

XZ: Deriving redshifts from X-ray spectra of obscured AGN[★]

C. Simmonds¹, J. Buchner^{1,2,3}, M. Salvato⁴, L.-T. Hsu⁵, and F. E. Bauer^{1,2}

¹ Pontificia Universidad Católica de Chile, Instituto de Astrofísica, Casilla 306, Santiago 22, Chile
e-mail: cpsimmonds@uc.cl

² Millennium Institute of Astrophysics, Vicuña, MacKenna 4860, 7820436 Macul, Santiago, Chile

³ Excellence Cluster Universe, Boltzmannstr. 2, 85748 Garching, Germany

⁴ Max Planck Institut für Extraterrestrische Physik Giessenbachstrasse, 85748 Garching, Germany

⁵ Institute of Astronomy & Astrophysics, Academia Sinica, Taipei, Taiwan

Received 11 May 2018 / Accepted 7 July 2018

ABSTRACT

Context. Redshifts are fundamental for our understanding of extragalactic X-ray sources. Ambiguous counterpart associations, expensive optical spectroscopy, and/or multimission multiwavelength coverage to resolve degeneracies often make estimation difficult in practice.

Aims. We attempt to constrain redshifts of obscured active galactic nuclei (AGN) using only low-resolution X-ray spectra.

Methods. Our method for determining redshifts from the X-ray spectrum (XZ) fits AGN X-ray spectra with a moderately complex spectral model incorporating a corona, a torus obscurer, and a warm mirror. Using the Bayesian X-ray Astronomy (BXA) package, we constrain redshift, column density, photon index, and luminosity simultaneously. The redshift information primarily comes from absorption edges in Compton-thin AGN, and from the Fe K α fluorescent line in heavily obscured AGN. A new generic background fitting method allows us to extract more information from limited numbers of source counts.

Results. We derive redshift constraints for 74/321 hard-band detected sources in the *Chandra* deep field South. Comparing with spectroscopic redshifts, we find an outlier fraction of 8%, indicating that our model assumptions are valid. For three *Chandra* deep fields, we release our XZ redshift estimates.

Conclusions. The independent XZ estimate is easy to apply and effective for a large fraction of obscured AGN in today's deep surveys without the need for any additional data. Compared to different redshift estimation methods, XZ can resolve degeneracies in photometric redshifts, help detect potential association problems, and confirm uncertain single-line spectroscopic redshifts. With high spectral resolution and a large collecting area, this technique will be highly effective for *Athena*/WFI observations.

Key words. galaxies: active – X-rays: galaxies – galaxies: distances and redshifts

1. Introduction

Measuring redshifts for distant active galactic nuclei (AGN) is notoriously difficult, yet crucial for constraining the accretion history of the Universe. Spectroscopic redshifts (hereafter *specz* values) are desirable but costly, and remain impossible for targets that are faint in the optical and near-infrared bands. Thus, it is common practice to rely on photometric redshifts (hereafter *photoz*'s), which when given a sufficient number of photometric bands can effectively provide very low-resolution spectra (e.g., Baum 1962; Koo 1985; Pello et al. 1996).

For galaxies, *photoz* values can be extremely accurate when the spectral energy distributions (SEDs) of various types have distinct sharp features (e.g., Lyman and Balmer breaks, or high equivalent width, EW, emission lines). This is complicated for AGN because the AGN component dilutes or hides the host galaxy features and has a number of high EW lines, leading to multiple degenerate redshift solutions (Salvato et al. 2009), and because the intrinsic variability of these sources makes the *photoz* gathered over years very uncertain (Simm et al. 2015).

Additionally, AGN selected in one waveband (in this case X-rays) have thus far had to be associated with counterparts to obtain *photoz* or spectroscopy. However, finding the correct counterpart when the X-ray position uncertainties are large can be ambiguous, especially in deep X-ray observations (see, e.g., Hsu et al. 2014; Salvato et al. 2018). Misassociations can affect both *photoz* and *specz* and are difficult to detect.

In this work we demonstrate that for obscured AGN there is sufficient information in the X-ray spectrum to compute a reliable redshift. Although typical redshift uncertainties are usually substantial ($\sigma(z) \sim 0.2$), they remain sufficiently accurate for verifying the association and checking for *photoz* outliers. The presented technique can be applied already with the low-resolution X-ray imaging spectrometers in current AGN surveys.

We describe in Sect. 2 the redshift information that is present in low-resolution, relatively low-count X-ray spectra. In Sect. 3 we explain our method in detail and quantify the redshift information gain. In Sect. 4 we test the reliability of our method with real survey data and provide redshift catalogues (additional fields in Appendix B). Synergies with ambiguous *specz* and *photoz* information are discussed in Sect. 5. Simulations are presented in Sect. 6 where we highlight the number of counts needed for current instruments (*Chandra*, *XMM-Newton*, *NuSTAR*, *Swift*) and future surveys with *eROSITA* and *Athena*.

[★] The XZ redshift estimates are only available at the CDS via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5) or via <http://cdsarc.u-strasbg.fr/viz-bin/qcat?J/A+A/618/A66>

2. Data

To demonstrate and validate our method, a meaningful test sample of obscured AGN with X-ray spectra and previously obtained specz and photoz information is needed. We primarily use data from the *Chandra* Deep Field-South (CDF-S; Xue et al. 2011; Rangel et al. 2013), which has extensive spectroscopic and multiwavelength photometric coverage, and thanks to its deep exposure features a large fraction of obscured sources.

The CDF-S survey (464.5 arcmin²) was reduced following the procedure of Laird et al. (2009). Here we work with the 326 hard (2–7 keV) detected X-ray sources in the 4Ms data from Rangel et al. (2013), whose X-ray identification numbers (IDs) we use throughout. Three Ms of additional X-ray data have recently been obtained (Luo et al. 2017). However, 4Ms spectra and associated data products are available from our previous efforts and these are sufficient for our purposes here, where we seek to generically demonstrate the utility of the method. Multiwavelength counterparts were found by Hsu et al. (2014) who also computed photoz data and collated specz values in a large catalogue. Because our goal is to apply an X-ray spectral model suitable for AGN, we discard sources whose photometry is best fitted with a stellar template in Hsu et al. (2014). X-ray spectra and related data products¹ were generated with ACIS EXTRACT (Broos et al. 2010) by Brightman et al. (2013). Our final X-ray sample, previously analyzed in Buchner et al. (2015), consists of 321 AGN, 264 of which are obscured ($N_{\text{H}} \geq 10^{22} \text{ cm}^{-2}$). We have optical redshift measurements for the entire sample (100% photoz; 56% specz). For photoz, we use a state-of-the-art photoz technique optimized for AGN (Hsu et al. 2014) and throughout use redshift probability distributions.

To test our technique in wider fields, we later also report results from the All Wavelength Extended Groth strip International Survey (AEGIS-XD; Nandra et al. 2015) and the Cosmological evolution Survey (C-COSMOS; Elvis et al. 2009). The X-ray spectra products were derived as in Brightman et al. (2013), with multiwavelength associations and redshift determinations by (Nandra et al. 2015; AEGIS-XD) and (Salvato et al. 2009; C-COSMOS). These samples were previously analyzed (see Buchner et al. 2015 for more details).

Our technique of determining redshifts from the X-ray spectrum (hereafter XZ) relies on fitting physically motivated templates commonly applied in obscured AGN. The intrinsic X-ray photon spectrum of AGN is to first order a power law $E^{-\Gamma}$ with photon index (typically $\Gamma = 1.9 \pm 0.15$; Nandra & Pounds 1994). Most AGN additionally feature photo-electric absorption (due to O, Fe, H, and He atoms along the line of sight) as well as Compton scattering (e.g., Rivers et al. 2013; Buchner et al. 2014). Figure 1 illustrates the resulting spectral shapes. Weak soft power law emission is also frequently observed (e.g., Brightman et al. 2014; Buchner et al. 2014), which can be explained by Thomson scattering of the intrinsic power law by a warm mirror which may bypass the obscurer.

In unobscured sources, the spectrum is essentially a featureless² power law, and little to no redshift information can be gained. However, in obscured sources the obscurer imprints absorption edges at characteristic energies. In heavily obscured AGN ($N_{\text{H}} \gtrsim 10^{24} \text{ cm}^{-2}$), the Fe K α line is prominent (Turner et al. 1998) and can pinpoint the redshift. We find in this

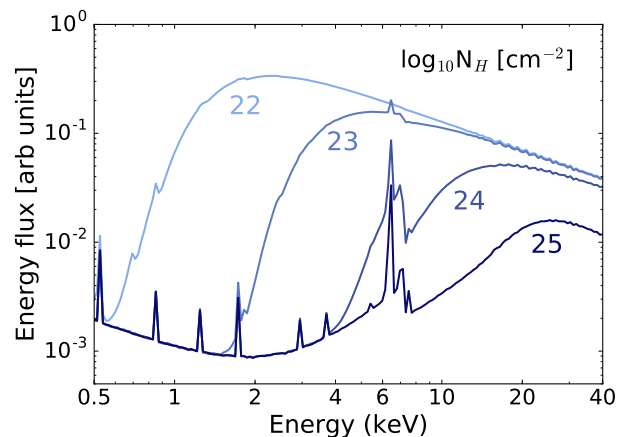


Fig. 1. AGN X-ray spectra from the TORUS model of Brightman & Nandra (2011). The intrinsic power law has a photon index of $\Gamma = 1.9$. Labeled numbers indicate the obscurer column density $N_{\text{H}} = 10^{22-25} \text{ cm}^{-2}$. At high column densities, the Fe K α line at 6.4 keV becomes prominent. A unobscured power law component normalized to 1% of the intrinsic power law has been added.

work that in mildly obscured AGN, given enough photons, the absorption edge depths and turnover energies can uniquely constrain the redshift, obscurer column density and photon index simultaneously, even using current CCD spectral resolution data.

Figure 2 shows the spectra of two illustrative examples in which the method provides precise and correct redshift estimates. In the heavily obscured case (left panel; R65), the redshift determination relies on identification of the Fe K α feature, and the redshift can be fit more or less by eye.

In the mildly obscured case (right panel in Fig. 2; R4), the Fe K α feature is not the source of redshift information. Instead, the location of the absorption edges rule out high-redshift solutions, as can be seen when comparing the two spectral fit curves. Because of the non-smooth instrument response the edges are difficult to identify by eye, and full forward modeling is needed to explore the possible redshifts, column densities, and spectral slopes. Full parameter space exploration with forward modeling also permits going to lower count spectra.

Figure 3 shows the redshift, photon index, and column density parameter space of the two examples. For R4 (mildly obscured; gray), the column density ($\log N_{\text{H}} = 22.57^{+0.04}_{-0.05} \text{ cm}^{-2}$) and redshift (XZ = 0.31 ± 0.02) of the source are both well constrained. The redshift uncertainty also agrees with specz (0.31). However, a degeneracy between redshift and column density (bottom middle panel) is visible.

The turnover shape in the model and the limited number of photons detected produces a degeneracy between redshift (decreasing the turnover energy) and column density (increasing the turnover energy). If this degeneracy is not resolved, large uncertainties remain on the redshift. However, given enough photon counts, the detection of absorption edges allows a relatively unique determination of redshift and column density. Furthermore, towards the Compton-thick limit of $N_{\text{H}} \sim 10^{24} \text{ cm}^{-2}$ the model shapes change (see Fig. 1) as the Fe K α absorption edge becomes stronger. Because of this behavior, very heavily obscured solutions (and associated low redshifts) can be excluded in the degeneracy (see gray “streak” of degeneracy in the lower middle panel of Fig. 3).

The two cases illustrate the redshift precision that can be attained. The mildly obscured source has 5690 source counts

¹ auxiliary response files (arf), redistribution matrix files (rmf), and background spectra (bkg).

² The equivalent width of the Fe K α line is below 100 eV (Turner et al. 1998), which is below the typical spectral resolution of most X-ray CCDs operating today.

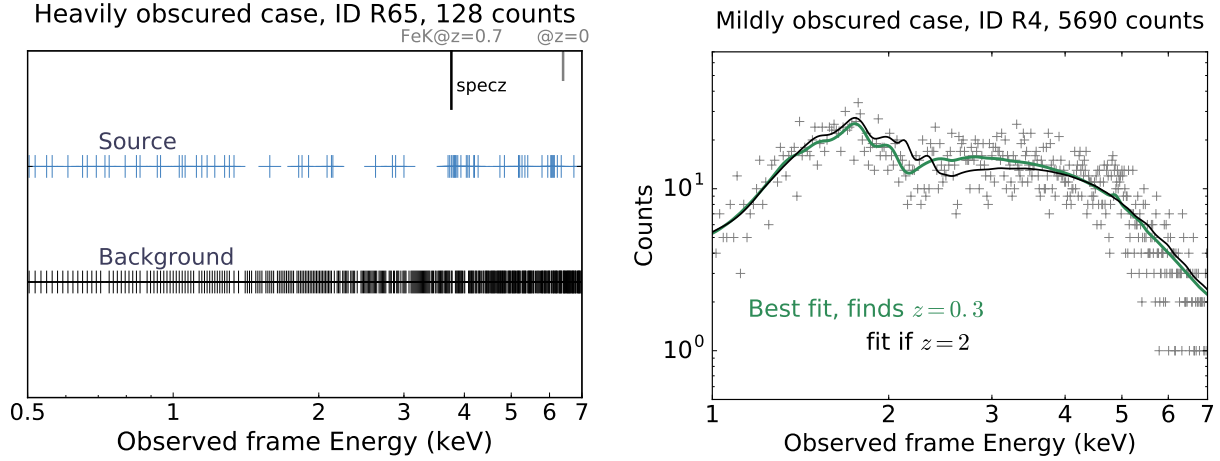


Fig. 2. Two extreme cases in which XZ provides the correct redshift. *Left panel:* unbinned X-ray source and background region counts for R65, a heavily obscured AGN ($\log N_{\text{H}} \approx 24.4 \text{ cm}^{-2}$) with a relatively low number of source counts. The counts are plotted in a horizontal line to show their energy in the observed frame. While the background count rate is nearly independent of energy, source counts clearly cluster at 4 keV. Assuming that feature corresponds to the rest-frame 6.4 keV Fe $K\alpha$ feature, we can compute the corresponding redshift $z = 6.4 \text{ keV}/E_{\text{obs}} - 1$ to the energy (*top axis*). We find $z = 0.69$, which is consistent with the optically derived spectroscopic redshift. *Right panel:* binned X-ray spectrum of source R4, a mildly obscured AGN with a high number of counts. The green curve is the best fit, yielding $\log N_{\text{H}} \approx 22.6 \text{ cm}^{-2}$ and $\Gamma \sim 1.9$ and finds the spectroscopic redshift $z = 0.31$ (see Fig. 3). The black curve shows the fit if the redshift is fixed to $z = 2$, which does not follow the data points as well.

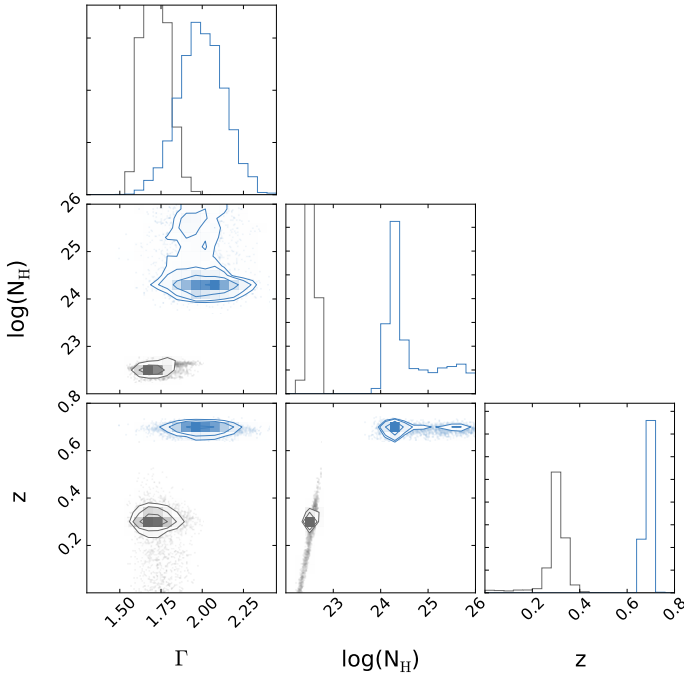


Fig. 3. Corner plot showing the parameter space over photon index, column density, and redshift for the heavily obscured AGN R65 (blue) and the mildly obscured AGN R4 (gray), see Fig. 2. Contours indicate 1, 2, and 3σ uncertainties. Any posterior samples outside are plotted as gray dots. *Bottom right panel:* shows the marginalized redshift probability distribution. In the *bottom middle panel* a modest degeneracy between column density and redshift is present for R4.

while the heavily obscured case (R65; blue) has only 128 source counts. Nonetheless, the redshift error for the latter is slightly smaller ($XZ = 0.69 \pm 0.01$), with a unique column density solution ($\log N_{\text{H}} = 24.39^{+0.02}_{-0.19} \text{ cm}^{-2}$). We note that the photon counts, source obscuration level, and source redshift all contribute to the redshift uncertainties.

3. Methodology

This section outlines our spectral model assumptions and parameter exploration algorithms used, as well as how redshift probability distributions are computed from X-ray spectra.

3.1. Spectral model assumptions

To fit the source we use the torus model from (Brightman & Nandra 2011), which models the circumnuclear material around the X-ray emitting corona and computes the spectrum with photo-electric absorption, Compton scattering, and fluorescent line emission (most importantly Fe $K\alpha$). We combine this model with a soft scattering power law with a normalization of up to 10%. We also model the galactic absorption towards the source, such that

$$\text{model} = \text{tbabs} * (\text{torus} + \text{powerlaw})$$

This is the `torus+scattering` model described in Buchner et al. (2014), which fits the CDF-S field data reasonably well. We also tested replacing the torus model with a simple photo-electric absorber model (`wabs`) and find comparable but slightly worse results for Compton-thin sources. This indicates that the Fe $K\alpha$ line is not necessary to determine redshifts in these sources.

Non-informative priors are assumed for all parameters (uniform on $\log N_{\text{H}} = 20 - 26$ and $z = 0 - 7$; log-uniform for normalizations), except for the photon index which has an informative Gaussian prior with mean 1.9 and standard deviation 0.15 (e.g., Nandra & Pounds 1994).

Solar abundances are assumed throughout. In obscured AGN, XZ primarily measures the total column density of oxygen and iron, and XZ succeeds only when the column density is high (e.g., $N_{\text{H}} > 10^{22}$). Extremely subsolar oxygen-to-iron abundance ratios could hide edges and lead to poor or invalid redshift estimates. In general, if sources do not adhere to our model assumptions we should see a large outlier fraction in our test sample.

To model the contribution of background spectra, we use a new automatic fitting technique, *BStat*. Modeling the background carefully by incorporating prior knowledge of background shapes seen in other observations also permits us to work with very few source counts. To create empirical background models with just a few parameters we use a simple machine learning method on archival observations (described in Appendix A). We carry this out for several detectors including those on *Chandra*, *XMM*, and *NuSTAR*. To reduce the number of parameters, we fit the background model to the background spectrum first and freeze its shape parameters. The source parameters and background normalization are then fitted jointly. In low-count sources, our approach extracts more information from the X-ray spectrum because the machine-learned background model exploits instrument-specific prior information.

3.2. Quantifying redshift information

Full exploration of the parameter space is crucial but not trivial. The parameter space of AGN obscuration presents multiple solutions due to the different spectral behavior in the Compton-thin and Compton-thick regimes (see Fig. 1) and due to the spectral turnover degeneracy involving redshift, column density, photon index, and source flux normalization. This presents difficulties for commonly employed minimization algorithms or Markov chain Monte Carlo techniques, and necessitates a global search algorithm (Buchner et al. 2014). We use the analysis software BXA (Buchner et al. 2014), which connects the Monte Carlo nested sampling algorithm MultiNest (Feroz et al. 2009) with the fitting environment CIAO/Sherpa (Fruscione et al. 2006). BXA yields probability distributions on all physical parameters left free in the fit, including the photon index Γ , the line-of-sight column density N_{H} , and most importantly the redshift z . These three parameters are often degenerate and their uncertainties large when the X-ray spectrum has insufficient information. To avoid overconfidence in the best fit, throughout we use the full posterior probability distribution for XZ. As we later show, the combination of XZ redshift probability distributions with those from photoz can be very powerful even in the presence of large uncertainties.

The first useful way to characterize the redshift constraints is via the mean and standard deviation of the probability distribution. If there is a single peak and the standard deviation is $\sigma(z) < 0.1$, we can say that we have constrained the redshift relatively well. However, the situation is less clear when there are multiple redshift solutions and upper or lower redshift bounds.

Thus, we define a redshift uncertainty measure that can deal with multiple redshift solutions. How much redshift information was gained from a measurement can be expressed by how different the posterior is from the (uninformative) prior. If at least some redshift intervals have zero posterior probability (i.e., they are excluded), we have learned something. We use the Kullback–Leibler definition of the information gain (IG):

$$\text{IG} = \int \text{Posterior}(z) \times \log_2 \frac{\text{Posterior}(z)}{\text{Prior}(z)} dz \text{ bits}$$

A value of $\text{IG} = 0$ bits means there was no information gained, i.e., the posterior is equal to the prior. A high value of IG means a larger amount of information gained. We define three classes and provide simplifying labels:

- (A) $\text{IG} > 2$ bits: Redshift constrained
- (B) $\text{IG} = 1\text{--}2$ bits: Some redshift information
- (C) $\text{IG} < 1$ bit: Little to no redshift information in the X-ray spectrum.

Table 1. Sample statistics.

	CDF-S	AEGIS-XD	COSMOS
X-ray hard-band detected	326	574	1016
Stars removed	5	11	6
Total ^a	321	563	1010
XZ: have redshift information	20%	14%	6%
XZ: redshift constrained	23%	11%	3%
Have photoz	100%	53%	98%
Have specz	56%	17%	49%
Obscured ($\log N_{\text{H}} \geq 22 \text{ cm}^{-2}$)	82%	34%	33%

Notes. Sources that have redshift information are those with $1 \leq \text{IG} < 2$ bits, while sources with constrained redshift are those with $\text{IG} \geq 2$ bits.

^(a) Total number of sources selected.

The borders at 1 and 2 bits crudely correspond to redshift uncertainties of $\sigma(z) = 0.85$ and 0.4, respectively, if there is only a single Gaussian solution.

4. Results

We first present some basic statistics about number and types of sources where XZ is able to extract redshift information from the X-ray data alone. Table 1 shows the sample statistics and the relative effectiveness of XZ, photoz, and specz. We find 74 (out of 321) CDF-S sources with constrained redshifts ($\text{IG} > 2$) from the X-ray spectra alone, which corresponds to $\sim 23\%$ of the sample. Table 2 shows the properties of all 138 CDF-S sources with $\text{IG} \geq 1$. The XZ method works not only for heavily obscured sources, but also in mildly obscured ones if the count statistics are good enough.

To build a better understanding of the IG quantity, we compare it to the simpler redshift error σ in Fig. 4. In a large fraction of CDF-S sources the IG has a close correspondence to $\sigma(z)$: High σ correspond to low IG. As the color-coding in Fig. 4 highlights, high-IG sources tend to be obscured. Our three IG classes are delimited by black vertical lines. Nearly all of the sources found in region A (redshift constrained, $\text{IG} \geq 2$ bits) have $\sigma \leq 1$. For region B (some redshift information, $\text{IG} = 1\text{--}2$ bits), there is large scatter in σ . The highest standard deviations here correspond to sources with multiple redshift solutions. In these cases simple error bars do not faithfully represent the uncertainties. Finally, sources in region C have little or no redshift information in their X-ray spectrum.

We test XZ against specz and photoz measurements in the following subsections. For this, we define three subsamples: The first two are composed of region A sources with constrained redshifts from XZ, which we compare to specz in subsample I (48 AGN; Sect. 4.1), and photoz in subsample II (26 AGN; Sect. 4.2). Sources with some XZ redshift information (region B, $1 \leq \text{IG} < 2$ bits) are compared to photoz in subsample III (64 AGN; Sect. 4.3). In all cases the specz and photoz information is taken from the compilation of Hsu et al. (2014).

4.1. Constrained XZ vs. specz (subsample I)

We first test the reliability of our XZ method against specz measurements.

Figure 5 plots XZ against specz for subsample I. For $\sim 60\%$, the specz resides within the 68% quantile of XZ. In many cases the uncertainties are very small, providing relatively precise redshift measurements. The inset of Fig. 5 shows the deviations between the specz and XZ median. For 75%, these are within $\pm 0.15(1+z)$ with a median absolute deviation of $0.04(1+z)$.

Table 2. Redshift catalogue of CDF-S sources with $IG \geq 1$ bit (excerpt).

R13 ID	Counts	RA	Dec	XZ	photoz	$\log_{10} N_{\text{H}}$	specz	Quality	IG (bits)	Sample
4	5690	53.1621	-27.95005	0.31 (+0.02, -0.03)	0.28 (+0.01, -0.01)	22.57 (+0.04, -0.05)	0.2983	0	4.32	I
5	1112	53.05775	-27.93325	1.89 (+1.20, -1.46)	2.88 (+1.98, -1.17)	23.00 (+0.48, -0.77)	-	-	1.27	III
6	2030	53.10342	-27.93336	0.78 (+0.48, -0.48)	1.06 (+0.10, -0.02)	22.36 (+0.26, -0.33)	1.05	2	2.41	II
7	2321	53.17007	-27.92966	3.35 (+0.05, -0.05)	2.97 (+0.08, -0.11)	23.93 (+0.05, -0.05)	1.221	2	3.32	II
13	1085	53.19605	-27.89266	1.97 (+0.93, -1.36)	2.60 (+0.16, -0.19)	22.92 (+0.32, -0.67)	-	-	1.50	III
16	622	53.14668	-27.88828	1.81 (+0.16, -1.03)	1.85 (+0.09, -0.09)	23.22 (+0.10, -0.54)	-	-	2.66	II

Notes. *Column 1:* identification number of [Rangel et al. \(2013\)](#). *Column 2:* total number of detected source counts. *Columns 3 and 4:* J2000 right ascension and declination in degrees. *Column 5:* redshift (XZ). *Column 6:* redshift (photoz) *Column 7:* neutral hydrogen column in units of cm^{-2} from XZ fit. *Column 8:* reported specz values from literature compiled by [Hsu et al. \(2014\)](#). *Column 9:* quality of specz, a value of 0 (2) represents secure (insecure) specz's. *Column 10:* information gained. *Column 11:* subsample. All quoted errors are 1σ .

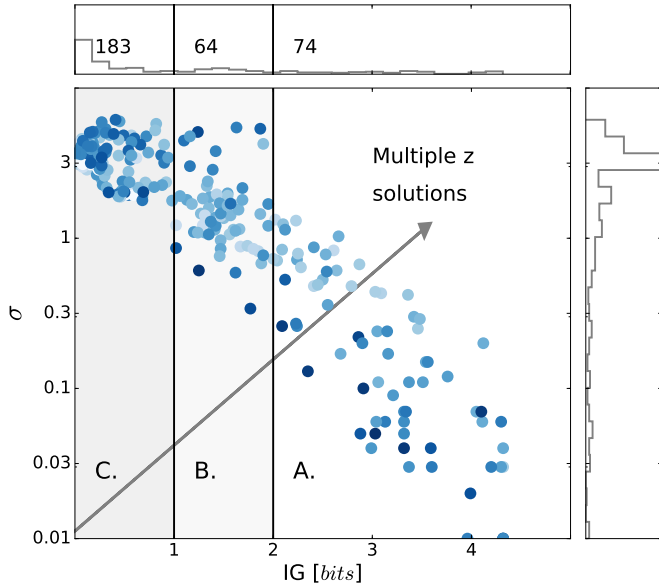


Fig. 4. Comparison of two redshift uncertainty measures, the information gain (IG; x -axis) and the standard deviation $\sigma(z)$ (y -axis). Generally, a larger $\sigma(z)$ implies a smaller IG. However, the IG also captures multiple solutions (examples are shown in subsequent figures) which can inflate the simple $\sigma(z)$ measure (symbolized by the gray arrow). The color-coding indicates the obscuration (median of uncertainties); darker colors represent higher column densities. Heavily obscured sources typically have only a single solution and are at lower σ . The vertical lines delimit three regions, A ($IG \geq 2$ bits, “constrained redshift”), B ($1 \leq IG < 2$ bits, “has redshift information”), and C ($IG < 1$ bit, “no redshift information”). The top and right panels show normalized histograms.

Thus, for the vast majority of this sample, XZ provides redshifts consistent with specz.

The ability to correctly recover specz values implies that the specz counterparts are correct and that the XZ model assumption are generally valid. We define outliers as those where the $1-\sigma$ XZ error bars do not cover the specz value and the XZ median is outside $\text{specz} \pm 0.15(1 + \text{specz})$. With this criterion, we find seven outliers among the 48 well-constrained sources (R49, R68, R156, R213, R446, R458, and R465). We discuss them in more detail in Sect. 5, where we find that, at least in some cases, the specz has to be corrected.

The dashed error bars in Fig. 5 indicate multiple-peaked redshift solutions. In these cases, error bars oversimplify the redshift probability distribution. To show the full distribution, a violin plot visualization is presented in Fig. 6. There we compare the XZ redshift probability distributions (cyan violin plots) and

the specz values (black vertical lines). Multiple solutions can be seen for several sources (e.g., R179, R181, R428).

4.2. Constrained XZ vs. photoz (subsample II)

Next, we consider sources without specz measurements, and validate XZ against photoz. This generally probes sources that are fainter in the optical/near-infrared wavelengths. The mutual validation of photoz and XZ can confirm redshifts or point out potentially problematic cases.

Figure 7 plots XZ against photoz for subsample II. While the majority of the sources show consistent redshifts ($XZ \approx \text{photoz}$) within the errors, there are six outliers out of 26 sources (using the same definition as above). This is a higher fraction than in subsample I. In order to analyze these cases in more detail, we show in Fig. 8 the normalized XZ and photoz probability distributions (gray and cyan, respectively). Disagreements are seen, for example in R115, R263, and R471. An interesting case is R227, which has two solutions in XZ, one of which overlaps with the photoz uncertainties. The combination of the constraints implies a redshift just above 2, demonstrating the synergy of the two techniques.

4.3. XZ redshift predictions (subsample III)

Finally, we consider XZ constraints that are weaker, but that can be combined with photoz constraints to distinguish between degenerate or multiple redshift solutions. Figure 9 shows the redshift probability distributions for subsample III ($IG = 1-2$ bits), as before the gray violin plots represent photoz, while the cyan plots show XZ. In several cases (e.g., R5, R185, R194, R306) X-ray spectral fitting can confirm results found via photometry and distinguish between multiple photoz redshift solutions. In one of the cases, R22, no photoz is presented because this X-ray source has an ambiguous association; for one of these opt/NIR counterparts (Hsu et al. ID 105068) the photoz is highly uncertain, while the other (61367) is constrained to $\text{photoz} = 2.11 \pm 0.06$. Unfortunately, the wide XZ error is not able in this case to provide much clarity.

5. Discussion

We now discuss differences in XZ, photoz, and specz redshift estimates.

We find a few noteworthy cases which we put into two distinct categories: (1) cases in which XZ and photoz agree, but neither agrees with specz, and (2) cases in which photoz and specz agree, but neither agrees with XZ.

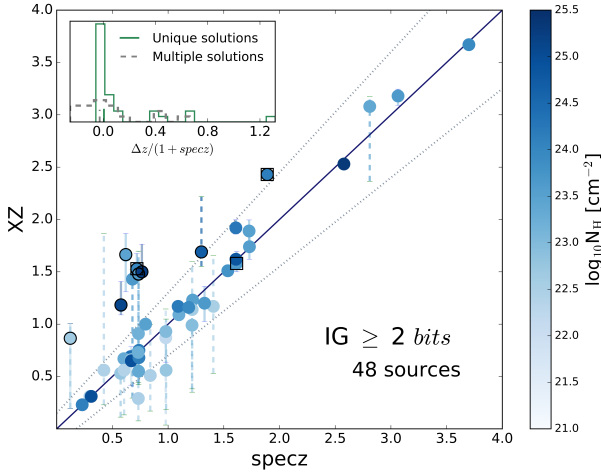


Fig. 5. Comparison of XZ and specz for constrained cases ($IG \geq 2$). The filled circles show the median of the XZ probability distribution while the error bars correspond to 1σ -equivalent quantiles and color-coded by the median of the obscuration level. Dashed error bars indicate XZ found multiple solutions. The black solid line shows $XZ = \text{specz}$, while the gray dotted lines show $XZ = \text{specz} \pm 0.15(1 + \text{specz})$. Points outside this region and where the XZ error bars do not cover the specz value are circled in black. Square symbols show insecure specz values (Luo et al. 2017). The inset shows the relative difference between XZ and specz.

5.1. Potential specz issues

Some disagreements with specz may be caused by incorrect spectral line identification, especially when only a single line has been observed. In some such cases, the redshift is nevertheless classified as “secure” in published catalogues.

We identified several cases where $XZ = \text{photoz} \neq \text{specz}$. These include sources with constrained redshifts (R49, R156, R446; from subsample I) and others with more limited redshift information ($IG > 1$ in R157, R334). For these five sources, we present the (dis)agreements between XZ, photoz, and specz in Fig. 10.

Through a literature search on these sources we found updated specz for two, R156 and R334, in Balestra et al. (2010) and Barro et al. (2014), respectively. We indicate these in Fig. 10 as vertical black lines. These new specz values agree with XZ. This makes us question the reliability of specz for the remainder of this group, in particular for R446. For R446 we do not have an updated specz; however, the redshift reported in Szokoly et al. (2004) relies on only a single spectral line. In Fig. 11 we show the evolution of the observed wavelength of two spectral lines according to redshift. The specz (0.72) was derived from a spectroscopic line assumed to be Mg $_{\text{II}}$. However, XZ and photoz indicate $z \sim 1.56$. This would be consistent if the line was C $_{\text{III}}$. Given that the literature search provided updated specz values for two sources, which agree with XZ, we caution that specz can sometimes be incorrect and this can be detected by our method.

5.2. Potential association issues

We find four constrained sources where $XZ \neq \text{photoz} = \text{specz}$, namely R68, R213, R458, and R465. For these outliers, Fig. 12 shows the redshift probability distributions, along with their associated specz. The disagreement likely indicates that some assumptions of the XZ method are invalid in these cases. From these we estimate that the outlier fraction is $\sim 8\%$ of the sample, indicating a high reliability of the XZ method.

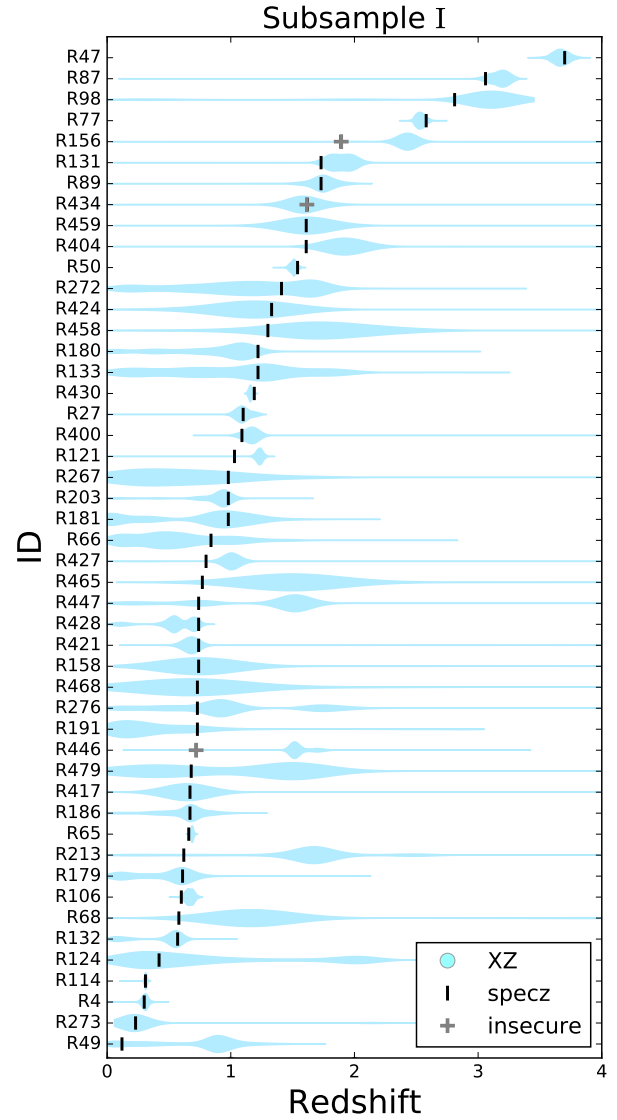


Fig. 6. XZ probability distributions (cyan violin plots) for subsample I compared to specz values (black vertical lines). Sources are ordered by increasing specz. Insecure specz values are shown as gray crosses.

For this group, in principle the XZ could be correct if both the specz and photoz are incorrect. This could occur because of incorrect optical counterpart associations. For our sample, this is not very probable because the *Chandra* positional error with approximately one hundred counts is usually not large enough to allow multiple counterpart associations.

5.3. Potential XZ issues

In six constrained cases XZ disagrees with photoz (from subsample II): R37, R56, R115, R227, R263, and R471 (see Fig. 8). All of these sources fall within the sensitivity range of the method, with either a high count number or a high level of obscuration. It is necessary to investigate these sources further to confirm the correct redshift since one (or both) of the methods fails in these cases. As discussed above, photoz (like specz) can suffer from possible counterpart misassociations, while XZ is far less likely to do so (given the X-ray source density). Both photoz and XZ can yield multiple solutions and sometimes the correct solution is outweighed by an incorrect one, leading to incorrect redshifts.

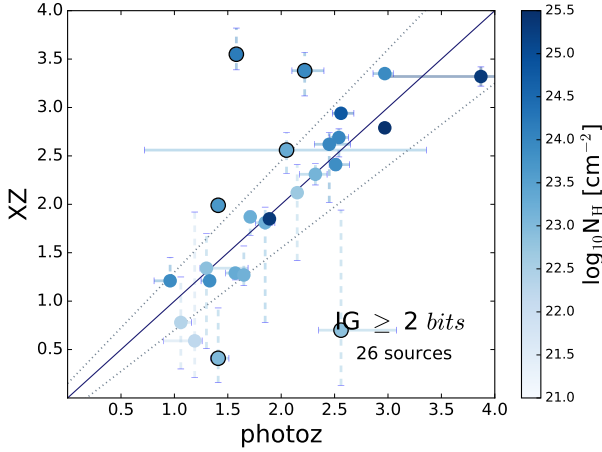


Fig. 7. As in Fig. 5, but comparing XZ to photoz (subsample II).

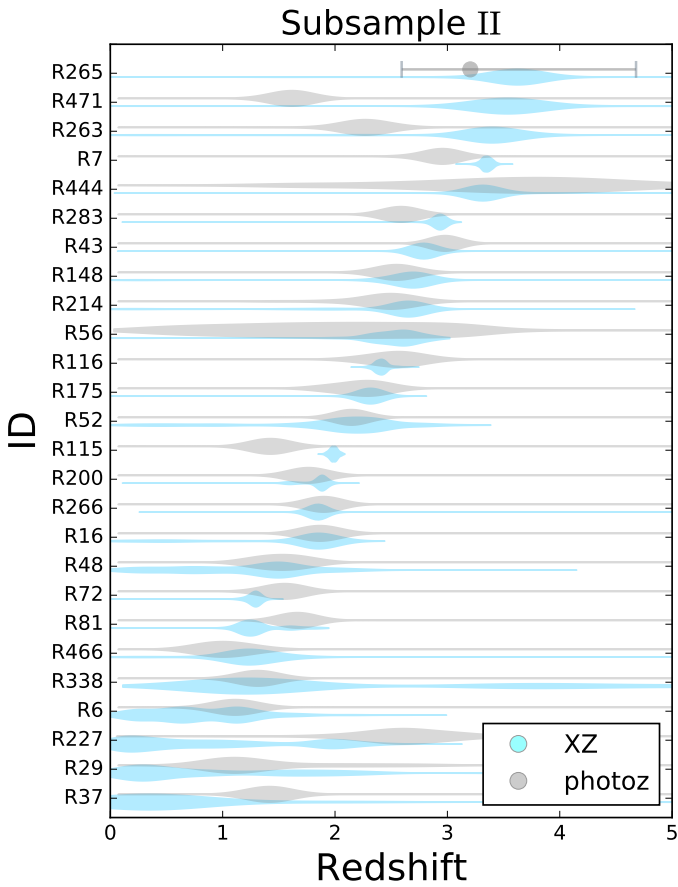


Fig. 8. XZ (cyan) and photoz (gray) probability distributions for subsample II ($IG \geq 2$ bits, no specz). XZ and photoz can mutually confirm redshifts and narrow uncertainties (e.g., R56, R116, R227, R444) or select between multiple photoz solutions (e.g., R12, R333). Some cases show disagreements (e.g., R115, R273, R471). For R265 the photoz was taken from Luo et al. (2010). Sources ordered by increasing XZ median values.

Figure 13 shows an example of the multiple solution issue, R447. While the photoz = 0.74 ± 0.01 and specz = 0.735 agree, the XZ (nominally $1.48^{+0.08}_{-0.87}$) has two solutions, with the less probable solution covering the same range as the photoz and specz. Thus, the most probable redshift solution reported by XZ is not necessarily the correct one, which is why it is advisable to combine the results with photoz in order to get more reliable results. In

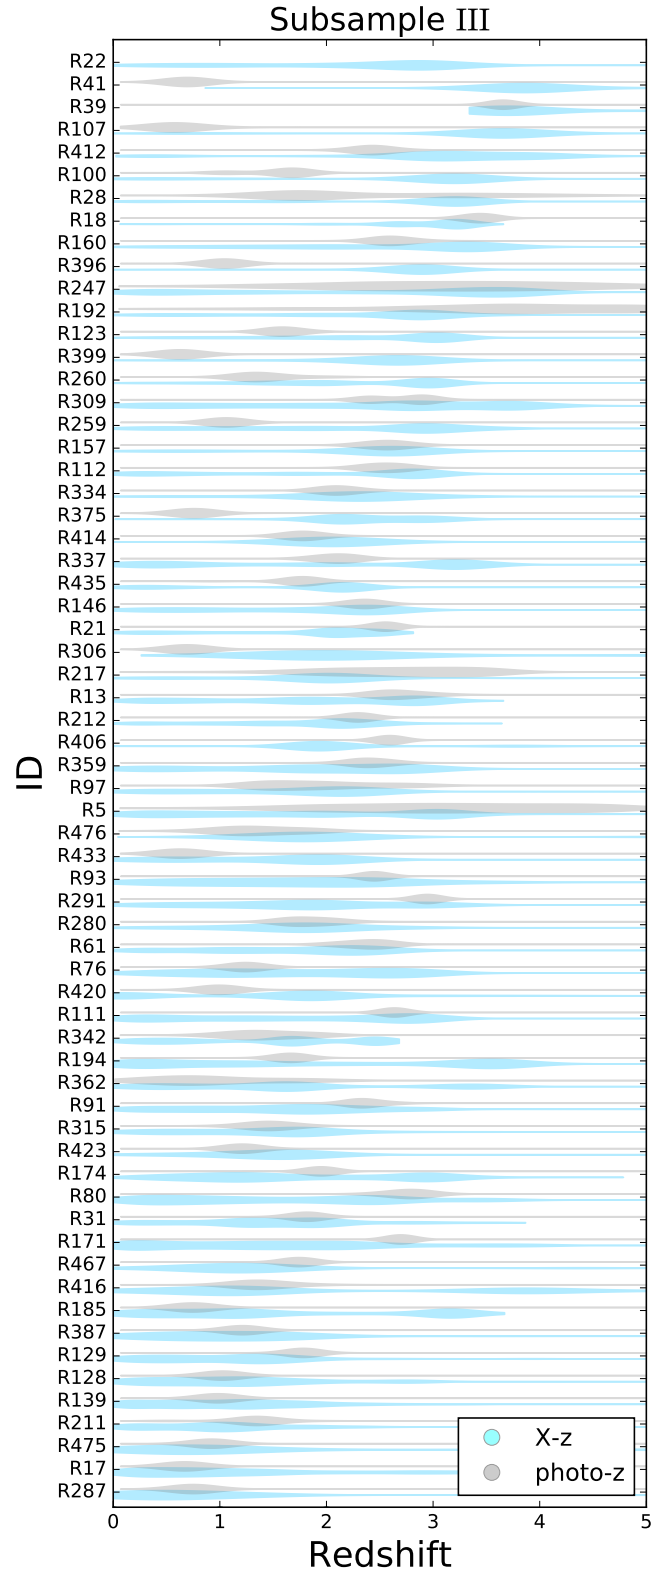


Fig. 9. XZ (cyan violin plots) and photoz (gray violin plots) probability distributions for subsample III ($1 \text{ bits} \leq IG < 2 \text{ bits}$).

addition, XZ can be affected by AGN variability. This is an open or unclear issue, as variability might either weaken or strengthen features needed for XZ depending on the sense of the variability (i.e., increasing or decreasing NH or intrinsic continuum).

Finally, there could be fundamental issues with the XZ approach. XZ is model-based and thus invalid model

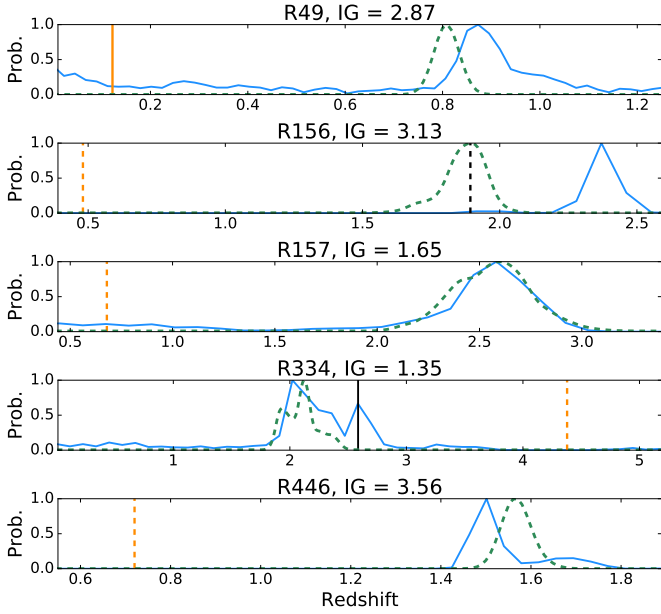


Fig. 10. Normalized redshift probability distributions where XZ (blue solid) and photoz (green dashed) agree but are different than the specz values (orange vertical lines). Black vertical lines in R156 and R334 show updated specz values from [Balestra et al. \(2010\)](#) and [Barro et al. \(2014\)](#), respectively, which agree with the photoz and XZ. Dashed vertical lines show insecure specz values.

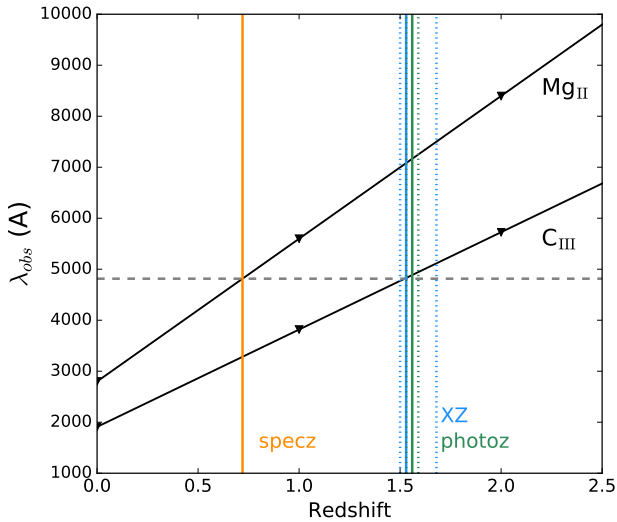


Fig. 11. Observed wavelength of Mg_{II} and C_{III} spectroscopic lines as a function of redshift. The horizontal dashed line shows the observed wavelength of a single detected optical spectral line in source R446. The vertical lines indicate the reported specz (orange), photoz (green), and XZ (blue). The dotted vertical lines show the errors of photoz and XZ, respectively. The specz was derived assuming Mg_{II} , but agrees with XZ and photoz if we assume C_{III} instead.

assumptions could lead to incorrect results. The model could be extended to be more realistic, e.g., by allowing different obscurer metallicities, obscurer geometries, multiple AGN components, and/or soft energy contaminants. However, the low outlier fraction of **8%** when validating against specz indicates that the simple adopted model is adequate in general.

XZ is most useful when obscured AGN with moderate count statistics make up a sizable fraction of the sample.

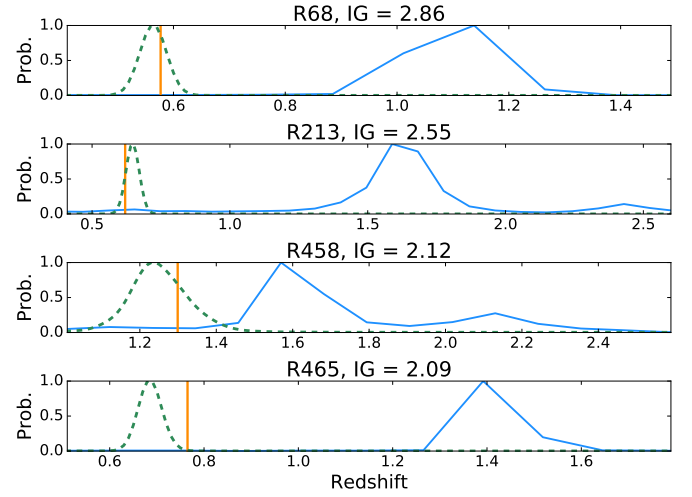


Fig. 12. Same as Fig. 10, but for cases where photoz (green dashed) and specz (orange vertical line) agree and are noticeably different than XZ (blue curve).

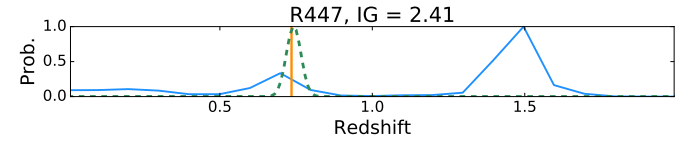


Fig. 13. Same as Fig. 10, but for case R447 where XZ (blue) presents multiple solutions, one of which agrees with photoz (green dashed) and specz (orange vertical line).

Table 3. Simulation parameter grid.

Spectral Parameter	Allowed values
Photon Index Γ	1.9
Column density N_{H}	$10^{22}, 10^{23}, 10^{24}, 10^{25} \text{ cm}^{-2}$
Viewing angle θ_{op}	45°
Opening angle θ_{view}	87° (edge-on)
Scattering fraction f_{scat}	1%
Redshift z	0.1 to 5 in steps of 0.1
Instruments	see Table 4

Notes. Source fluxes were normalized for each instrument to target reasonable source count ranges, as listed in Table 4.

Table 4. Simulation setup for tested instruments.

Instrument	Counts	N	E (keV)	Exp (ks)
Athena/WFI	10, 20, ..., 150	3000	0.2–10	200
Chandra/ACIS	10, 20, ..., 300	6000	0.5–8	4000
eROSITA	100, 200, ..., 2000	4000	0.2–10	1
NuSTAR	10, 20, ..., 300	6000	3.3–78	400
Swift/XRT	25, 50, ..., 400	3200	0.5–5	100
XMM/PN+MOS	19, 38, ..., 285	3000	0.3–8	20

Notes. *Column 1:* instrument name. *Column 2:* Grid of source net counts used. For XMM, the counts refer to the number of PN counts, but data from PN, MOS1, and MOS2 are fitted simultaneously. *Column 3:* Total number of simulations. *Column 4:* Energy range considered. *Column 5:* Exposure time.

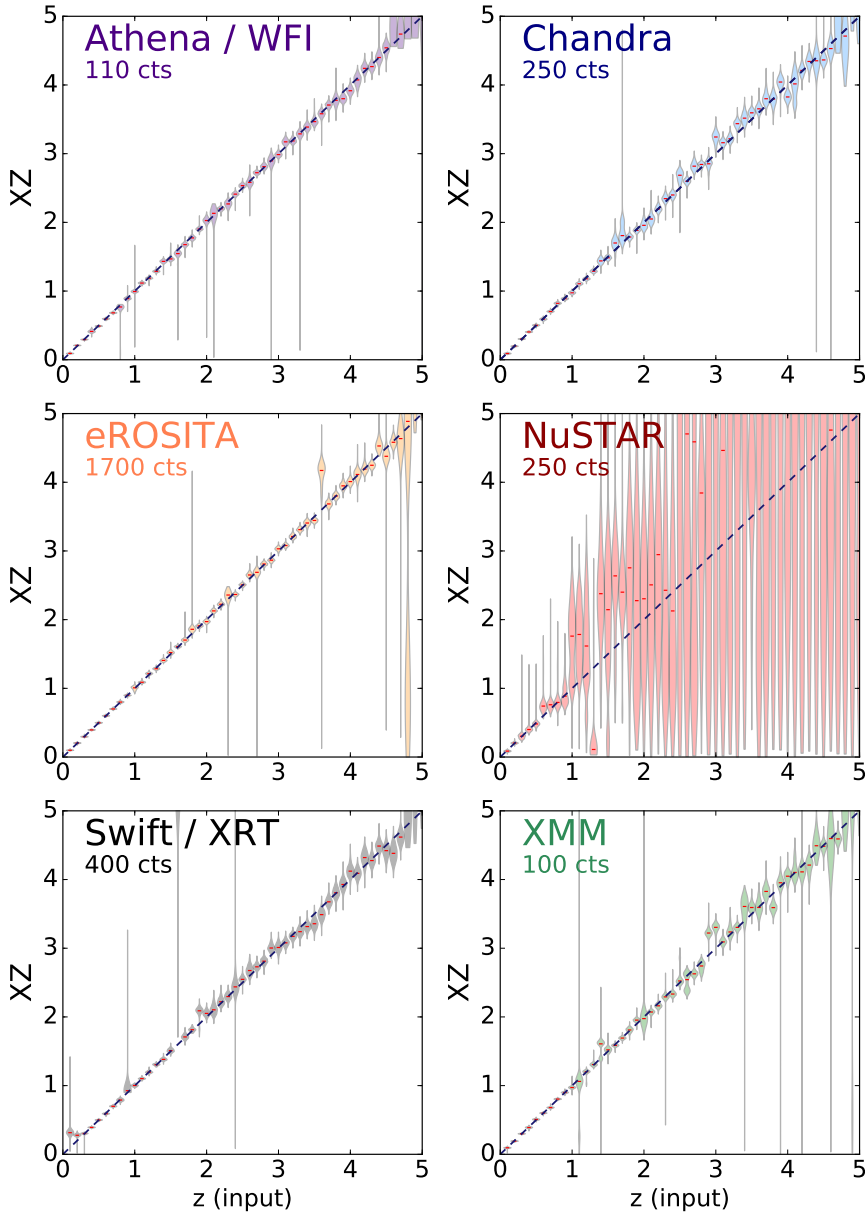


Fig. 14. Recovery of redshifts by the XZ method. Output XZ redshift probability distributions are shown as violin plots for each input z (x -axis). In all cases $N_{\text{H}} = 10^{24} \text{ cm}^{-2}$ is shown. The panels correspond to different instruments, with a total number of source counts chosen at the approximate threshold where the XZ method starts to work optimally.

Because of the observed increase in the obscured fraction at low luminosities and high redshifts (e.g., Ueda et al. 2003; Hasinger et al. 2005), XZ is most useful in deep surveys like the CDF-S. Consistent with this, we find lower percentages of constrained redshifts when testing our method on data from the COSMOS and AEGIS-XD fields (see Table 1). The detailed results from these two surveys are presented in Appendix B, including redshift catalogues and photoz/specz plots.

6. Sensitivity analysis for different X-ray Instruments

6.1. Simulation setup

To answer how many source photon counts are needed for good redshift constraints, we performed extensive realistic simulations for multiple instruments. We generate source and background spectra and fitted them as with the real observations. The recovery is tested across different redshift ($z = 0 - 5$) and

column density regimes ($N_{\text{H}} = 10^{22-25} \text{ cm}^{-2}$). The simulation parameter grid is listed in Table 3.

For each such scenario (z , N_{H}), several instruments are considered, including *Chandra*, *XMM-Newton*, *Swift*, *NuSTAR*, *eROSITA*, and *Athena*. Importantly, we vary the expected number of source counts to find the relation between number of counts and XZ redshift uncertainty. We also tested different background levels by varying the exposure time, but find that these show small differences (also thanks to the background model fitting), and that N_{H} and the number of counts remain the most important variables. For *Chandra*, for example, we adopted a relatively long exposure time representative of the CDF-S, which gives the worst case of high background levels. For each instrument, Table 4 shows the counts, the exposure time, and the total number of simulations created and fitted.

Simulations were created in *Sherpa* using the `fake_pha` command. We used the same spectral model for simulation as in the fitting. To simulate background spectra, we let `fake_pha` re-sample a representative `bkg` file for each instrument. For upcoming missions *Athena* and *eROSITA* we simulated a background

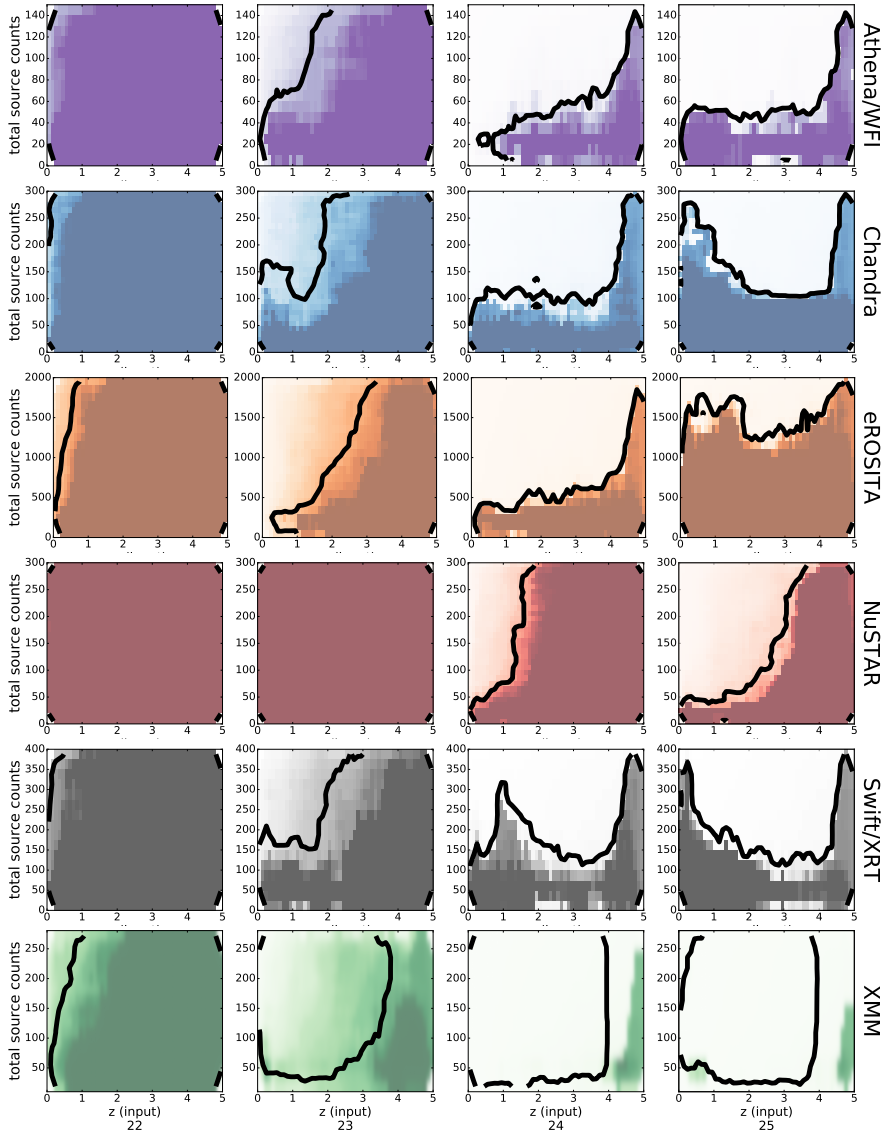


Fig. 15. Sensitivity heat maps for each instrument. Redshift uncertainty (colors) are plotted against source redshift (x -axis) and total source counts (y -axis). Above the thick contour lines, $IG > 1$ bit, i.e., the number of counts is sufficient to provide some redshift information. From *left to right panels*, the columns show $\log N_{\text{H}} = 22, 23, 24,$ and 25 cm^{-2} . XMM total source counts refer to total counts across PN/MOS1+2.

following the information provided by the respective mission websites³.

6.2. Redshift recovery

The recovery of input redshifts is presented in Fig. 14. The XZ uncertainties are presented for various instruments and redshifts. Here, a total source count number was chosen for each instrument representing the approximate threshold where the XZ method starts to be useful. Figure 14 demonstrates the accuracy and precision of XZ. The correct retrieval of the redshift (y -axis) at each input redshift (x -axis) is expected because the model used to simulate data is the same as the fitted model. While optimistic, Fig. 14 demonstrates that the parameter space is reliably explored. It also shows that when the X-ray spectrum contains no redshift information, the XZ uncertainties increase dramatically, but the results are not incorrect. For instance, at $z > 1$ any discriminatory information from the Fe $K\alpha$ line is redshifted out

of the *NuSTAR* observing window, and the uncertainties increase drastically.

6.3. Minimal counts for XZ constraints

We now investigate the minimum number of counts needed to obtain redshift information ($IG \geq 1$). Figure 15 shows redshift uncertainty (darker: larger errors, white: tightly constrained) as a function of total source counts (y -axis), input source redshift (x -axis), instrument (rows), and column density (columns). The thick black contour line shows $IG = 1$ bit, i.e., above it (higher counts), some redshift information can be extracted by XZ from the X-ray data alone. Considering *Chandra* for instance, we find that the sensitivity of the XZ method depends on both the count number and the neutral hydrogen column density value along the line of sight. In general, the best results are obtained with heavily obscured column densities. The required count number varies from instrument to instrument, from 50 to 500 minimal counts. Compared to *Chandra* the number of counts required is higher with eROSITA and *Swift* (both are less sensitive at $>4 \text{ keV}$). In *NuSTAR*, only Compton-thick AGN are usable because the Fe $K\alpha$ line is the primary source of information.

³ http://www.mpe.mpg.de/1341347/eROSITA_background_v12.pdf for eROSITA, <https://www.cosmos.esa.int/web/athena/resources-by-esa> for Athena

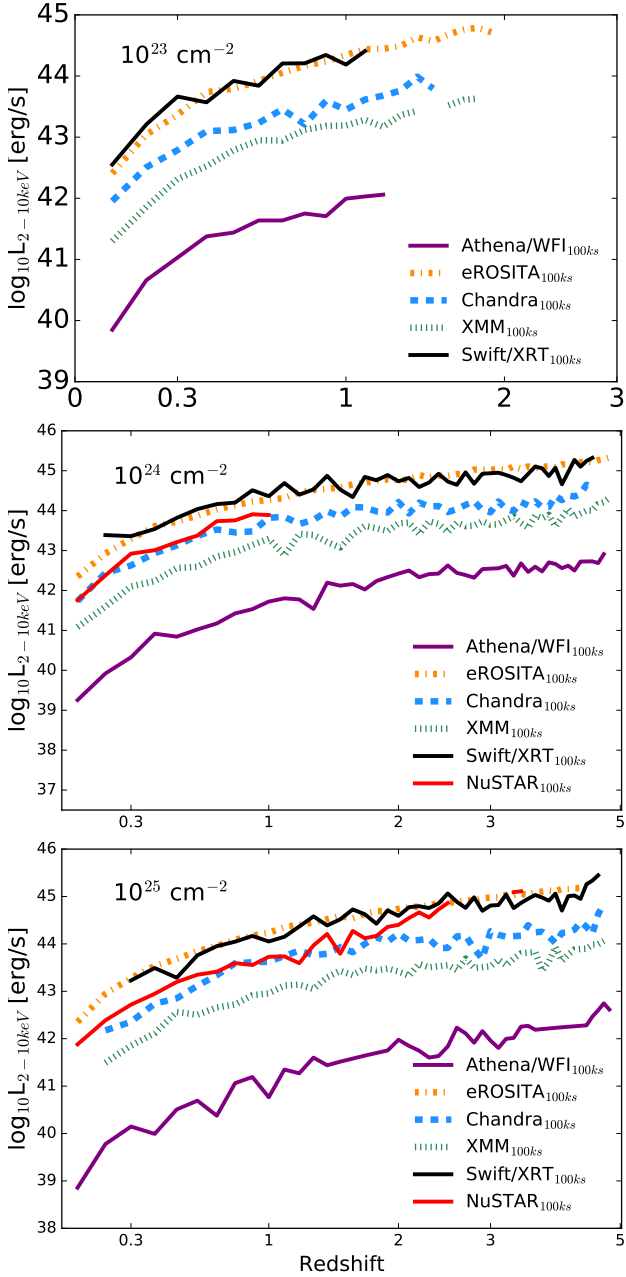


Fig. 16. 2–10 keV luminosity (observed-frame, absorbed) vs. redshift plot that shows the sensitivity limit of the XZ method for $N_{\text{H}} = 10^{23}$ (top panel), 10^{24} (middle panel) and 10^{25} (bottom panel) cm^{-2} . In order to compare the response of XZ for different instruments, we normalize the luminosity values to 100 ks. Both current and future missions are shown; XZ is particularly sensitive to the upcoming mission *Athena*.

The required number of counts imply minimal exposure times and luminosities for the different instruments. For *Chandra*, for example, we assume a relatively long exposure time (see Table 4), which gives the worst case of high background levels. We tested different exposure times, but find that these show small differences (also thanks to the background model fitting), and that N_{H} and number of counts remain the most important variables. To compare the efficiency of the XZ technique for each instrument, we assume 100 ks exposures and determine the luminosity limit that can be probed. Figure 16 shows this limit as a function redshift for each instrument. The panels correspond to different column densities $N_{\text{H}} = 10^{23-25} \text{cm}^{-2}$. At higher

luminosities, *Swift* and eROSITA are found to be similarly sensitive. However, the eROSITA all-sky survey (eRASS) will only reach exposures of several ks (Merloni et al. 2012), thus its luminosity limit will be substantially higher. At these fluxes, only a few AGN are expected across the full sky (Kolodzig et al. 2013), and have likely already been identified, for example in the WISE all-sky survey (Wright et al. 2010).

We have also considered applying the XZ method to gamma-ray bursts, which frequently show Compton-thin obscuration up to $N_{\text{H}} = 10^{23} \text{cm}^{-2}$ (e.g., Buchner et al. 2017). *Swift* (Gehrels 2004) detects gamma-ray bursts and observes X-ray spectra with XRT. Obtaining redshift constraints is crucial in order to probe massive star formation in the early Universe. Today, this requires timely follow-up and favorable observing conditions. XZ could improve this process by augmenting photometric redshifts. We tested XZ on a large sample of X-ray afterglow spectra from the *Swift* archive (described in Buchner et al. 2017). However, the emission in these sources is generally too short to accumulate the necessary counts for XZ to provide constraints with *Swift/XRT*. For some ultra-long gamma-ray bursts the counts are sufficient, but these sources present unusual spectral features that are poorly understood and likely arise from spectral variability (Piro et al. 2014; Evans et al. 2014). In any case, these sources are a small subpopulation of all gamma-ray bursts. An observatory with higher collecting area is required to obtain XZ constraints for gamma-ray bursts.

The best performance is seen with the next-generation *Athena* mission (Nandra et al. 2013), which will feature a mirror with large effective area (2m^2), wide energy coverage (0.3–12 keV) and good resolution ($<150 \text{eV}$) even for the Wide-field Imager (WFI; Meidinger et al. 2017). With this mission, surveys as deep as the CDF-S 4Ms can be performed twenty times faster (Rau et al. 2016), and cover larger areas ($\sim 100 \text{deg}^2$), allowing a deep and complete census of the AGN population at high redshift. Previously, Castelló-Mor et al. (2011) developed a redshift extraction technique based on finding the Fe $K\alpha$ line in Fourier-transformed spectra. In 100 ks exposures, this succeeds in unobscured AGN with luminosities of $10^{44} \text{erg s}^{-1}$ at $z = 1$. In Fig. 16 we estimate that XZ can measure redshift of obscured AGN, which constitute the bulk of the population in deep observations, down to luminosities as low as $10^{41} \text{erg s}^{-1}$. This indicates that for a substantial fraction of sources detected in such a survey, expensive follow-up may be unnecessary, as redshift information can be extracted directly from *Athena/WFI* X-ray spectra.

7. Conclusions

Deep X-ray fields probe accretion in the early Universe, and are dominated by obscured AGN. The deep observations require time-consuming follow-up campaigns to obtain reliable redshift estimates. In this regime we demonstrated that for a sizable fraction of sources, the X-ray spectrum contains substantial redshift information in the absorption edges and, if Compton-thick, the Fe $K\alpha$ fluorescent line. Through simulations, we characterized in detail the obscuration level and number of counts required to find constrained and reliable solutions. However, failure to meet such conditions does not lead to incorrect results, only to large uncertainties.

Sidestepping association problems, our XZ method is promising for estimating redshifts or confirming/reinforcing photoz and specz estimates. The XZ estimate is easy to apply and does not require additional data. In addition to direct redshift estimates in some high-count sources, it can discriminate

between multiple photoz solutions in a large fraction of obscured AGN in today's deep surveys. We demonstrate that disagreements of the independent XZ estimate with photoz/specz estimates can help to flag problematic measurements, for example due to incorrectly identified spectroscopic lines, to contaminated photometric data due to blending or variability, or to ambiguous counterparts. With high spectral resolution and a large collecting area, this technique will be highly effective for *Athena*+*WFI* observations.

Acknowledgements. We thank the referee for helpful comments that have improved this paper. We thank Laurent Douchy for creating Docker container images. We also thank Alex Markowitz and Claudio Ricci for providing example background spectra for XRTE, *Suzaku*, and *Swift*/XRT. We acknowledge support from the CONICYT-Chile grants Basal-CATA PFB-06/2007 (JB, FEB), FONDECYT Regular 1141218 (FEB), FONDECYT Postdoctorados 3160439 (JB), EMBIGGEN Anillo ACT1101 (FEB), and the Ministry of Economy, Development, and Tourism's Millennium Science Initiative through grant IC120009, awarded to The Millennium Institute of Astrophysics, MAS (JB, FEB). This research was supported by the DFG cluster of excellence Origin and Structure of the Universe.

References

- Akaike, H. 1974, *IEEE Trans. Autom. Control*, **19**, 716
- Balestra, I., Mainieri, V., Popesso, P., et al. 2010, *A&A*, **512**, A12
- Barro, G., Trump, J. R., Koo, D. C., et al. 2014, *ApJ*, **795**, 145
- Baum, W. A. 1962, in *Problems of Extra-Galactic Research*, ed. G. C. McVittie, *IAU Symp.*, **15**, 390
- Brightman, M., & Nandra, K. 2011, *MNRAS*, **413**, 1206
- Brightman, M., Silverman, J. D., Mainieri, V., et al. 2013, *MNRAS*, **433**, 2485
- Brightman, M., Nandra, K., Salvato, M., et al. 2014, *MNRAS*, **443**, 1999
- Broos, P. S., Townsley, L. K., Feigelson, E. D., et al. 2010, *ApJ*, **714**, 1582
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2015, *ApJ*, **802**, 89
- Buchner, J., Schulze, S., & Bauer, F. E. 2017, *MNRAS*, **464**, 4545
- Castelló-Mor, N., Barcons, X., & Ballo, L. 2011, *Adv. Space Res.*, **48**, 1304
- Civano, F., Elvis, M., Brusa, M., et al. 2012, *ApJS*, **201**, 30
- Elvis, M., Civano, F., Vignali, C., et al. 2009, *ApJS*, **184**, 158
- Evans, P. A., Willingale, R., Osborne, J. P., et al. 2014, *MNRAS*, **444**, 250
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, in *SPIE Conf. Ser.*, **6270**, 62701V
- Gehrels, N. 2004, in *Gamma-Ray Bursts: 30 Years of Discovery*, eds. E. Fenimore, & M. Galassi, *AIP Conf. Ser.*, **727**, 637
- Hasinger, G., Miyaji, T., & Schmidt, M. 2005, *A&A*, **441**, 417
- Hsu, L.-T., Salvato, M., Nandra, K., et al. 2014, *ApJ*, **796**, 60
- Kolodzig, A., Gilfanov, M., Sunyaev, R., Sazonov, S., & Brusa, M. 2013, *A&A*, **558**, A89
- Koo, D. C. 1985, *AJ*, **90**, 418
- Laird, E. S., Nandra, K., Georgakakis, A., et al. 2009, *ApJS*, **180**, 102
- Luo, B., Brandt, W. N., Xue, Y. Q., et al. 2010, *ApJS*, **187**, 560
- Luo, B., Brandt, W. N., Xue, Y. Q., et al. 2017, *ApJS*, **228**, 2
- Meidinger, N., Barbera, M., Emberger, V., et al. 2017, *SPIE Conf. Ser.*, **10397**, 103970V
- Merloni, A., Predehl, P., Becker, W., et al. 2012, ArXiv e-prints [arXiv:1209.3114]
- Nandra, K., Barret, D., Barcons, X., et al. 2013, ArXiv e-prints [arXiv:1306.2307]
- Nandra, K., Laird, E. S., Aird, J. A., et al. 2015, *ApJS*, **220**, 10
- Nandra, K., & Pounds, K. A. 1994, *MNRAS*, **268**, 405
- Pello, R., Miralles, J. M., Le Borgne, J.-F., et al. 1996, *A&A*, **314**, 73
- Piro, L., Troja, E., Gendre, B., et al. 2014, *ApJ*, **790**, L15
- Ptak, A., & Griffiths, R. 2003, in *Astronomical Data Analysis Software and Systems XII*, eds. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook, *ASP Conf. Ser.*, **295**, 465
- Rangel, C., Nandra, K., Laird, E. S., & Orange, P. 2013, *MNRAS*, **428**, 3089
- Rau, A., Nandra, K., Aird, J., et al. 2016, in *Proc. SPIE*, **9905**, 99052B
- Ricci, C., Trakhtenbrot, B., Koss, M. J., et al. 2017, *ApJS*, **233**, 17
- Rivers, E., Markowitz, A., & Rothschild, R. 2013, *ApJ*, **772**, 114
- Salvato, M., Hasinger, G., Ilbert, O., et al. 2009, *ApJ*, **690**, 1250
- Salvato, M., Buchner, J., Budavári, T., et al. 2018, *MNRAS*, **473**, 4937
- Simm, T., Saglia, R., Salvato, M., et al. 2015, *A&A*, **584**, A106
- Szokoly, G. P., Bergeron, J., Hasinger, G., et al. 2004, *ApJS*, **155**, 271
- Turner, T. J., George, I. M., Nandra, K., & Mushotzky, R. F. 1998, *ApJ*, **493**, 91
- Ueda, Y., Akiyama, M., Ohta, K., & Miyaji, T. 2003, *ApJ*, **598**, 886
- Willis, J. P., Pacaud, F., Valtchanov, I., et al. 2005, *MNRAS*, **363**, 675
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868
- Xue, Y. Q., Luo, B., Brandt, W. N., et al. 2011, *ApJS*, **195**, 10

Appendix A: Automatic background fitting

This section describes a new technique, BStat, for characterizing spectra of background regions. With principal component analysis (PCA) this technique learns the typical background shape and variations from a large sample. The technique is applied for each telescope and detector separately and forms a general solution to empirical background modeling.

To obtain training data, we download background spectra for *Chandra* ACIS and *XMM-Newton* PN&MOS from the Xassist website⁴ (Ptak & Griffiths 2003). We extracted *NuSTAR* point source spectra with corresponding backgrounds at the locations of the *Chandra* hard-band detected sources in the COSMOS, CDF-S, and AEGIS-XD surveys handled in this paper using the standard *NuSTAR* pipeline. Most of these sources are undetected. We further obtained *Swift* XRT background spectra⁵ (Buchner et al. 2017) and the *Swift*/BAT survey (Ricci et al. 2017). This yielded a sample of 1895 (*Chandra*/ACIS), 575 (*NuSTAR*), 414 (*Swift*/XRT), 1859 (*XMM*/PN), and 1982 (*XMM*/MOS) background spectra.

Principal component analysis (PCA) identifies the eigenvectors (in this case, background components) as primary axes of variance (background diversity). We stacked the background spectra so that each background spectrum has at least 10 000 total counts. This avoids non-Gaussianity in the Poisson counts. On this sample we apply PCA and extract the mean and the 20 most important features (background spectra vectors). Figures A.1 (*Chandra*), A.2 and A.3 (*XMM* MOS & PN), A.5 (*Swift*/XRT), and A.4 (*NuSTAR*) illustrate some examples of these components.

We then use in *Sherpa* these components to fit the background. Our fitting procedure starts by having all component normalizations free and performing a normal simplex fit to the background spectrum. Then, the least important component is disabled (normalization set to zero) and another fit is done. This is repeated until only the mean background component remains. For each of these fits we note the Akaike Information criterion (AIC; in this case the C statistic plus twice the number of fit parameters; Akaike 1974; Buchner et al. 2014), and choose the fit with the highest AIC. For fit robustness against local minima, we try to increase the number of components again from there until the AIC does not increase further. Choosing the fit with the highest AIC selects only the K most important

components where the background spectrum has enough information to justify activating these components, while the remainder of the spectrum is aligned with the typical (based on the training sample) background spectrum for this detector. Using the AIC thus avoids over-fitting.

Throughout the study we work with background counts without convolving through the response as some background components originate on the detectors and this is an empirical model. Working in linear space (on the counts) would be ideal because some background components are additive. However, this can lead to negative backgrounds when we fit. Instead, we transform into logarithmic space with $\log(\text{counts} + 1)$. This has the benefit that all possible combinations of background spectral features are positive counts because in linear space the (positive) spectral components are multiplied together. However, a drawback is that Gaussian lines of varying strength can not be characterized well. We address this by adding Gaussian lines on top when needed (i.e., as long as justified by an increase in the AIC). Overall this fitting procedure works robustly and captures the background spectrum well.

When fitting a science target we add the background model to the source model. We scale the background contribution proportionally to the source and background region sizes. In the source fitting we keep the background shape fixed. However, we allow the background normalization to vary freely and simultaneously fit the source and background spectra. Freezing the background normalization as well sometimes leads to overly small uncertainties, in particular when few background counts have been detected.

Our technique decreases the fit uncertainties compared to *WStat*, the *xspec* default behavior. *WStat* does per-bin estimates of the background, i.e., does not assume any continuity between spectral bins. Additionally, zero or low-count count bins produce biases (e.g., Willis et al. 2005) and require rebinning. In contrast, our technique always works with continuous background spectral models across bins, and can handle even zero-count bins. Both in our approach and when using custom-built background models, source fit uncertainties are smaller than in *WStat* because prior knowledge about typical background shapes is exploited. Our machine learning approach, however, does this without requiring the user to hand-craft a model. Our automatic background fitting technique is released as part of BXA (Buchner et al. 2014).

⁴ <http://xassist.pha.jhu.edu/>

⁵ <http://www.swift.ac.uk/>

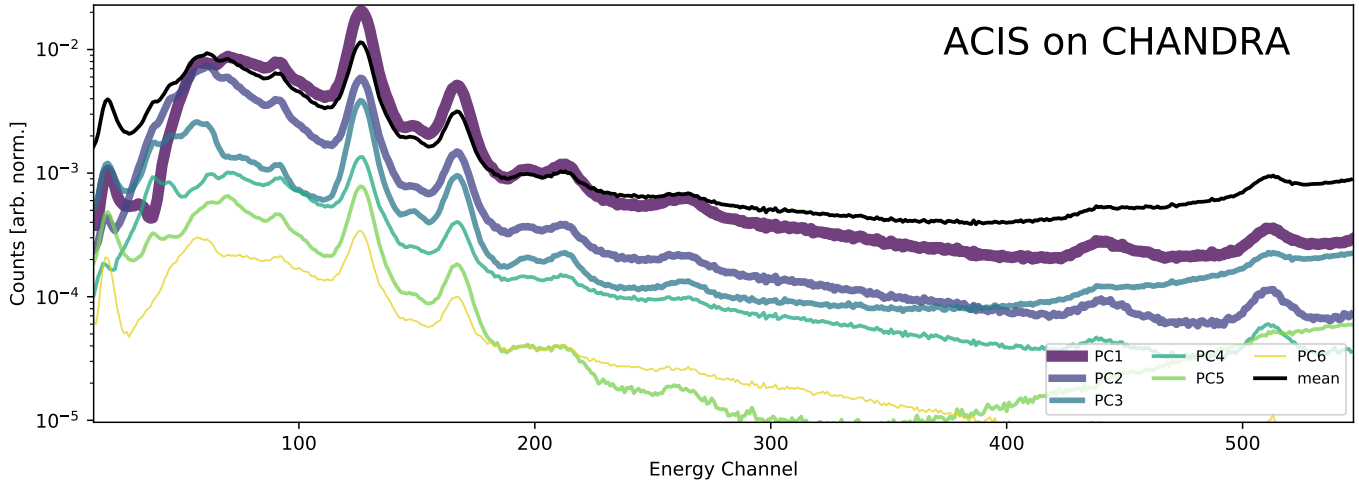


Fig. A.1. *Chandra*/ACIS background PCA components. The average spectrum plus the selected principle component is plotted in arbitrary units. All CCDs have been combined as they show broadly similar background shapes.

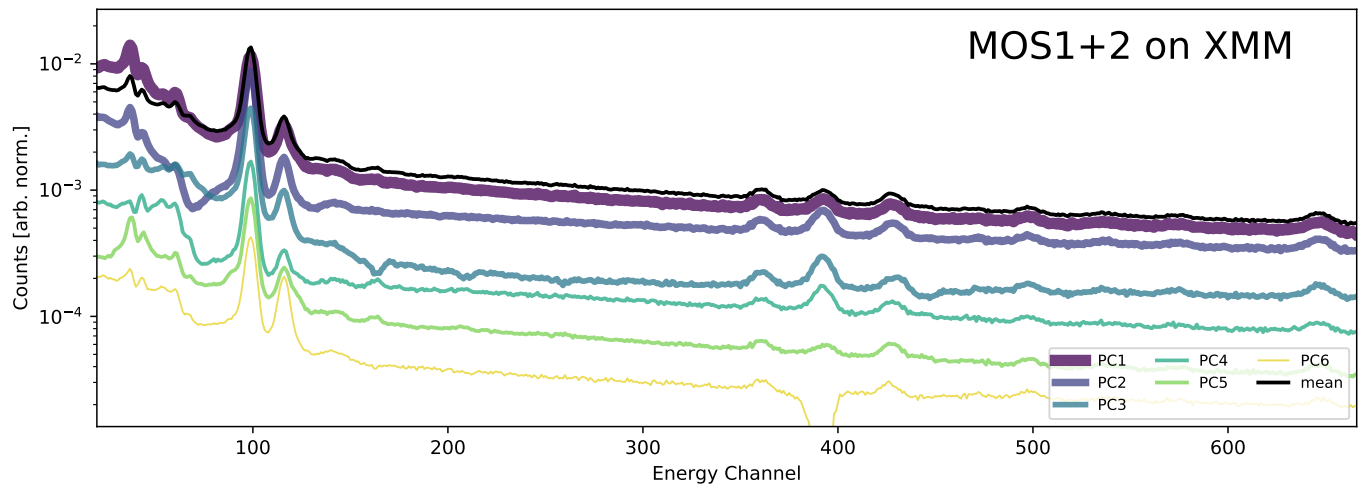


Fig. A.2. XMM/MOS background PCA components. MOS1 and MOS2 have been combined as they show similar background shapes.

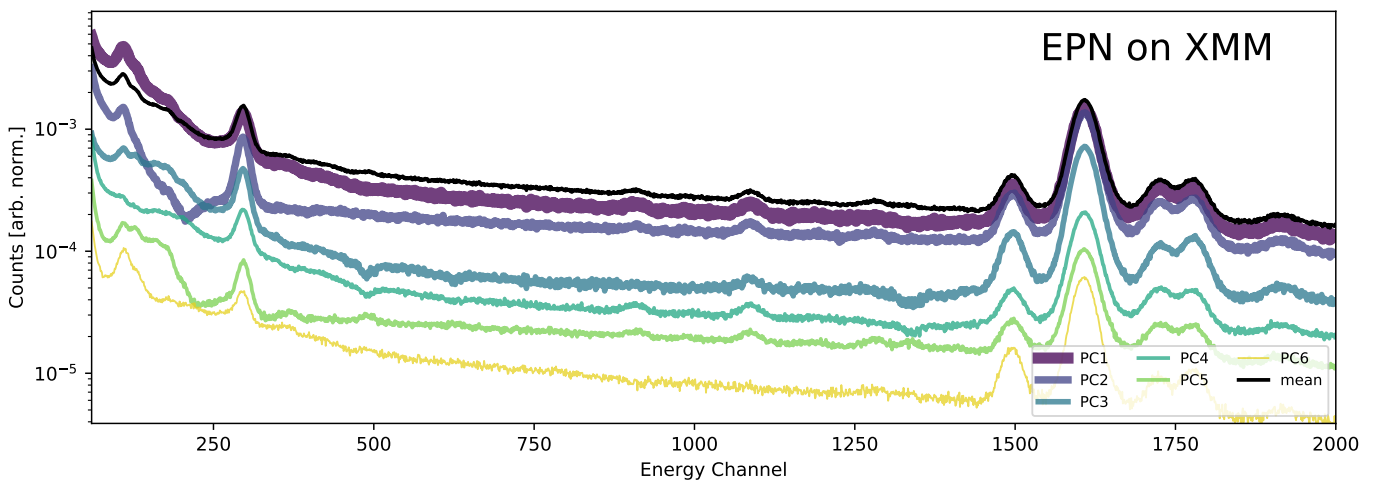


Fig. A.3. XMM/PN background PCA components.

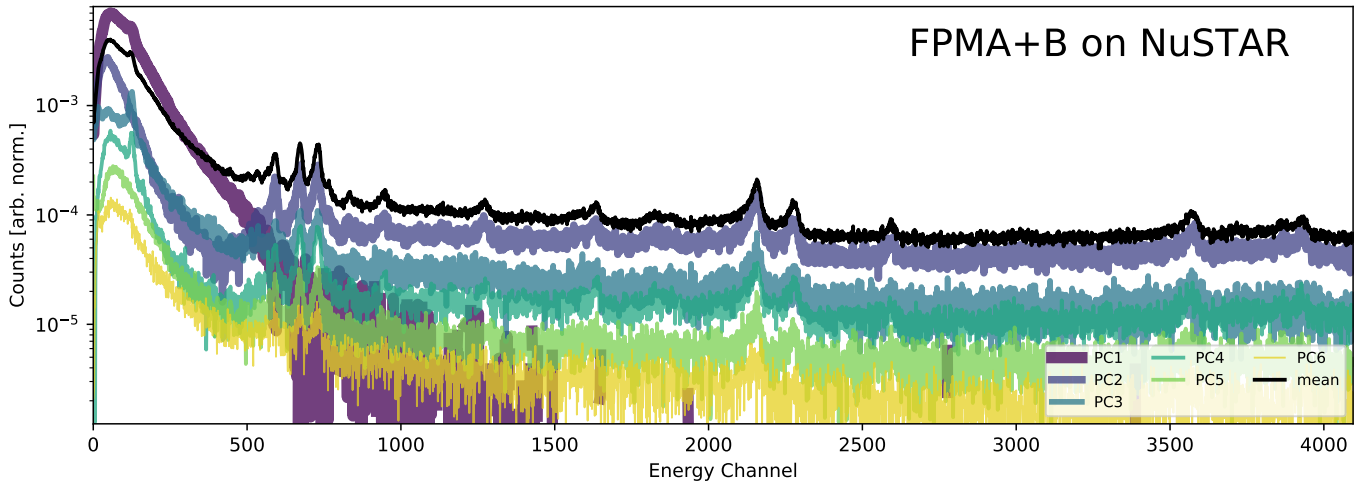


Fig. A.4. *NuSTAR* background PCA components. The FPMA and FPMB detectors have been combined as they show similar background shapes.

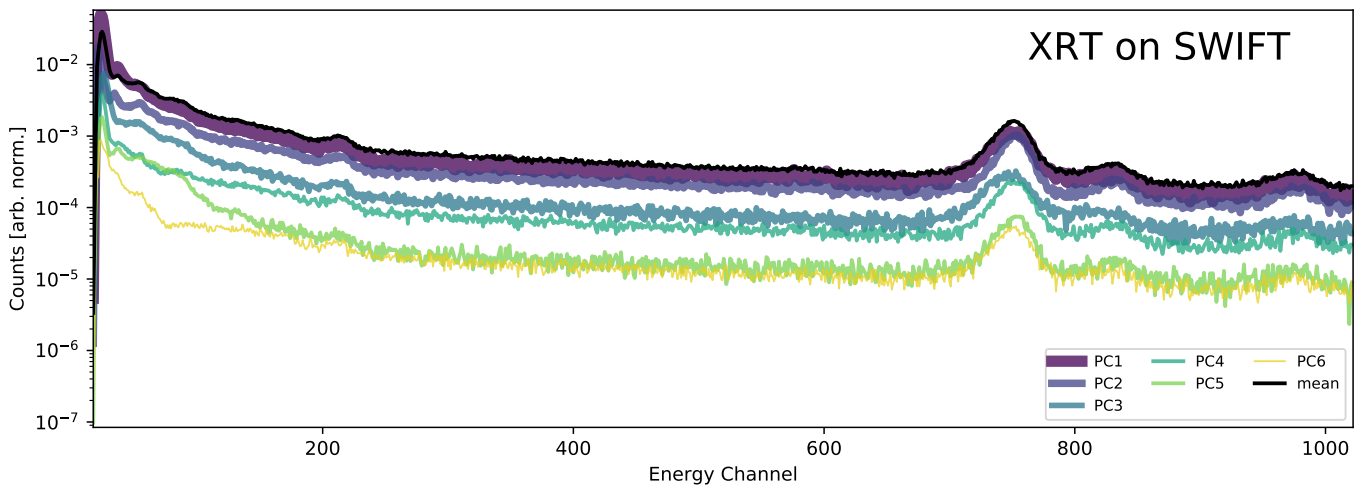


Fig. A.5. *Swift*/XRT background PCA components. WT and PC mode observations have been combined as they show similar background shapes.

Table B.1. Properties of COSMOS Sources with IG ≥ 1 bit.

ID	Counts	RA	Dec	XZ	Photoz	$\log_{10}N_{\text{H}}$	specz	IG (bits)
342	1128	149.7008	2.40256	1.68 (+2.09, -1.31)	0.39 (+0.75, -0.04)	22.71 (+0.70, -0.72)	0.375	1.43
227	1016	150.24247	1.76643	2.42 (+0.86, -1.58)	0.66 (+0.48, -0.04)	22.67 (+0.30, -0.67)	0.623	1.37
437	254	149.99335	2.37723	0.46 (+0.58, -0.31)	0.93 (+0.21, -0.05)	22.94 (+0.40, -0.27)	0.931	2.84
393	331	150.47461	2.09414	1.01 (+0.66, -0.54)	1.22 (+0.16, -0.12)	22.89 (+0.34, -0.35)	-99.0	2.21
726	248	150.2325	1.94266	0.19 (+0.02, -0.08)	0.18 (+0.96, -0.05)	23.46 (+0.08, -0.08)	0.219	4.32
386	425	150.10557	2.12631	3.17 (+0.25, -2.10)	1.03 (+0.12, -0.04)	23.42 (+0.24, -0.97)	-99.0	1.33
664	211	150.01203	2.52562	1.01 (+0.40, -0.64)	0.87 (+0.27, -0.05)	23.15 (+0.24, -0.45)	1.028	2.22
376	201	149.88424	2.33814	0.41 (+0.70, -0.29)	1.03 (+0.11, -0.08)	22.89 (+0.49, -0.25)	1.022	2.64
336	279	149.834	2.30884	1.21 (+2.99, -0.02)	1.1 (+0.09, -0.04)	25.11 (+0.63, -2.16)	1.162	2.72
459	163	149.87645	2.56878	1.38 (+0.26, -0.95)	1.65 (+0.13, -0.13)	23.03 (+0.19, -0.54)	-99.0	1.94

Notes. *Column 1:* identification number of [Elvis et al. \(2009\)](#). *Column 2:* total amount of detected source counts. *Columns 3 and 4:* J2000 right ascension and declination in degrees. *Column 5:* redshift (XZ). *Column 6:* redshift (photoz) with 1σ errors from [Salvato et al. \(2009\)](#). *Column 7:* neutral hydrogen column in units of cm^{-2} from XZ fit. *Column 8:* reported specz values from [Civano et al. \(2012\)](#). *Column 9:* information gained in units of bits. All quoted errors are 1σ .

Table B.2. Properties of AEGIS-XD Sources with IG ≥ 1 bit.

ID	Counts	RA	Dec	XZ	Photoz	$\log_{10}N_{\text{H}}$	specz	IG (bits)
2	292	214.5296	52.4231	1.68 (+0.62, -1.05)	1.76 (+0.15, -0.19)	22.96 (+0.28, -0.57)	-99.0	1.59
9	280	214.6242	52.45268	1.75 (+0.51, -0.34)	0.79 (+0.16, -0.10)	23.59 (+1.24, -0.29)	-99.0	1.54
15	386	214.538	52.46775	2.9 (+0.13, -0.91)	1.82 (+0.16, -0.14)	23.56 (+0.10, -0.33)	-99.0	1.85
21	903	214.4244	52.47308	1.1 (+0.85, -0.79)	0.98 (+0.10, -0.09)	22.18 (+0.38, -0.48)	1.14825	1.6
24	196	214.4684	52.48146	1.47 (+0.10, -0.66)	1.26 (+0.17, -0.15)	23.63 (+0.20, -0.42)	-99.0	2.35
32	201	214.4394	52.49756	1.3 (+5.56, -0.81)	0.9 (+0.10, -0.12)	23.59 (+0.77, -0.53)	0.872816	1.19
34	184	214.5468	52.50349	3.25 (+0.07, -0.10)	2.9 (+0.19, -0.22)	25.33 (+0.46, -0.41)	2.91	2.89
38	934	214.4415	52.50898	1.86 (+0.43, -1.35)	0.96 (+0.10, -0.10)	23.09 (+0.19, -0.70)	0.985102	1.77
39	159	214.4748	52.5095	2.38 (+0.83, -1.52)	0.94 (+0.16, -0.20)	23.54 (+0.41, -0.78)	-99.0	1.14
56	94	214.5373	52.53091	2.32 (+0.20, -0.34)	2.2 (+0.69, -0.46)	24.75 (+0.71, -0.36)	-99.0	2.04

Notes. *Column 1:* identification number of [Nandra et al. \(2015\)](#). *Column 2:* total detected source counts. *Columns 3 and 4:* J2000 right ascension and declination in degrees. *Column 5:* redshift (XZ). *Column 6:* redshift (photoz) from [Nandra et al. \(2015\)](#). *Column 7:* neutral hydrogen column in units of cm^{-2} from XZ fit. *Column 8:* Reported specz values from [Nandra et al. \(2015\)](#). *Column 9:* information gained in units of bits. All quoted errors are 1σ .

Appendix B: COSMOS and AEGIS-XD results

Similar to Table 2 for CDF-S sources, we present XZ constraints catalogues for the COSMOS and AEGIS fields in Tables B.1 and B.2, respectively. The data products used are described in [Brightman et al. \(2013\)](#) and [Buchner et al. \(2015\)](#).

Figures B.1 and B.5 present the performance of COSMOS and AEGIS samples, respectively, where XZ constrained the redshifts. AEGIS, a relatively deep yet wide field has

the largest number of constrained redshifts across the three fields.

Similar to the CDF-S results, we plot the XZ constraints in three subsamples divided by information gain. Subsample I contains sources with specz and XZ-constrained redshifts (Figs. B.2 and B.6). Subsample II contains sources without specz and compares XZ-constrained redshifts to photoz (Figs. B.3 and B.7). Subsample III contains sources without specz and compares XZ with some redshift information to photoz (Figs. B.4, B.8, and B.9).

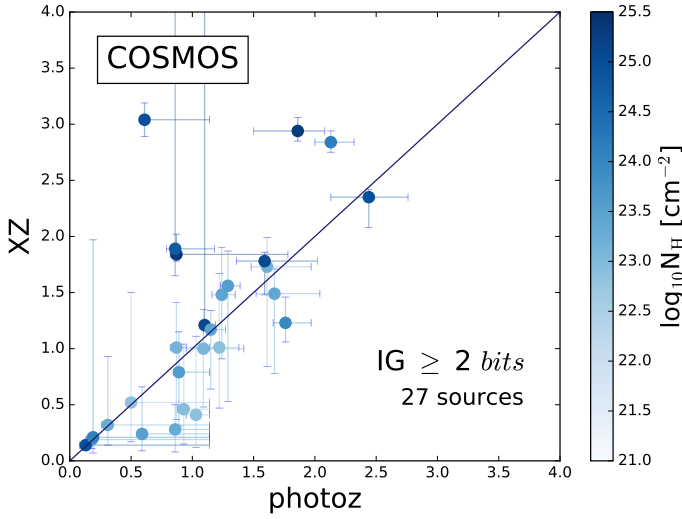


Fig. B.1. Plot of best-fit XZ values vs. photoz from Salvato et al. (2009).

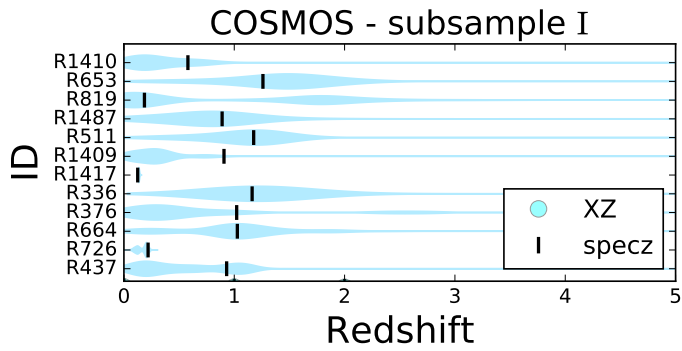


Fig. B.2. XZ (cyan violin plots) and specz (black vertical lines) probability distributions for subsample I ($IG \geq 2$ bits).

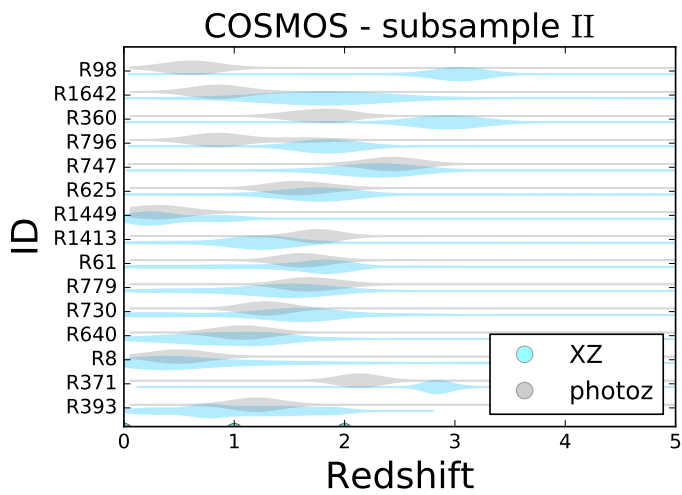


Fig. B.3. XZ (cyan violin plots) and photoz (gray violin plots) probability distributions for subsample II ($IG \geq 2$ bits).

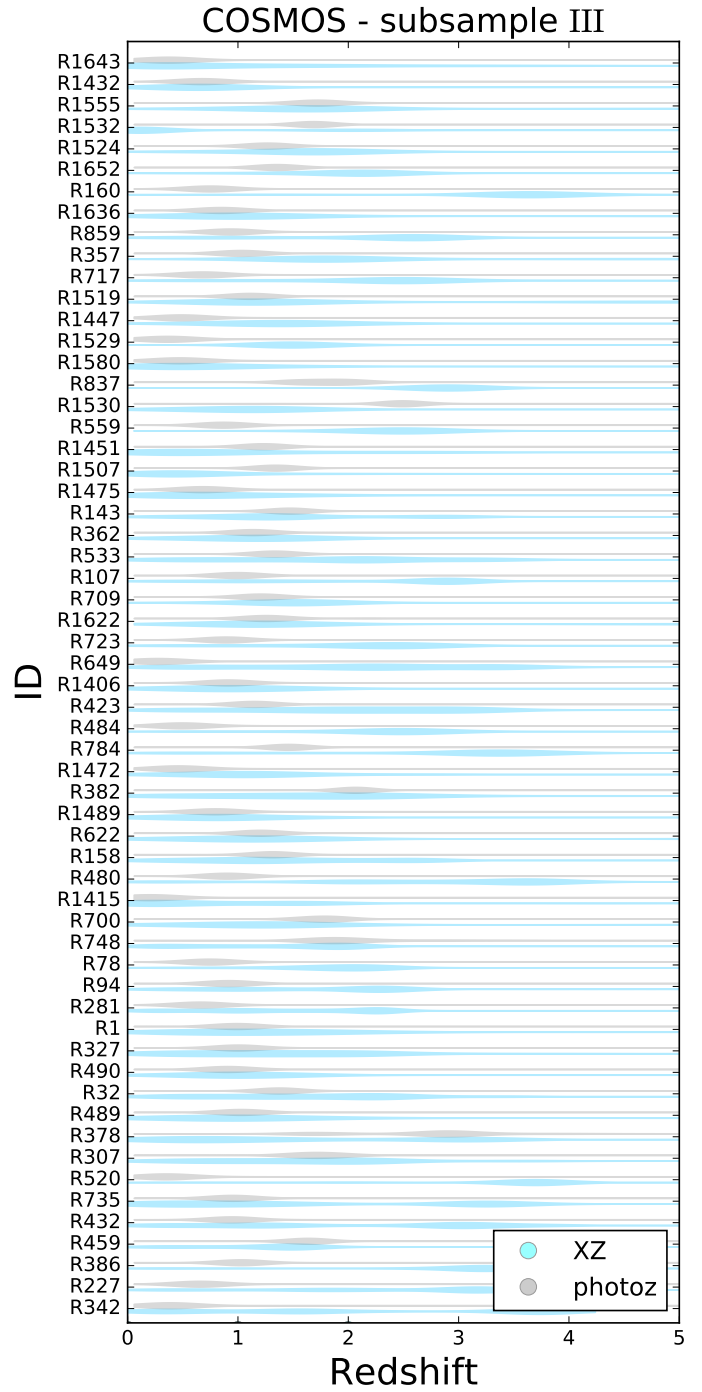


Fig. B.4. XZ (cyan violin plots) and photoz (gray violin plots) probability distributions for subsample III ($IG \geq 2$ bits).

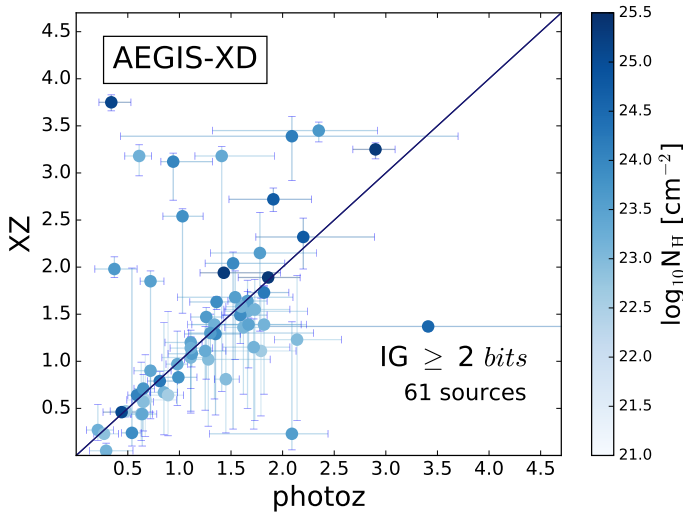


Fig. B.5. Plot of best-fit XZ values vs. photoz from Nandra et al. (2015).

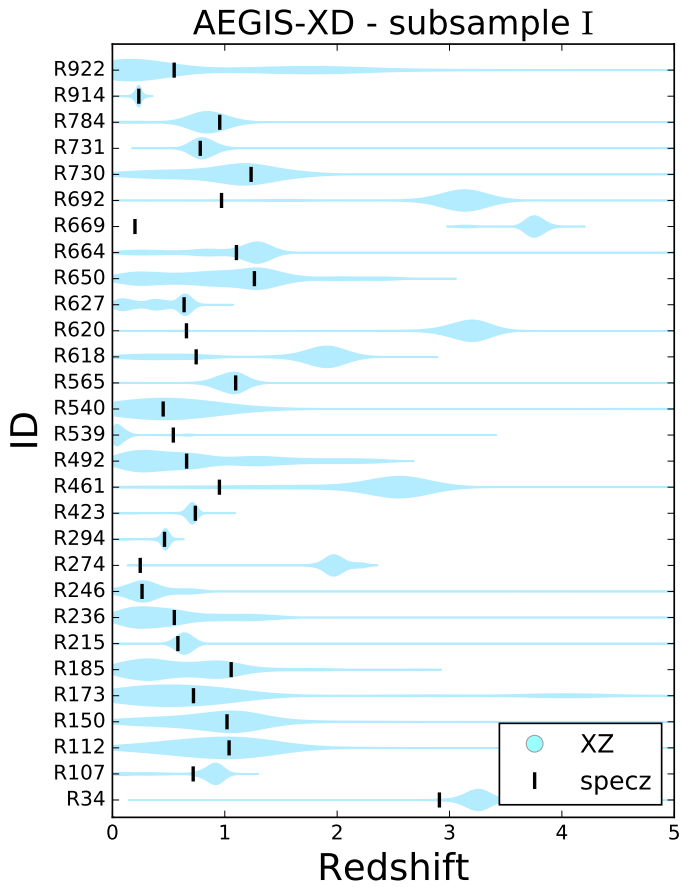


Fig. B.6. XZ (cyan violin plots) and specz (black vertical lines) probability distributions for subsample I ($IG \geq 2$ bits).

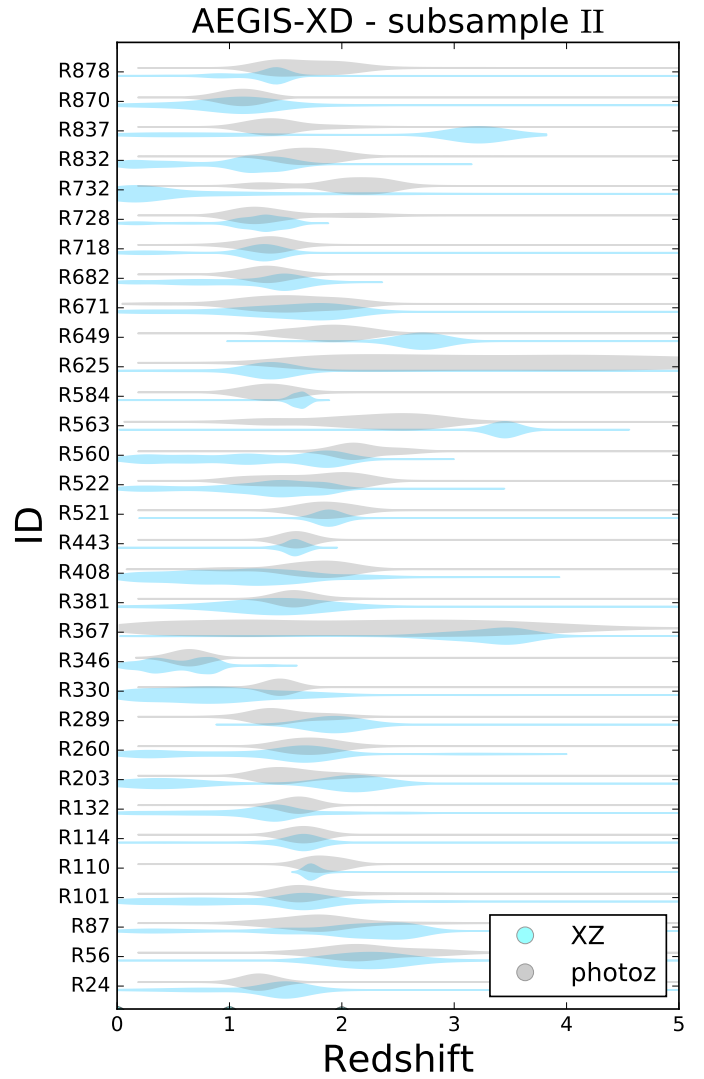


Fig. B.7. XZ (cyan violin plots) and photoz (gray violin plots) probability distributions for subsample II ($IG \geq 2$ bits).

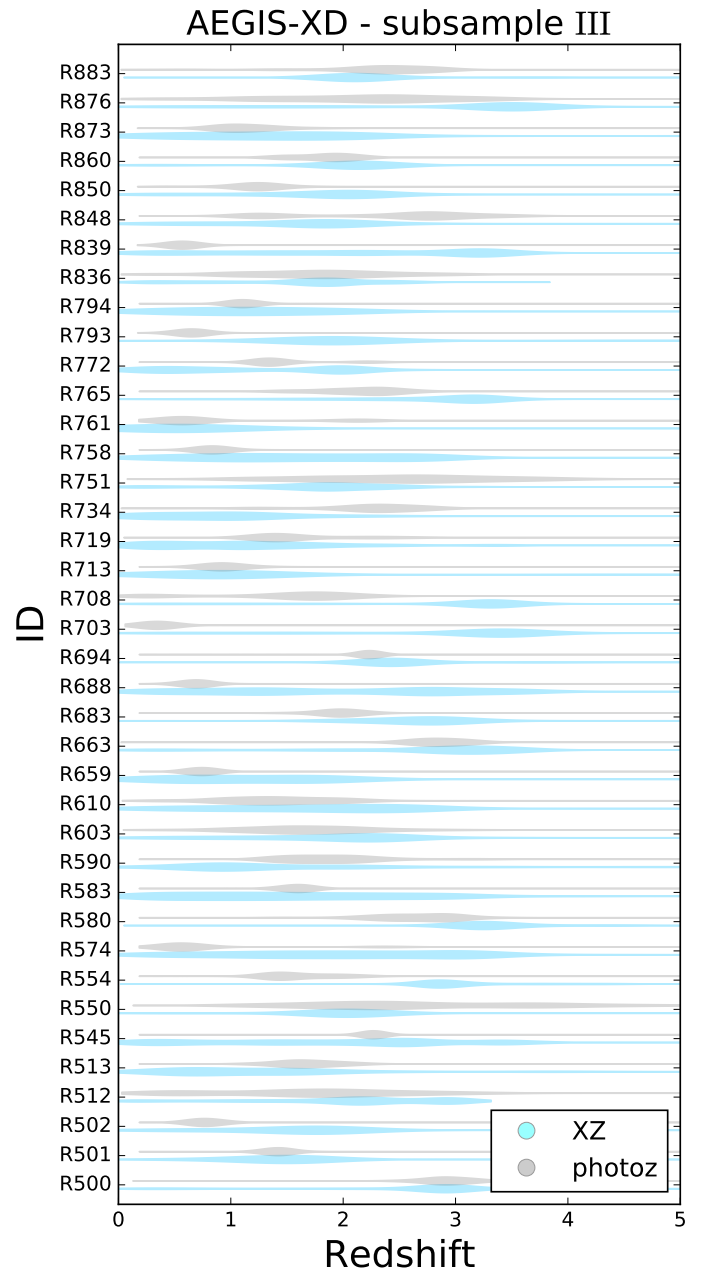
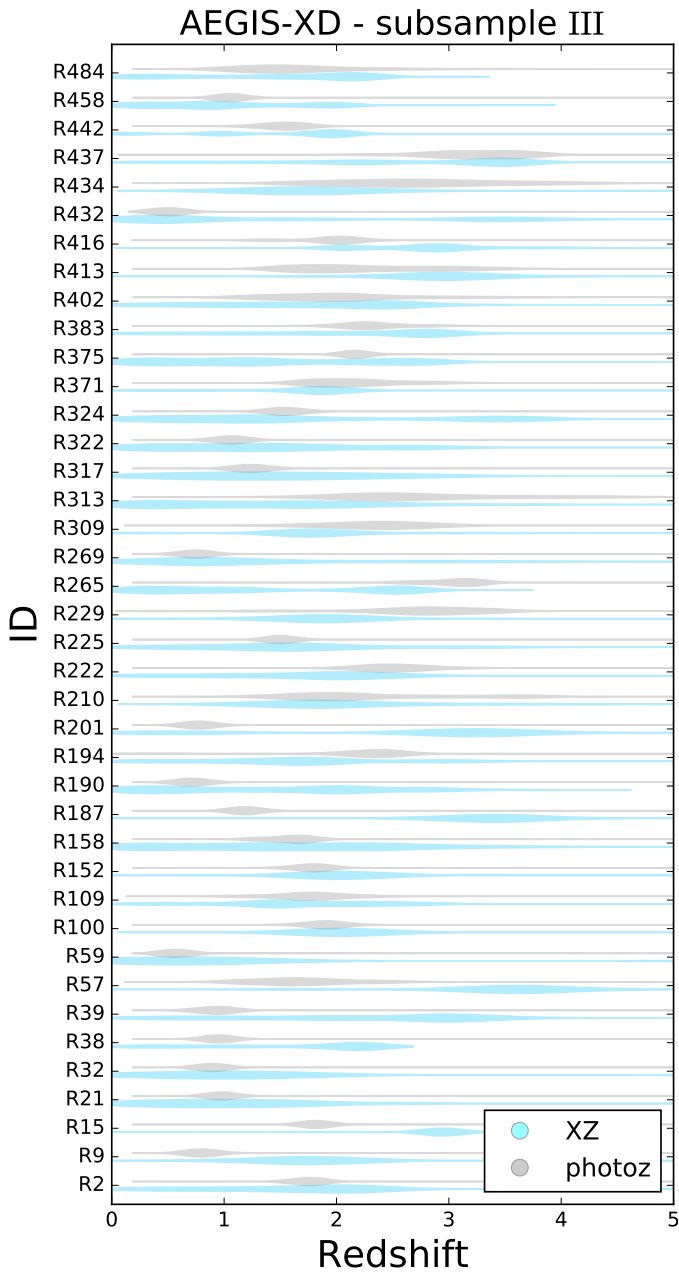


Fig. B.8. XZ (cyan violin plots) and photoz (gray violin plots) probability distributions for subsample III ($IG \geq 2$ bits).

Fig. B.9. Figure B.8 cont.