2007) and Hoeft et al. (2008), respectively.

Cluster name	z	n_H (10^{20} cm^{-2})	T_X^{lit} (keV)	T_X^{XCS} (keV)
XMMU J131359.7−162735	0.28	4.92	$3.57^{+0.12}_{-0.12}$	$3.45^{+0.19}_{-0.19}$
2XMM J100451.6+411627	0.82	0.89	$4.2^{+0.4}_{-0.4}$	$5.36^{+0.96}_{-0.78}$
XLSS J022045.4−032558	0.33	2.49	$1.7^{+0.3}_{-0.2}$	$2.27^{+0.92}_{-0.51}$
XLSS J022145.2−034617	0.43	2.52	$4.8^{+0.6}_{-0.5}$	$5.78^{+1.83}_{-1.20}$
XLSS J022404.1−041330	1.05	2.51	$4.1^{+0.9}_{-0.7}$	$3.74^{+0.51}_{-0.38}$
XLSS J022433.8−041405	0.26	2.46	$1.3^{+0.2}_{-0.1}$	$1.16^{+0.19}_{-0.24}$
XLSS J022457.1−034856	0.61	2.49	$3.2^{+0.4}_{-0.3}$	$3.87^{+1.23}_{-0.85}$
XLSS J022524.7−044039	0.26	2.49	$2.0^{+0.2}_{-0.2}$	$2.27^{+0.99}_{-0.53}$
XLSS J022530.6−041420	0.14	2.35	$1.34^{+0.14}_{-0.06}$	$1.22^{+0.12}_{-0.13}$
XLSS J022540.6−031121	0.14	2.66	$3.5^{+0.6}_{-0.5}$	$4.16^{+0.52}_{-0.51}$
XLSS J022616.3−023957	0.06	2.67	$0.63^{+0.03}_{-0.03}$	$0.62^{+0.04}_{-0.04}$
XLSS J022722.4−032144	0.33	2.61	$2.4^{+0.5}_{-0.4}$	$2.82^{+1.02}_{-0.64}$
XLSS J022739.9−045127	0.29	2.59	$1.7^{+0.1}_{-0.1}$	$1.56^{+0.33}_{-0.18}$
XLSS J022803.4−045103	0.29	2.67	$2.8^{+0.6}_{-0.5}$	$2.38^{+2.13}_{-0.93}$

Table C4. Clusters used for the literature β comparison with the work of Alshino et al. (2010).

Cluster name	z	n_H (10^{20} cm^{-2})	β^{lit}	β^{XCS}
XLSS J022726.0−043216	0.31	2.67	$0.58^{+1.25}_{-0.14}$	$0.61^{+0.11}_{-0.21}$
XLSS J022356.5−030558	0.30	2.63	$0.43^{+0.08}_{-0.05}$	$0.43^{+0.13}_{-0.09}$
XLSS J022616.3−023957	0.06	2.67	$0.69^{+0.07}_{-0.05}$	$0.79^{+0.05}_{-0.04}$
XLSS J022045.4−032558	0.33	2.49	$0.54^{+0.03}_{-0.03}$	$0.55^{+0.06}_{-0.04}$

Table C5. Clusters used for the literature luminosity comparison with the work of Pacaud et al. (2007).

Cluster name	z	n_H (10^{20} cm^{-2})	$L^{lit}_{X,500}$ (10^{43} erg s^{-1})	XCS $L^{fit}_{X,500}$ (10^{43} erg s^{-1})
XLSS J022540.6−031121	0.14	2.66	$0.93^{+0.05}_{-0.06}$	$1.1^{+1.2}_{-1.1}$
XLSS J022616.3−023957	0.06	2.67	$0.025^{+0.002}_{-0.002}$	$0.019^{+0.050}_{-0.017}$
XLSS J022404.1−041330	1.05	2.52	$4.83^{+0.37}_{-0.34}$	$4.9^{+5.2}_{-4.5}$
XLSS J022206.7−030314	0.49	2.52	$2.89^{+0.15}_{-0.15}$	$2.3^{+2.5}_{-2.1}$
XLSS J022457.1−034856	0.61	2.49	$3.32^{+0.15}_{-0.15}$	$3.0^{+3.2}_{-2.8}$
XLSS J022045.4−032558	0.33	2.49	$0.38^{+0.03}_{-0.03}$	$0.35^{+0.42}_{-0.30}$

This paper has been typeset from a TEX/LATEX file prepared by the author.