



THE UNIVERSITY *of* EDINBURGH
School of Physics
and Astronomy

Identifying Hidden Supermassive Black Holes with X-ray Emissions

MPhys Project Report

D. J. Hirshon

Submitted for the 40pt MPhys Project course PHYS11016
March 24, 2024

Abstract

We estimate the effective column density of 13,554 sources from the *ExSeSS* catalogue, made up of *Swift* XRT observations. We calculated a small selection of these values by fitting the spectra of 378 high-count sources. The remaining 13,176 low-count sources were classified using a KNN algorithm trained on the spectral fitting results. Our integrated number counts strongly agree with AGN synthesis models at moderate fluxes, but we observe an excess of sources at bright fluxes. These results provide new constraints to future population synthesis models. We find 9 candidate Compton-thick AGN, and observe no relationship between Compton-thick AGN and galaxy merger events.

Supervisor: Professor J. Aird

Personal statement

I spent the first two weeks of the project reading literature on the history of observational X-ray astronomy, techniques for identifying obscured AGN, and background on the catalogues I would be using. I had weekly meetings with my supervisor and group meeting with his cohort. These group meetings were useful for understanding the broader field of AGN research and for gaining practical advice. Following the first group meeting, I switched to Zotero and Jupyter Notebooks for citation management and coding.

The majority of work done in the project was performed using my Jupyter notebook. I started by importing the data catalogue created by Jack Delaney—one of my supervisors PhD students—and implemented literature methods suggested by my supervisor. For purposes of data handling and visualization, I learned how to use the Python libraries `pandas`, `seaborn`, and `plotly` over my more frequent `numpy` and `matplotlib`. The biggest problems in the first four weeks were appropriately understanding the error calculations performed in the catalogues, and understanding how to appropriately use `WebPIMMs`, an online website for calculating count rates in one band given the count rates in another band. These were not large hurdles, and I produced a simple model for identifying obscured AGN by the end of the fourth week.

The next goal was to develop more complicated models of measuring obscuration using `XSPEC`, a software from the `HEASOFT` package. This was performed on *cuillin*, the Institute for Astronomy’s computing cluster. I initially had difficulties importing the software into Python, but this was resolved within the week.

The next four weeks were spent on a wide range of issues including model initialization, data visualization, and switching to `BXA`, a Bayesian X-ray analysis library. I tried to use three models: the first being the Aird et al. (2015) model, and the second and third being table models available online (Baloković et al. 2018; Buchner et al. 2019). The table models would ideally be used within the project, but I was unable to get them working as the kernel crashed whenever I operated on them. The data visualization issue was resolved by combining `matplotlib` code with `PyXspec`, and binning the data so that errors appeared reasonable. Following success with the first model, I tried to switch to `BXA` in order to improve results. My initial attempts at getting `BXA` to work were failures—the solutions were highly degenerate, and took very long to run computationally. I decided not to use `BXA` within this project due to these setbacks.

I finished the first semester by using the Aird et al. (2015) `PyXspec` model to measure the level of obscuration for all high-count sources in the data.

In the beginning of the second semester I started implementing the KNN machine learning algorithm to generalize the high-count results to the low-count sources. I spent a few weeks trying to optimize the model, and also re-performing spectral fits to correct any issues.

On several occasions I looked at the most obscured sources (which had slight variation depending on which version of spectral fitting i was on), and examined their optical and infrared images in order to determine any relationship with galaxy mergers.

After this I took about two weeks to calculate the intrinsic infrared and X-ray luminosity. The infrared luminosity derivation took less time, but I had to get the X-ray luminosities through XSPEC, which took a while. Once I had these, I removed poorly fit sources using the MIR-XRAY relation and the reduced cash statistic. (Reduced cash statistic is a bit of an assumption).

After this I spent a few weeks calculating the $\log N - \log S$ plots within each obscuration bin, and comparing the results to other literature. It took a while to get this sorted because of mistakes with error propagation, but also due to challenges with binning sources.

The last few weeks of the project were spent double checking previous work, e.g. making sure I used the Aird 2015 XSPEC model the whole way through, that spectral fitting had been done correctly, that the KNN test sample did not contain the training sample within it, and that the $\log N - \log S$ binning was done correctly. I started gathering N_H distributions within other literature in order to make better comparisons.

In the final week I removed stellar contaminants from the sample. This should have been done earlier but I mistakenly thought I had done this at the start of the project. I also re-validated the KNN algorithm using a different statistic, and compared the results to a few dummy classifiers.

Acknowledgments

I want to thank my supervisor James Aird for all his help and guidance over the course of this project, he was always readily available for my questions—regardless of how relevant they were—and shared a lot of enthusiasm with me on the history of X-ray astronomy. I’d also like to thank Jack Delaney, who produced the *ExSeSS* catalogue. He provided great advice and was always very friendly. I also want to thank everyone else in the “Aird Crew”, who provided many interesting discussions, questions, and miscellaneous advice throughout the year.

This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester.

Contents

1	Introduction	1
1.1	Modern AGN Model	1
1.2	X-ray Astronomy	2
1.3	Quantifying Obscuration	3
1.4	The <i>ExSeSS</i> Catalogue	4
1.5	Motivations	5
2	Data Preparation	6
3	XSPEC and Our Model	7
4	Simple Method of Quantifying Obscuration	10
4.1	Band and Hardness Ratios	10
4.2	Initial Estimations Using Hardness Ratios	11
4.3	Two-dimensional Hardness Ratios	12
5	Spectral Fitting of hard-cut Objects	14
5.1	Performing Spectral Fitting	14
5.2	Validating our Spectral Fits	15
5.2.1	The Cash Statistic	15
5.2.2	The X-ray and Mid-Infrared Luminosity Correlation	18
6	Generalizing Results to Low-Count Data Using a KNN Algorithm	21
6.1	KNN Algorithms	21
6.2	Model Evaluation	21
6.3	Generalisation Results	24
6.4	Comparing Obscured Fractions to Other Surveys	26

7	Comparison to AGN Population Synthesis Models	28
7.1	AGN Population Synthesis Models	28
7.2	Calculating Number Counts	29
7.3	Comparing to the Gilli and Ueda models	29
8	Relationship Between Obscured AGN and Galaxy Mergers	33
8.1	Do Obscured AGN Indicate a Galaxy Merger Event?	33
8.2	Optical and Infrared Imaging of our Compton-thick Sources	33
9	Conclusion	36
	Appendices	42
A	Intrinsic Mid-Infrared Luminosity Derivation	42
B	Additional Spectra	44

1 Introduction

1.1 Modern AGN Model

Supermassive black holes are thought to be at the centre of most galaxies. Roughly 10% of these are known as Active Galactic Nuclei (AGN) because they are actively accreting material onto the black hole. These objects are incredibly luminous due to the constant influx of material by an accretion disk. They emit energy across the entire electromagnetic spectrum, however we will only be interested in their X-ray and infrared emissions within this report. Observational data is not uniform for all AGN, leading to many different naming conventions for distinct cases. We will not explore the menagerie of AGN classes within this report but direct an interested reader to Padovani et al. (2017) or Risaliti and Elvis (2004) for more details. A large topic of debate within this field is on AGN unification theory, which posits that the many different types of AGN are powered by the same underlying physical system, and are only different because of the observers viewing angle (Antonucci 1993).

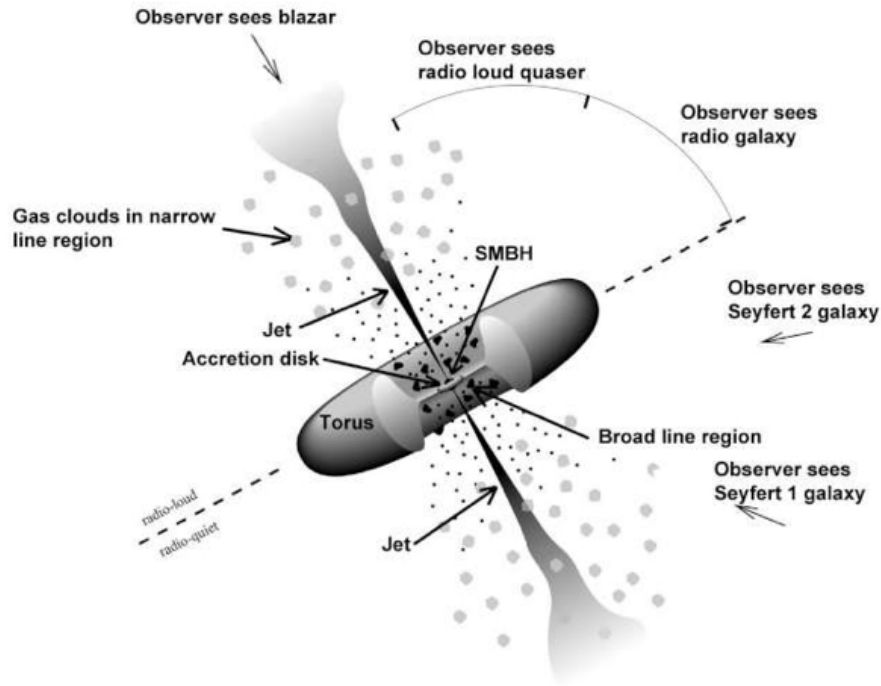


Figure 1. General model of AGN. The inclination angle changes the type of AGN observed, varying from Seyfert 1 galaxies at the bottom to blazars at the top. Image taken from NASA (2024)

The unified model of AGN shown in Figure 1 is composed of several elements: the accretion disk, obscuring torus, broad line region, narrow line region, and radio jets. We will focus on only two components within this figure, namely the accretion disk and the obscuring torus, as well as a component not shown in Figure 1: the hot corona. We note here that the details provided in this section are only a brief summary of the surrounding literature (Hickox and Alexander 2018; Netzer 2015). The accretion disk is a ring of particles that swirls around the

black hole centre. Portions of the accretion disk fall inwards and are captured by the black hole, fueling both its growth and its emissions. Thermal emissions from the viscosity within an accretion disk tend to produce photons at UV-optical wavelengths, but low-energy X-ray emissions can be produced in the inner-most and therefore hottest regions of the accretion disk. The next component we will discuss is the hot corona, which is typically understood to be a dense collection of hot electrons located above the black hole and accretion disk. High-energy X-ray emissions are produced within the corona due to the inverse Compton scattering of photons from the accretion disk, whereby photons gain energy from their collisions with higher-energy electrons. The obscuring torus is located beyond the accretion disk and is a clumpy region of dense gas. It’s much colder than the accretion disk and absorbs incoming photons, obscuring the light that we see.

It is worth specifying that the above model describes an AGN in radiative mode. AGN have two different accretion modes, radiative and jet mode, where the latter is radiatively inefficient. In this second mode the AGN has a smaller torus and ejects its energy primarily through radio jets (Heckman and Best 2014). We will not expand further on jet mode AGN as our selection methods are biased towards radiative mode AGN.

1.2 X-ray Astronomy

X-ray astronomy is a relatively young branch of astronomy, with focused research only starting in 1962 (Giacconi 2009). X-ray space telescopes began launching in the next decade, and have since become known as a viable way for measuring the obscuration of active galactic nuclei (AGN). This is because X-ray emissions from a supermassive black hole are unaffected by line-of-sight absorption and are typically unaffected by contamination from the host galaxy except in cases of extreme obscuration (Hickox and Alexander 2018; Lansbury et al. 2017), where obscuration refers to the column density of hydrogen surrounding an AGN.

To precede the rest of this section we define some of the nomenclature used within this report, as well as any oddities of X-ray astronomy. Instead of writing X-ray data in the form of wavelength or frequency, of which is popular in many other regions of astronomy, X-ray astronomers typically discuss X-rays by their energy values. These energies are labelled as “soft” or “hard” depending on if they are low or high energy, respectively. These are relative terms depending on the data used, so the exact values vary by literature. Within this report, we work with data from 0.3–10 keV and are thus unable to probe extremely hard X-rays. We define our soft band as being between 0.3–1 keV, the medium band as being between 1–2 keV, the hard band as being from 2–10 keV, and the total band as being from 0.3–10 keV. The reasons for quantifying these values will become useful for measuring obscuration later on. An additional term is the count rate, which is defined as the rate at which photons are registered. This is a separate measurement for each energy band, allowing one to compare the rate of photon detections between soft and hard bands. The last relevant detail is that X-rays are typically a much more sparse dataset than other wavebands; there is a substantial amount of uncertainty from each source, with Poisson statistics used frequently. To differentiate between our most secure sources and our most uncertain ones, we define our

sources with over 100 individual photon counts as being high-count sources, and sources with less as being low-count sources. This will be discussed in more detail in Section 2.

1.3 Quantifying Obscuration

The level of obscuration of an AGN refers to the effective column density of hydrogen within the obscuring torus. Within the context of this report, the most obscured source we would be able to identify is “Compton-thick” (CT). A Compton-thick source has $N_H \geq 1.5 \times 10^{24} \text{ cm}^{-2}$, and is defined by a hydrogen column density that exceeds the inverse of the Thomson scattering cross-section. It is difficult to observe more obscured sources without relying on multi-wavelength catalogues in our selection process, as their X-rays will be absorbed within the torus. Since our 0.3–10 keV observations are biased against detecting CT-AGN, all sources we find at this level of obscuration will require further follow-up to confirm their obscured status. Less obscured sources, such as Compton-thin AGN, are defined as being less obscured than CT-AGN but more obscured than the Galactic gas density. Within this report we partition our obscuration into five distinct classes, shown in Table 1, to better distinguish between our sample. Note the difference between the definition of a Compton-thick source given earlier and the one in the table—this is to fit better in line with the literature, which frequently compares estimations in logarithmic bins, e.g. $\log N_H = 20 - 21$ or $\log N_H = 24 - 25$ (Ricci et al. 2015). The effect of each class on a typical spectrum is shown in Figure 2.

Obscuration	$\log N_H$
Unobscured	< 21
Lightly obscured	≥ 21
Obscured	≥ 22
Heavily obscured	≥ 23
Compton-thick	≥ 24

Table 1. Table of different obscuration bins. Note that this is specific to this current report. There are higher levels of obscuration than Compton-thick, but we will not be able to probe this regime with our data. Additionally, all labels from lightly obscured to heavily obscured AGN could be grouped together as Compton-thin.

Measuring the levels of obscuration is important because photons are absorbed within the obscuring torus, making heavily obscured AGN more difficult to observe. Without constraining the value for N_H , it is difficult to determine the intrinsic properties of an AGN. This places large uncertainties on the total population of CT-AGN. Researchers have tried to estimate an intrinsic fraction of CT-AGN—the fraction of CT-AGN relative to the total AGN population—using a myriad of techniques. These include modelling the X-ray luminosity function (XLF) to predict the total level of CT-AGN (Aird et al. 2015), as well as correcting for selection biases (Burlon et al. 2011; Ricci et al. 2015). Through the usage of these techniques, the intrinsic fraction was predicted to be $\sim 10\text{--}30\%$ as of a decade ago. More recent studies such as Carroll et al. (2023) have arrived at a higher intrinsic fraction

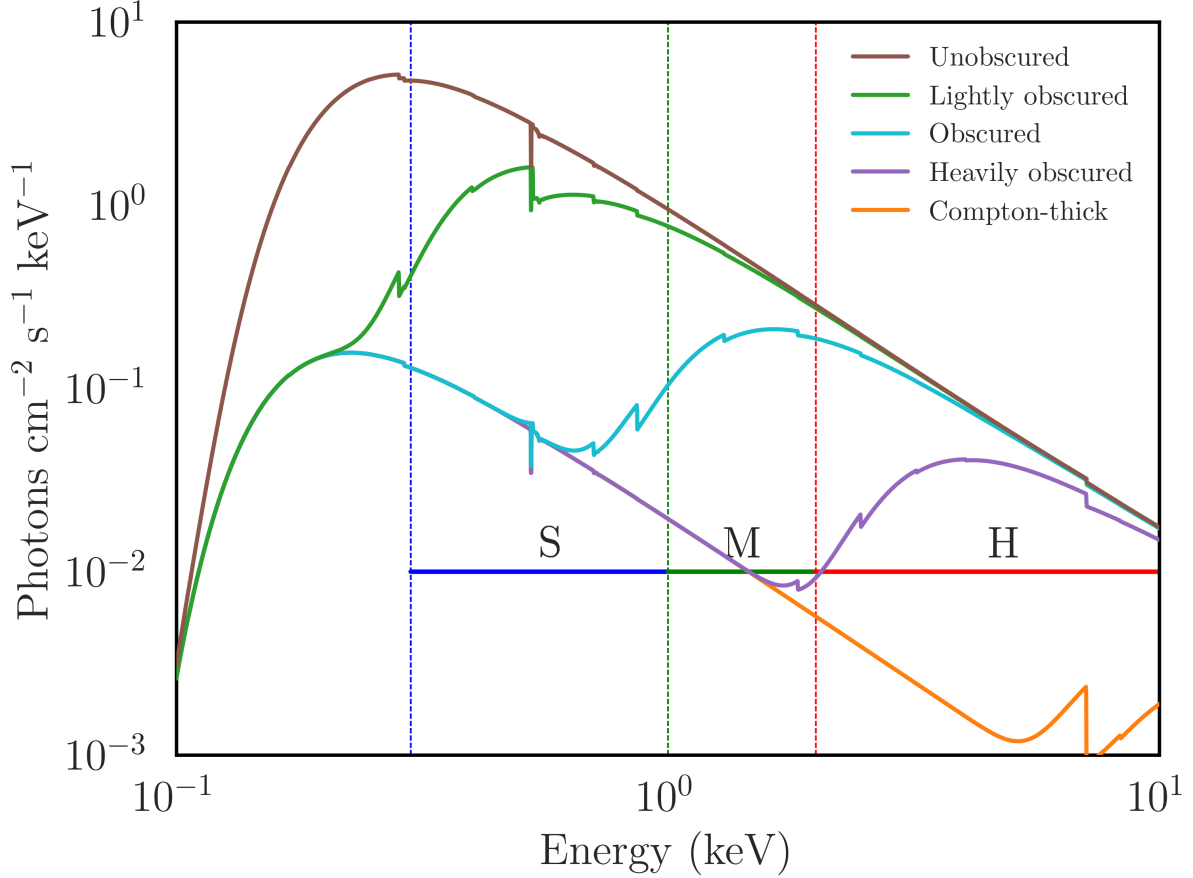


Figure 2. Effect of obscuration on the number of photons observed per unit area per unit time per keV. The brown line represents an unobscured source ($\log N_H = 20$), the green line a lightly obscured source ($\log N_H = 21$), the cyan line an obscured source ($\log N_H = 22$), the purple line a heavily obscured source ($\log N_H = 23$), and the orange line a Compton-thick source ($\log N_H = 24$). The template spectrum is constructed assuming a redshift of 0 and a photon index of 1.9. The blue, green, and red horizontal lines indicate the soft, medium, and hard bands, respectively.

of ~ 50 - 60% due to incorporating mid-infrared data in their selection process. Constraining this value will help resolve a disconnect between the measured X-ray background luminosity and the contribution from unobscured sources (Akylas et al. 2016). Identifying more CT sources may also support AGN and galaxy co-evolution models, where the most obscured sources might indicate periods of intense black hole growth (Kocevski et al. 2015).

1.4 The *ExSeSS* Catalogue

This report makes use of the Extragalactic Serendipitous Swift Survey (*ExSeSS*) catalogue created by Delaney et al. (2023). This catalogue is constructed from the serendipitous detections of an earlier survey, the Swift-XRT Point Source (*2SXPS*) catalogue (P. A. Evans et al. 2020), which compiles all the data collected with the *Swift* X-Ray Telescope (XRT). It covers the 0.3–10 keV energy range, which is harder than *eROSITA* (0.2–8 keV) (Predehl et

al. 2021) but softer than *XMM-Newton* (0.1–15 keV) and *NuSTAR* (3–79 keV) (Harrison et al. 2013; Jansen et al. 2001). This instrument is attached to the *Swift* Burst Alert Telescope (BAT), which scans the sky for Gamma-ray bursts. When BAT detects a Gamma-ray burst, the *Swift* satellite points XRT in the direction of the burst in order to obtain an X-ray counterpart and improve the position. In addition, there are times when *Swift* XRT is pointed at certain targets rather than acting from a Gamma-ray burst. The *ExSeSS* catalogue is constructed by removing the targets of both “triggered” and “pointed” observations, leaving only the serendipitous sources. This makes the *ExSeSS* catalogue a patchwork collection of X-ray observations. It contains X-ray data from 79,342 sources, as well as spectroscopic and photometric redshift predictions. In addition to X-ray data, the *ExSeSS* catalogue is cross-matched with data from the Wide-field Infrared Survey Explorer (WISE). This satellite has 4 channels which cover wavelengths of 3.4, 4.6, 12, and 22 μm (Wright et al. 2010). This is particularly useful because X-ray and infrared observations probe different regions of an AGN, and develop a more complete picture when used together. We will explore the relationship between X-ray and infrared emissions in Section 5.2.2 as a means of validating our X-ray spectral fits.

ExSeSS contains observations ranging from 0.3–10 keV and covers a total area of 2086.6 deg^2 . These values are especially relevant as they probe sufficiently hard energies over larger areas than other surveys. There are only a few collections of data that probe the same energy range—the *Chandra* source catalogue and 4*XMM* (I. N. Evans et al. 2010; Webb et al. 2020)—but these only cover 800 and 1283 deg^2 respectively. X-ray telescopes at higher energy ranges like *NuSTAR* and *Integral* have been able to identify numerous obscured AGN (Gandhi et al. 2014; Lansbury et al. 2017), but suffer from poor angular resolution and low area coverage. Thus, while these higher energy telescopes can procure better data for obscured AGN identification, the size of the samples are limited. This issue has led some researchers to suggest the usage of wider-field surveys at low energies to identify candidate CT-AGN, which can then be verified by higher-energy telescopes (Akylas et al. 2016). The benefit of using the *ExSeSS* catalogue is thus its edge in being a large wide-field survey with good sensitivity across the 2–10 keV band, having been cross-matched with infrared observations, and completed with photometric and spectroscopic redshift estimations.

1.5 Motivations

Within this report, we will determine the effective column density of a large number of sources. A small subset of these sources will have interesting physical properties which could then be subject to detailed follow-up in future works. This includes using either a higher energy X-ray telescope such as *NuSTAR* to confirm whether a source is Compton-thick, or observing with *JWST* to study optical and infrared properties. Improving our understanding of CT-AGN compared to unobscured AGN would uphold or challenge the unified model of AGN.

In addition to this, the *ExSeSS* catalogue probes a unique parameter space. The results we find in this report will therefore provide a new constraint to pre-existing AGN population

synthesis models. These models seek to replicate the physical properties of the Universe as observed in X-ray wavebands and will be discussed further in Section 7. In addition to testing older models, our observations are capable of calibrating future ones, as no current synthesis model uses surveys from *Swift* XRT.

We also experiment with a novel method for quantifying AGN obscuration within Section 6. We do this with a K-Nearest Neighbours machine learning algorithm, combining the observed data of Section 4 with the spectral fits performed in Section 5. A successful application of a K-Nearest Neighbours machine learning algorithm would be a more computationally efficient means of estimating obscuration that could be improved by subsequent observations.

The structure of this report is as follows: In Section 2 we discuss our implemented constraints on the *ExSeSS* catalogue, as well as any supplementary data we take from the *2SXPS* catalogue. In Section 3 we explain the physical motivations of our *XSPEC* model, as well as the conditions at which it begins to break down. In Section 4 we begin to estimate obscuration using photon count ratios, and how this is a function of redshift. In Section 5 we explain how we perform spectral fitting on our high-count sources. We discuss the quality of our fits and also examine methods for external validation. In Section 6 we train a KNN machine learning algorithm on our spectral fitted sources and apply it on our low-count sources. In Section 7 we compare the number counts of our sample to the expected number counts of AGN population synthesis models, comparing the values at different levels of obscuration. We conclude with Section 8, where we examine the optical and infrared properties of our Compton-thick candidates to determine whether or not they indicate a relationship with merger events.

2 Data Preparation

We initially constrained the *ExSeSS* catalogue to ensure that the dataset we worked with produced sensible results. We mandated that all sources must have the following: a detected count rate in the hard band, a redshift estimation, a flux detection in the total band, and a confidence value of `true`. The confidence value denotes that an object is likely an AGN. Upon applying these cuts, we reduce the *ExSeSS* catalogue from 63,564 sources to 13,972. We also removed all sources within the catalogue that were flagged as stellar objects.

The count rate and flux detection cuts were necessary as we could not determine the observed X-ray luminosity of objects without these values. The redshift cut was done to remove erroneous predictions, as the observed X-ray energy varies from the emitted energy according to the Doppler effect due to the expansion of the universe:

$$E_{observed} = \frac{E_{emitted}}{(1+z)} \quad (1)$$

or, in terms of wavelength:

$$\lambda_{observed} = (1+z)\lambda_{emitted}$$

Sources without a valid redshift will thus have uncertain column density predictions and ratio estimations, so we remove them here. It is often common to purge catalogues of X-ray sources with a total luminosity below 10^{42} erg s $^{-1}$ as fainter objects may be star-forming galaxies, where the emission comes from the combined effect of X-ray binaries instead of AGN (Hickox and Alexander 2018), but we did not implement this cut for two reasons. The first is that the *ExSeSS* catalogue should largely be made up of AGN sources, hence the `true` confidence value. The latter reason is that the most obscured sources within our sample were likely to have greatly diminished observed luminosities. Removing these sources from our dataset would thus have been detrimental.

Following this purge, we supplemented the remaining sources with data from the *2SXPS* catalogue (P. A. Evans et al. 2020). These additions include the Galactic absorption as calculated by Willingale et al. (2013), the measured background and source counts within the total and hard bands, and the ratio of count rates between different energy bands, otherwise known as hardness ratios.

Galactic absorption has a relatively small effect on X-ray compared to the obscuring torus but is important to include in order to capture all the physical properties of the system. The inclusion of source and background counts was important as it allowed us to partition our objects by their number of photon counts, rather than their count rates, and thus determine a tight sample of high-count sources with reduced uncertainties. Our high-count selection kept sources with at least 100 photon counts in the hard band, and is hereafter referred to as the hard-cut. This cut reduced our sample from 13,554 objects to 519.

The last addition, the hardness ratios, were taken because they were calculated using a Bayesian method outlined in Park et al. (2006) for sources with less than 100 photon counts. This prevents the hardness ratios from being pegged at 1 due to a lack of counts in the medium or soft bands, and improves the value for $\approx 95\%$ of the sources.

3 XSPEC and Our Model

X-ray astronomers frequently use *XSPEC*, an X-ray spectral fitting package, to analyze their data. This software interprets physical models that have calculated spectra, and can then be compared to observed spectra. The initial step when using this software is to assume a model. These models quantify the amount of radiation originating from an AGN, as well as the amount of absorption the photons will undergo. There is a large range of complexity between different models. A simple model will not capture all the physical properties seen in observed spectra, but will likely generalize well to a large collection of data. A complicated model may account for the physical properties better but may struggle to generalize across large surveys. Previous work by Buchner et al. (2014) determined that there are four key components necessary to capture the variety in AGN torus models: an intrinsic power law, re-processed radiation by a cold obscurer, an unabsorbed power law from scattering, and Compton-reflection. For these reasons, we adopted the model used by Aird et al. (2015), henceforth referred to as the A15 model, as it contained all of the required components. The

A15 model permits three ways that X-rays can reach our line-of-sight, shown pictorially in Figure 3. To help the reader understand our model we will reference **XSPEC** components as [*component name*] when we discuss the physical phenomena that they seek to emulate.

The main contribution to the observed spectra at lower levels of obscuration are the X-rays that move directly into our line-of-sight, suffering absorption from the torus. Their contribution to the spectra is affected by photoelectric absorption [*ztbabs*] and optically-thin Compton scattering within the torus [*cabs*]. It is also defined by a high energy cutoff around 300 keV, but this has a small impact on our 2–10 keV observations [*zhigect*]. The A15 model defines a scattering fraction, $f_{scatt} = 0.02$, which is then removed from this component, as roughly 2% of the photons can be scattered out of our line-of-sight.

The reverse is also true: photons travelling away from us can be scattered into our line-of-sight by ionised gas. For low redshift Compton-thick AGN, these photons are dominant as the previous component would be fully absorbed. As these X-rays travel outside the path of the torus, they do not suffer any non-Galactic absorption. The A15 model multiplies this value by the previously mentioned scattering fraction, as only 2% of these photons should move into our line-of-sight.

The final contribution tends to dominate at higher redshifts and is composed of photons that reflect off of cold matter before passing through the torus [*pexrav*]. Cold is a relative term here, where the accretion disk and torus are much colder than the hot corona. Reflection gives these photons a distinct spectral shape from the absorbed component. The A15 model fixes the angle of inclination of reflection to be 30° assuming the geometry of an infinite plane. Due to this angle, the assumption holds well for cases in which reflection occurs in the accretion disk but is poor when reflection occurs in the torus. This contribution dominates at higher redshifts, as the peak of this contribution is around 20–30 keV. From Equation 1, this component will then contribute at redshifts around $z \geq 2$.

In all three cases, the initial number of photons emitted are described by a power law intrinsic to the hot corona [*zpowerlw*]. They also are affected by Galactic absorption [*tbabs*].

For a more in-depth explanation of the definitions behind individual components, please see Aird et al. (2015) and the **XSPEC** documentation¹. We note here that we updated the *wabs* and *zwabs* components used within Aird et al. (2015) to their more recent versions, *tbabs* and *ztbabs*. Our model then takes the following form within **XSPEC**:

$$\text{tbabs} * [\quad 0.98 * \text{ztbabs} * \text{cabs} * \text{zpowerlw} * \text{zhigect} \\ 0.02 * \text{zpowerlw} \\ \text{ztbabs} * \text{pexrav}]$$

While the A15 model satisfies the minimum requirements for our purposes, we find it important to note its limitations and the areas where a different model could improve results.

¹<https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/Models.html>

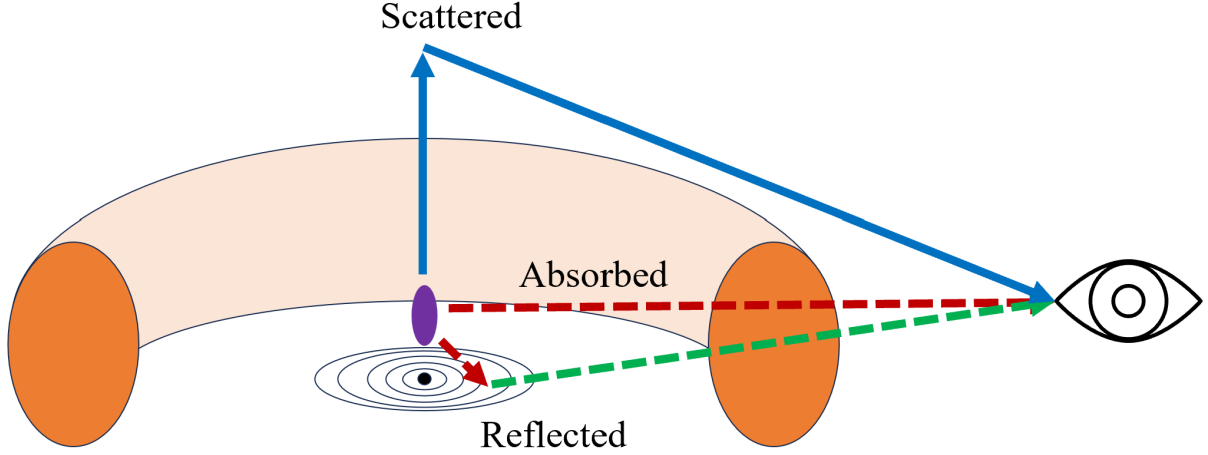


Figure 3. Depiction of the physical processes underpinning our *XSPEC* model. The X-rays originate from the hot corona (shown in purple). The blue arrow indicates X-rays that are scattered into our line-of-sight by ionised gas. The dashed red arrow is the contribution from X-rays that move directly through the torus (shown in orange) into our line-of-sight. The dashed green arrow is the contribution from X-rays which reflect off of the accretion disk.

One issue is with the cold-matter reflection component. This component assumes the geometry of an infinite plane, which is not as physically accurate as toroidal geometry. The ramifications of this is that it only allows for reflection off the accretion disk. This poorly constrains the parameter and is not physically accurate. More advanced models such as **borus02** (Baloković et al. 2018) or **UXCLUMPY** (Buchner et al. 2019) assume a toroidal geometry and hence contain more variance in the reflection parameter. Another limitation of the A15 model is its inability to make use of the 6.4 keV Fe $K\alpha$ emission line, which is a strong marker for Compton-thick levels of obscuration (Buchner et al. 2014; Hickox and Alexander 2018; Netzer 2015).

To conclude this section we include values for the photon index and Galactic absorption that we will use for our simplest models. The photon index characterizes the slope of the X-ray power law. It is related to the number of photons emitted per unit energy and time by:

$$N_\gamma \propto E^{-\Gamma}$$

where N_γ is the number of photons emitted per unit energy and time, E is the energy, and Γ is the photon index (Netzer 2015). We set our photon index to $\Gamma = 1.9$ as Γ typically has a Gaussian distribution centred around 1.9 (Fotopoulou et al. 2016; Ishibashi and Courvoisier 2010; Mateos et al. 2008; Nandra et al. 2024; Ricci et al. 2017a). We assumed the Galactic absorption to be $1 \times 10^{20} \text{ cm}^{-2}$ as this was roughly the average value found within our sample. We will assume these values in the next section in order to create simple obscuration estimates with hardness ratios, but we will be more thorough in Section 5 when we fit our model to observed spectra.

4 Simple Method of Quantifying Obscuration

4.1 Band and Hardness Ratios

Obscured AGN are likely to have higher count rates in the hard bands than in the medium and soft bands because the higher energy photons can penetrate the torus more easily. Heavily obscured candidates can thus be found by measuring the ratio of photons between hard and softer energies. There are two types of ratios, often labelled as band ratios:

$$BR_{HM} = \frac{H}{M} \quad \text{and} \quad BR_{MS} = \frac{M}{S} \quad (2)$$

and hardness ratios:

$$HR_{HM} = \frac{H - M}{H + M} \quad \text{and} \quad HR_{MS} = \frac{M - S}{M + S} \quad (3)$$

where H , M , and S are the count rates in the hard, medium, and soft bands, respectively. These can be re-configured to compare the medium and soft bands. Prior studies tend to choose one of these ratios as a preference, e.g. Alexander et al. (2008) and Lansbury et al. (2017) for band ratios, or Giacconi et al. (2001) and Silverman et al. (2008) for hardness ratios. Within this report, we will focus on hardness ratios to be consistent with the data taken from the *2SXPS* catalogue. Hardness ratios are bounded from -1 to +1, where a more positive value typically indicates a more obscured source. These assumptions only hold at low levels of obscuration, as CT-AGN will have most of their X-ray emissions suppressed at energies below 10 keV, and will thus have hardness ratios near zero. At higher redshifts, the suppressed emissions from CT-AGN will be replaced by higher energy emissions that have been redshifted down by the Doppler effect expressed in Equation 1. Subsequently, the hardness ratio for CT-AGN will become more positive as redshift increases.

A high band ratio indicates a more obscured source, while a hardness ratio closer to +1 is likely an obscured source. These estimations hold until the Compton-thick regime, at which point more complicated methods are required.

Using **XSPEC**, one can calculate the expected count rates from a source at designated obscuration values. A rough approximation for the obscuration can then be found by comparing the observed count rate ratios to the **XSPEC** estimation. If a source is above the **XSPEC** track, then it is more obscured than the obscuration value used for the **XSPEC** calculation. This is a rough approximation for the obscuration of the source but can be a useful method of identifying a small sample of highly obscured targets. This has been done by Lansbury et al. (2017), where they compared the band ratio between 8–24 and 3–8 keV, and identified several Compton-thick sources above a band ratio of 1.7. We note that the previous technique was used on data gathered by *NuSTAR*, and thus probes much higher energies. This allowed them to avoid degeneracies between unobscured and Compton-thick sources. This is not the case for us, so it is difficult to distinguish Compton-thick sources from less obscured sources

at low redshift. To ensure we do not predict an abundance of Compton-thick sources, we implement an additional condition that the Compton-thick hardness ratio threshold must be higher than any other classification threshold before a source can be classified as Compton-thick.

4.2 Initial Estimations Using Hardness Ratios

We first produced tracks in XSPEC at four different levels of obscuration by calculating the count rates within our soft, medium, and hard bands. We varied the redshift from 0 to 7 and calculated the hardness ratios at column densities of $\log N_H = 21, 22, 23$, and 24. We then interpolated these values across the redshift values of our sample.

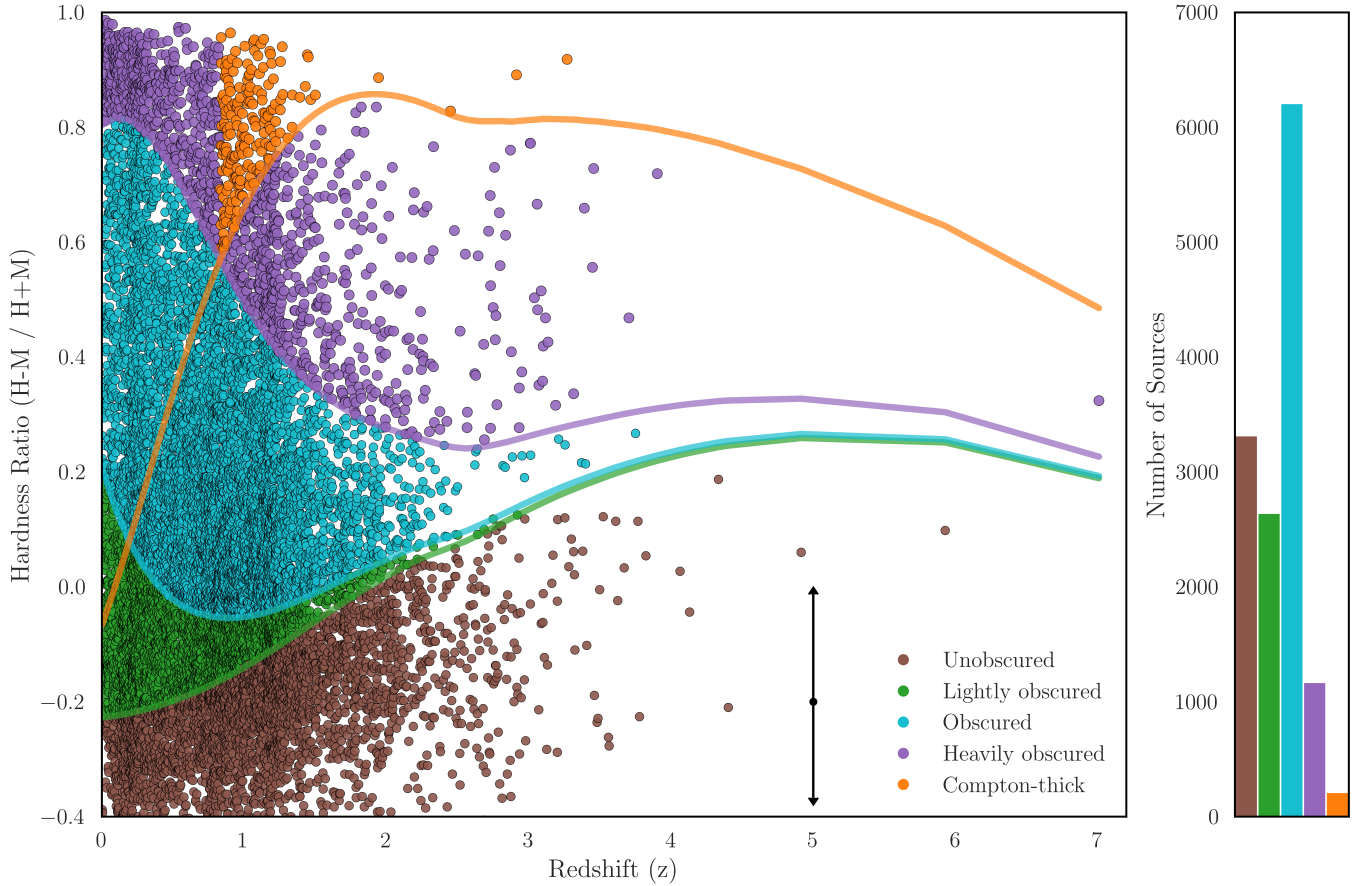


Figure 4. Left: Hardness ratio values for the ExSeSS sample. Color corresponds to the level of obscuration, and will be used consistently within this report. The black point with error bars is the mean error of all sources. Right: Number of sources within each obscuration bin. Note that the image has been cropped for size, and that there are sources within our sample with hardness ratios below -0.4.

The hardness ratio obscuration estimations are shown in Figure 4. There was a turnover around $z = 1$, at which point it became possible to classify objects as Compton-thick. At

lower redshifts, it was not feasible to distinguish between Compton-thick levels of obscuration and lighter levels of obscuration. There was also a degeneracy between $\log N_H = 20 - 22$, as their hardness ratios converge at higher redshifts. The shape of the distribution was roughly what should be expected from a 2–10 keV survey as there were far more sources at moderate and low levels of obscuration than high.

We identified 214 sources with an observed hardness ratio exceeding the threshold for a Compton-thick column density at their respective redshift values. This was clearly an over-estimation of the true number of Compton-thick sources within our sample, as a 2–10 keV survey should be biased against detecting Compton-thick AGN.

Hardness ratios are a very rudimentary method of classifying obscuration and allow for a brief look into the expected distribution of a sample. It should be noted, however, that the results do not indicate the true properties of the sample. The degeneracies between classifications are unaccounted for, and the large errors on each source make the classifications untrustworthy.

4.3 Two-dimensional Hardness Ratios

To improve the previous classifications, we began to include the hardness ratio between the medium and soft bands. To resolve the large errors on each source, we examined only the hard-cut sources discussed in Section 2. These are objects with over 100 photon counts in the hard band. In addition, we partitioned the data into different redshift bins to track the evolution of the expected tracks with cosmic time. We distributed equal numbers of sources in each bin, putting 58 sources in the first eight bins and 55 sources in the highest redshift bin.

We then modified the tracks from the previous section, setting the redshift equal to the average redshift within a bin. We calculated the HR_{HM} and HR_{MS} values of a source at different levels of obscuration from $\log N_H = 20$ to 24. For each source within the hard-cut, we assigned it an obscuration value equivalent to the closest XSPEC value. The results of this are shown in Figure 5.

Hardness ratios are affected by redshift and obscuration. Viewing hardness ratios in the above manner highlights the distinct regions that each level of obscuration occupies. The degeneracies between Compton-thick and unobscured sources only remains at redshifts below $z = 0.1$, where Compton-thick sources are intertwined with unobscured sources. The degeneracies between $\log N_H = 20$ and 21 remained, as these sources have very slight differences in their hardness ratios. Across all redshifts, the populations of heavy, moderate, and light obscuration were distinct from one another.

The observed separation of obscured populations in a two-dimensional hardness ratio space largely motivates Section 6, where we will incorporate these values into a machine learning algorithm. While hardness ratios have helped us identify this behaviour, they are far too simple to use as a reliable estimation of N_H values. We did not account for the error on a point when assigning obscuration based on the nearest point, and assumed values

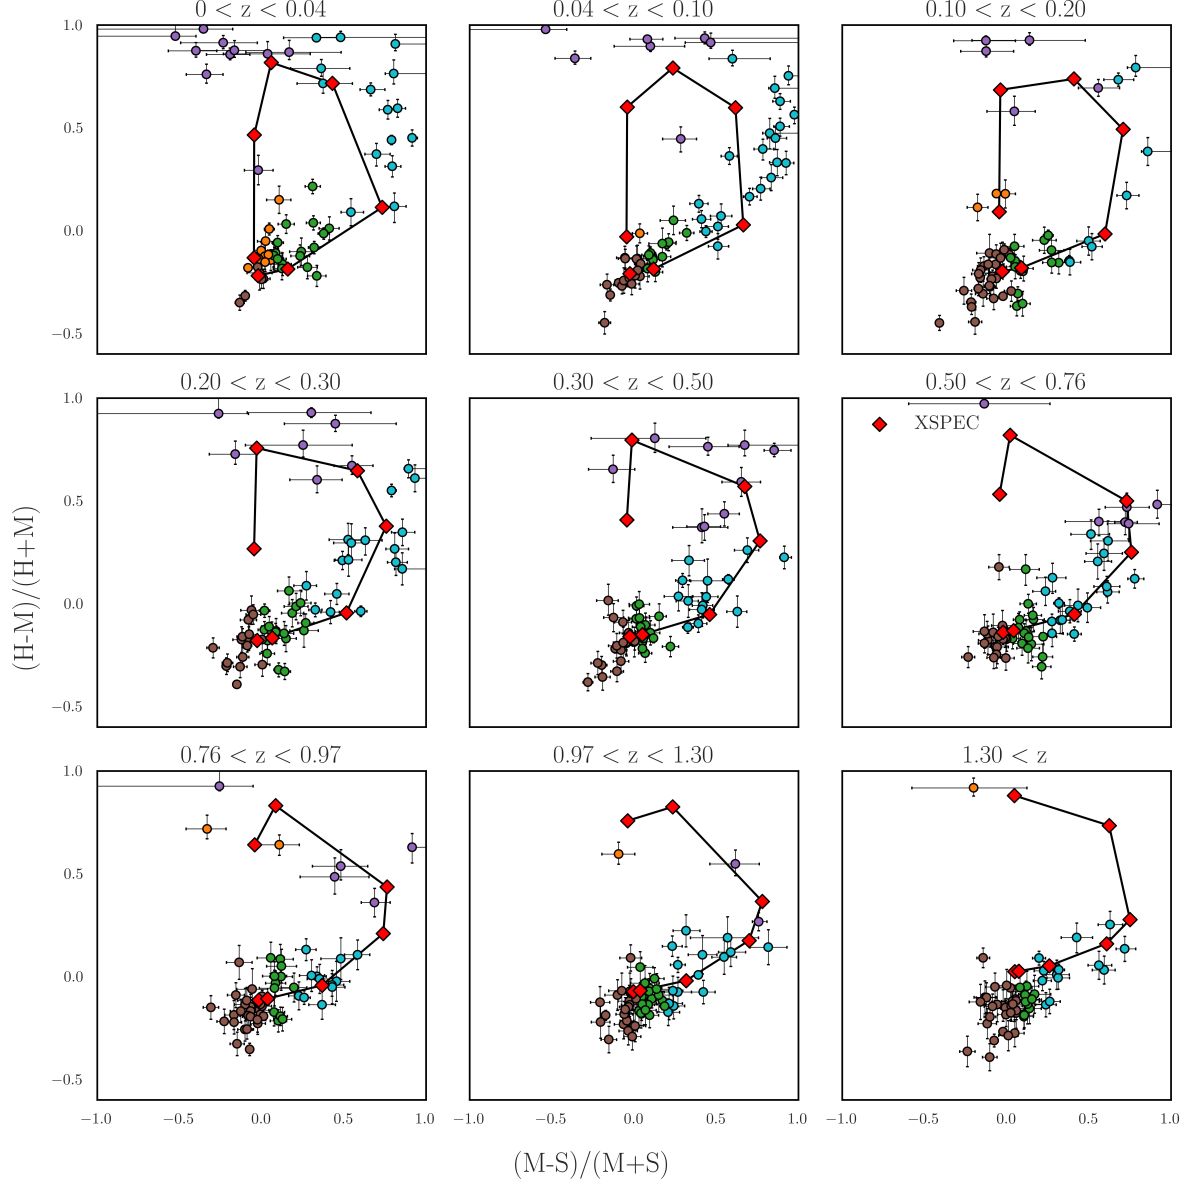


Figure 5. Two-dimensional hardness ratio plots across nine different redshift bins. Red diamonds indicate the HR_{HM} and HR_{MS} values calculated in XSPEC. These values assumed $\log N_H$ values of 20, 21, 22, 22.5, 23, 23.5, and 24, with the unobscured sources typically originating at $HR_{HM} = HR_{MS} = 0$. The circular objects are coloured based on their classification, which depends on the closest XSPEC value. As redshift increases, Compton-thick and unobscured sources become distinguishable. There are also distinct regions populated by each class: unobscured and lightly obscured are located in the lower centre, moderate obscuration on the right, heavy obscuration in the upper centre (and in the upper right at high redshifts), and Compton-thick sources moving from the centre to the upper centre.

for the photon index, galactic absorption, and redshift. While it is easier to distinguish between Compton-thick sources and less obscured sources at lower redshift, we still observe a degeneracy between unobscured and lightly obscured objects. In the next section, we will fully utilize the physical model developed in Section 3 by fitting it to the observed spectra

of our sources, which will improve the effective column densities we calculate.

5 Spectral Fitting of hard-cut Objects

5.1 Performing Spectral Fitting

Hardness ratios are a rough approximation because they only take into consideration the count rates across different energies. For more accurate measurements, one should fit their assumed model to the observed spectrum from a source. This is the most accurate way of determining the effective column density using only X-ray observations, hence its widespread usage (Carroll et al. 2023; Gandhi et al. 2014; Sengupta et al. 2023). The downsides to this technique are that it is heavily contingent on the models used, and is less accurate for sources with low photon counts. We previously discussed in Section 3 that the A15 model should work well until Compton-thick levels of obscuration. Compton-thick sources will make up a very small subset of our sample, so these limitations will be negligible for the majority of our sources. Additionally, we have limited our sample to objects with more than 100 photon counts. While more data would improve the spectral fitting, this should be sufficient given the limitations of our data.

We performed our fits using **XSPEC**, loading in the previously described A15 model as well as the observed spectra of the hard-cut objects. We allowed the effective column density to vary between 0.01 and 250, in units of 10^{22} cm^{-2} , the photon index between 1.6 and 2.2, and the normalization between 1^{-5} and 1.

The column density boundaries are chosen due to the physical properties of absorption and the limitations of our model. The lower limit of N_H is set to be 0.01 as any lower values of obscuration would be indiscernible from galactic absorption. The upper limit of N_H is due to our model breaking down. At higher densities, the rate of Compton scattering will increase within the torus, which will lower the energy of X-ray photons. The A15 model does not account for this, and so we set our upper limit as only slightly higher than Compton-thick. The photon index limits are motivated by other observations, which we previously discussed to have a Gaussian distribution centred around 1.9.

We then set the initial values of several parameters to their known values: the Galactic absorption to the values calculated by Willingale et al. (2013), and the redshifts to the values gathered by Delaney et al. (2023). We also fixed the reflective parameter at -1 to avoid double counting the contribution from accretion-disk reflection. Lastly, we fixed the angle of inclination for cold-matter reflection at 30° following the work of Aird et al. (2015).

We found our optimal fits by minimizing the cash statistic (C) using the default **XSPEC** fitting method, which is a Levenberg–Marquardt algorithm, to perform our fits. We will discuss the details of the C statistic in Section 5.2.1, but note briefly that it is more accurate than χ^2 for estimating parameters at low photon counts. The Levenberg-Marquardt algorithm is a damped least squares fitting procedure and is susceptible to converging at local minima. It

is less optimal for C statistic minimization than for χ^2 minimization (Buchner and Boorman 2024), but we used it since it produced less degenerate results than other fitting methods.

To allow the algorithm to reach a global minimum instead of a local minimum, we repeated fits at a variety of initial conditions, choosing the fit with the lowest C statistic. We froze all parameters other than the column density and performed a fit across a range of N_H values. After this, we saved the N_H value that corresponded to the minimum C statistic and unfroze the photon index and normalization parameters. We then repeated our fitting procedure using the optimal N_H calculated previously and instead varied the photon index over values from 1.8 to 1.9. To calculate the errors, we used the built-in XSPEC function to calculate the 90% confidence range intervals².

The results of two spectral fits are shown in Figure 6, highlighting the differences between a low redshift unobscured source and a high redshift Compton-thick source. There are several key differences between the two, such as the higher count rate of soft photons at low redshift, and the increased contribution of the reflected component at high redshift. Spectra between these redshift ranges and at different levels of obscuration can be found in Appendix B.

We found that our average effective column density was $\log \langle N_H \rangle = 23.5 \pm 1 \text{ cm}^{-2}$, and our average photon index was $\langle \Gamma \rangle = 1.88 \pm 9$. Our column densities were significantly above what was expected, being two orders of magnitude above the results of Fotopoulou et al. (2016), which was an *XMM-Newton* survey. This is alarming as *XMM-Newton* probes the same energy range, and we should therefore have similar values for our N_H distributions. Our photon index, on the other hand, was in solid agreement with typical values (Fotopoulou et al. 2016; Mateos et al. 2008; Nandra et al. 2024). The main reason as to why our average column density was so high could be attributed to the large number of Compton-thick sources found by spectral fitting. We classified 65 objects as Compton-thick following our spectral fitting, 20 as heavily obscured, 108 as obscured, 119 as lightly obscured, and 207 as unobscured. These values seem to overestimate the number of Compton-thick sources within our sample.

It is therefore essential that we validate our spectral fitting through external metrics to ensure that a fit is good before accepting its result. In the next two sections, we will discuss how we can verify our spectral fitting. We will begin by discussing the C statistic and how we can use it to measure whether a spectrum is being overfit or underfit. We will then discuss the relationship between mid-infrared and X-ray luminosities, and how we can use an empirical relationship to remove bad fits.

5.2 Validating our Spectral Fits

5.2.1 The Cash Statistic

Poisson statistics are commonly used over Gaussian statistics in X-ray astronomy due to the lack of abundant data. As a result of this, the quality of a fit is usually determined

²<https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/node79.html>

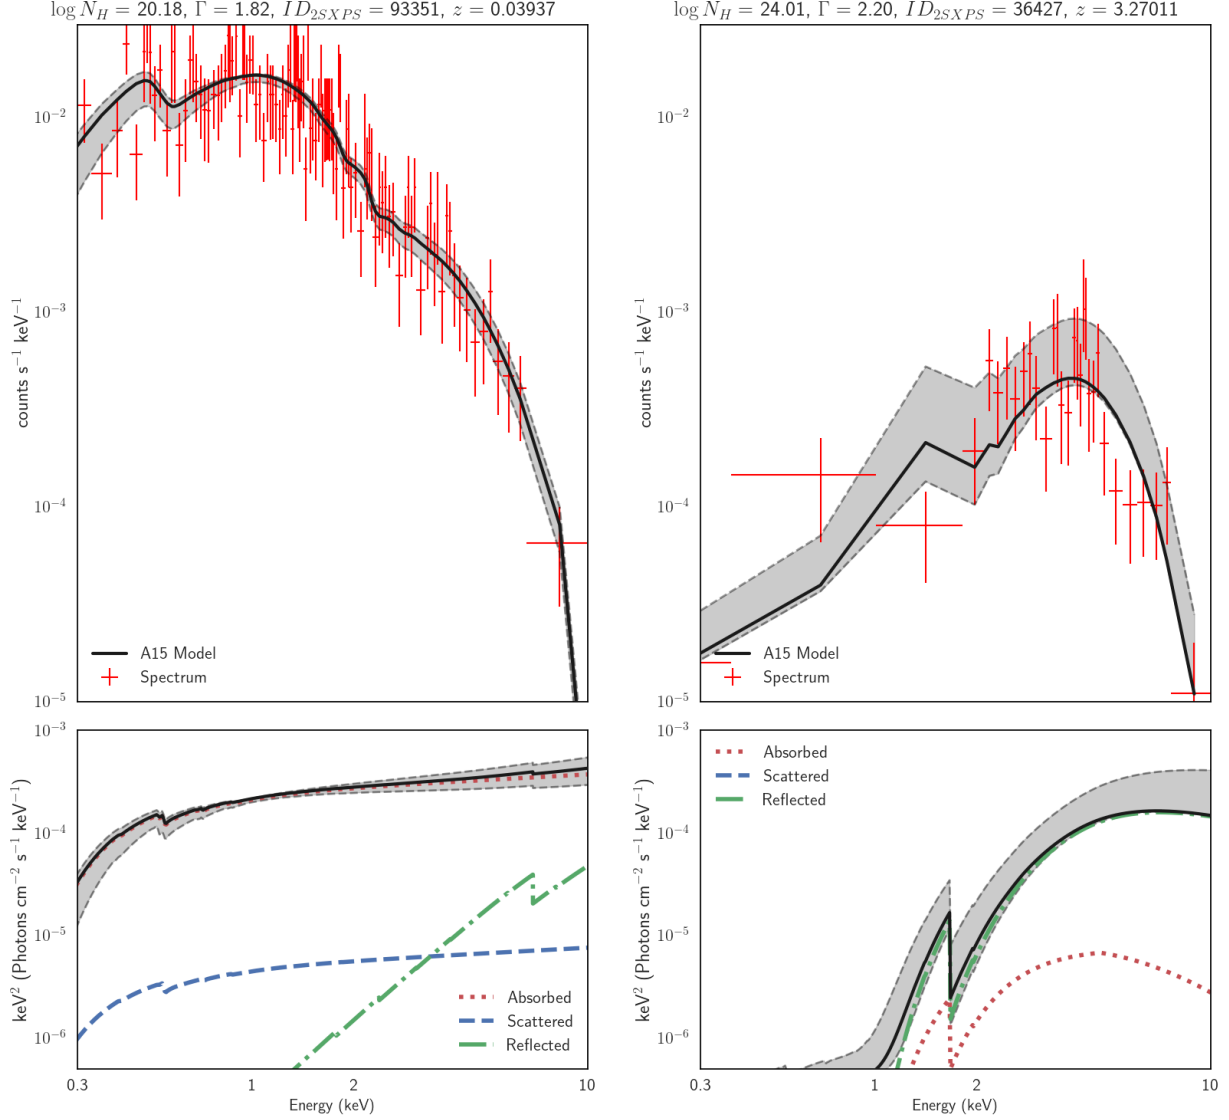


Figure 6. Spectral fitting results for two different X-ray sources. The left side shows a low redshift, unobscured source, while the right side shows a high redshift, Compton-thick source. The top boxes indicate the folded A15 model (black) plotted on top of the observed spectra from 2SXPS (red). The bottom boxes show the relative contribution of all three model components. The low redshift unobscured source is predominantly composed of the absorbed component, e.g. photons which pass directly through the torus. The high redshift Compton-thick source is mainly made up of the reflected component, as the Compton hump at 20–30 keV has been redshifted down.

by the cash statistic, C , rather than through χ^2 , which assumes a Gaussian measurement model. For the purposes of spectral fitting, the cash statistic is well defined by Buchner and Boorman (2024) as:

$$C = \sum_i c_i \log \lambda_i - \lambda_i$$

where i denotes the detector channel, c_i is an integer number of photon counts randomly sampled from a Poisson distribution, and λ_i is the expected number of detected count events.

The C statistic has been designed for parameter estimation since its conception (Cash, W. 1978). The first assumption of the C statistic is that the form of a model is known, but the exact values of its parameters are unknown. This is in line with our A15 model: we are confident in the general form but are looking to find the parameter values. A benefit of the C statistic is that it is unbiased at all photon counts, unlike the χ^2 statistic which has a strong bias at low photon counts and a weaker bias at high photon counts (Humphrey et al. 2009). In addition, the C produces much better fits for parameter estimation than χ^2 (Bonamente 2020; Nousek and Shue 1989). The pitfall of the C statistic is that it does not have an equivalent reduced χ^2 value. It was therefore thought unable to measure the goodness of fit (Humphrey et al. 2009; Nousek and Shue 1989), however, Kaastra (2017) challenged this perception by finding that the C statistic could be used to measure the goodness of fit. Other works, such as Ricci et al. (2017a), used C divided by the degrees of freedom to roughly mimic the value of a reduced χ^2 . They adopted a combination of C statistic and χ^2 , using C statistics for sources with < 200 photon counts and χ^2 in the high-count regime. While their $\frac{C}{dof}$ had a much larger spread than the χ^2_{red} , it was still centred on the expected value of 1. We could have followed their example and used χ^2 statistics for sources with over 200 photon counts, but we opted to use the C statistic for all sources because it calculates better parameter values.

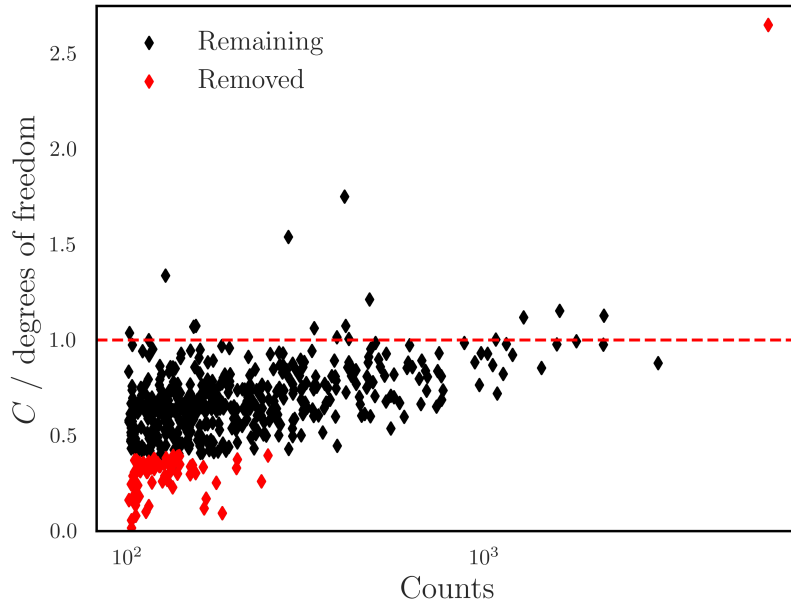


Figure 7. Plot showing our sample of sources remaining after removing overfit and underfit sources. Objects that were removed are shown in red, while those that are kept are shown in black. The dashed red line is the expected reduced χ^2 value for a good fit.

We followed the work of Ricci et al. (2017a) and validated our spectral fitting by dividing the measured C statistic by the degrees of freedom. The expected ratio between the C statistic and the degrees of freedom should be centred on 1, and we show the ratio of all objects in Figure 7. Objects above 1 are underfitted, while objects below 1 are overfitted. The vast

majority of objects had a low number of counts and tended to be overfitted, with only one source suffering from being underfitted. We removed all objects with a $\frac{C}{dof}$ below 0.4 or above 2.5. This removes 68 objects from our spectral-fitted sample.

As our XSPEC fitting was susceptible to converge at local minima, it was not surprising that we had vastly more overfit objects than underfit. The overfit objects were largely due to an improperly fit "soft excess". A soft excess denotes a larger number of soft photons than expected from a power-law spectrum (Crummy et al. 2005). While the overfit objects typically latch onto an abundance of soft photons, the single underfit object was unable to converge at appropriate values. Upon checking the coordinates of this source on SIMBAD, we find that it is also known as QSO B1207+39 and has a spectroscopic redshift of 0.617 as determined by Abazajian et al. (2009). This discrepancy in redshift could explain the underfit nature of this source.

There are several reasons as to why we had a non-negligible population of poorly fit objects. These include the assumptions of our model, the number of photon counts available in our sources, and the fitting algorithm we used. In Ricci et al. (2017a), they fit 9 different models to their observed spectra. One method of improving our fitting would be to use a range of models, some simpler and others more complex. Another way of improving fits would be to use a global solver instead of a local solver, such as Bayesian X-ray Analysis (BXA) (Buchner et al. 2014). We did not use this global solver due to encountering a large number of degenerate solutions but suggest that future spectral fitting should try to use this solver. In addition, we have assumed in this report that the C statistic divided by its degrees of freedom can be a good measure of fit quality. This assumption might not be accurate, which may indicate the high number of poorly fit objects. Another reason might be that the redshift estimates within *ExSeSS* are inaccurate. Our sources often disagreed with the redshifts listed on SIMBAD, which indicates that our estimations might not be accurate. We note that resolving this discrepancy is a task of high importance for any future work involving the *ExSeSS* catalogue.

5.2.2 The X-ray and Mid-Infrared Luminosity Correlation

Within this report, we have only made use of X-ray observations, which has enabled us to estimate obscuration through the usage of hardness ratios and even directly calculate them through spectral fitting. While this has largely been beneficial, it would be detrimental if we did not correct our sample to remove $\log N_H$ estimations for poorly fitted objects. Spectral fitting has a highly degenerate solution space, which means its results usually require external verification before they can be met without skepticism. Popular methods include using a multitude of XSPEC models or supplementing data with multiwavelength catalogues.

In this section, we will examine the relationship between X-ray and Mid-Infrared (MIR) luminosities. We will first discuss where infrared emission comes from and how it is related to X-ray emission, and will then calculate the intrinsic luminosities of our sources by applying a series of corrections. We end this section by removing our sources which do not lie on the X-ray—MIR relation, to improve the overall quality of our spectral fits.

The infrared emission largely stems from the dusty torus, but also in part by the outskirts of the accretion disk. A fuller picture of AGN can therefore be constructed by comparing the infrared emissions born by the torus with the X-ray emissions that penetrate it. It has previously been reported that there is a strong correlation between intrinsic X-ray and infrared luminosities for unobscured AGN, particularly between 2-10 keV and 6 μm (Lutz et al. 2004). This has been explored extensively by Chen et al. (2017), Fiore et al. (2009), Gandhi et al. (2014), and Stern (2015) each deriving unique empirical solutions (hereafter referred to as C17, F09, G14, S15, respectively). The relations differ the most at higher luminosities due to a lack of data and hence problems of extrapolation, but ultimately show a tight relationship between X-ray and MIR luminosities.

The above relations were calculated assuming unobscured AGN, but an approximate relation for CT-AGN has been used by Lansbury et al. (2017) where they scaled the X-ray luminosity down by the typical amount of suppression. In this manner, we can investigate our Compton-thick sources to see if their observed luminosities lie on the expected tracks of a $N_H = 10^{24} \text{ cm}^{-2}$ source. More importantly, we can verify the results of our spectral fitting by mandating that spectral fitting results be kept only if a source’s intrinsic X-ray luminosity lies within one order of magnitude of the MIR—X-RAY relation.

We calculated the observed and intrinsic X-ray luminosities of our sources using XSPEC, and separately interpolated the 4.6 and 15 μm WISE observations to get the intrinsic 6 μm luminosity, applying a K-correction to account for the different redshift of our sources. The full derivation of the 6 μm luminosity can be found in Appendix A. We note that we made several assumptions in this derivation which may differ from the actual intrinsic 6 μm luminosity.

Class	Spectral Fitting	CSTAT Correction	MIR Relation
Compton-thick	65	61	6
Heavily obscured	20	13	13
Obscured	108	82	80
Lightly obscured	119	99	92
Unobscured	207	196	187
TOTAL	519	451	378

Table 2. Training sample size within each classification after initial spectral fitting and subsequent verification. The spectral fitting column shows the classifications of the 519 objects selected by mandating at least 100 photon counts in the hard band. The CSTAT column shows the sizes following the removal of outlier $\frac{C}{\text{dof}}$ values, while the MIR relation denotes the classes remaining within one order of magnitude of the MIR—X-ray relation.

We then compared the intrinsic X-ray and infrared luminosities to the expected tracks produced by F09, G09, S15, and C17, and removed all sources that did not lie within one order of magnitude of the predictions. We show in Figure 8 the intrinsic X-ray luminosity versus the intrinsic infrared luminosity after correcting for absorption. All remaining Compton-thick candidates have an observed luminosity around the value expected for an $N_H = 10^{24} \text{ cm}^{-2}$ source, indicated by the black arrow. We see similar success with our less obscured sources, which fit the unobscured relation with a low amount of spread. While this correction has

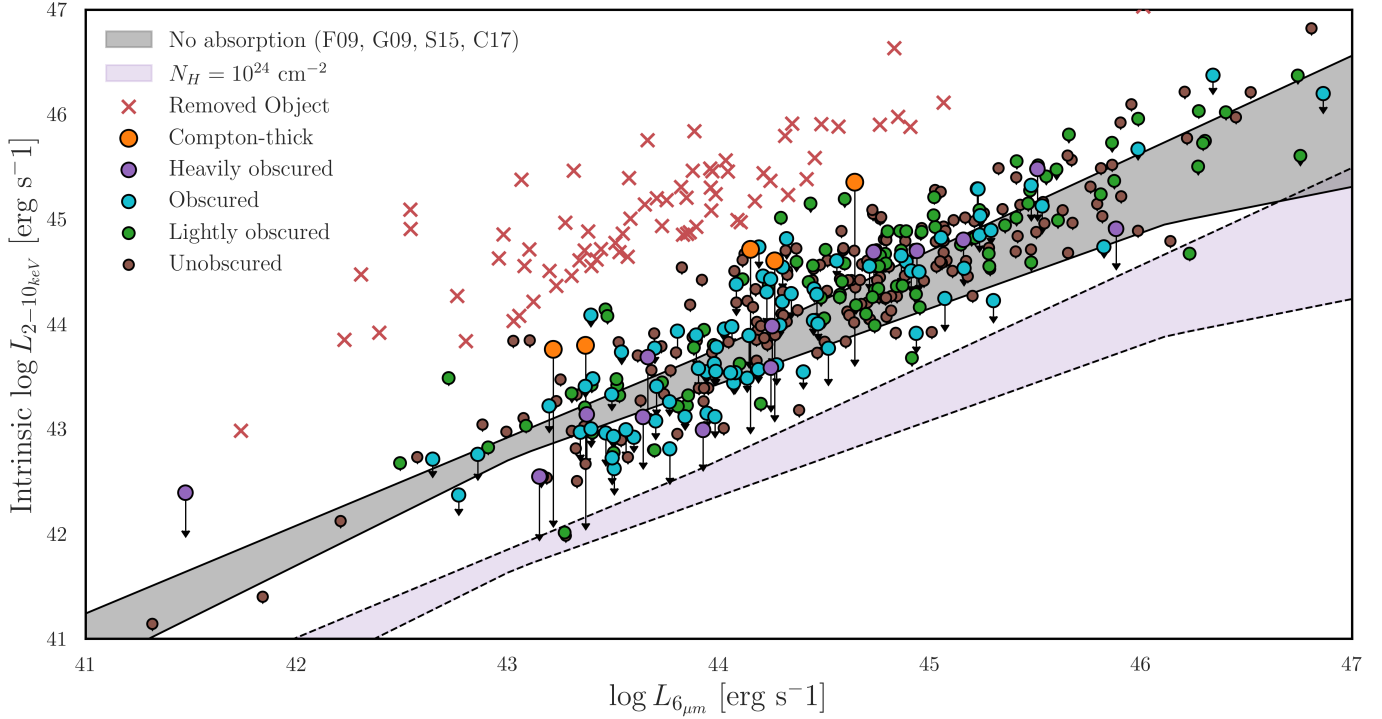


Figure 8. Evaluating the empirical relationship between X-ray (2–10 keV) and mid-infrared (6 μm) luminosities. These luminosities have been corrected to their intrinsic values. The downward arrows indicate the observed luminosities of each source before correcting for obscuration. The grey shaded region indicates an empirical relationship for unobscured sources (Chen et al. 2017; Fiore et al. 2009; Gandhi et al. 2014; Stern 2015), while the purple shaded region shows an approximate relation for obscured sources. The red crosses indicate objects which are an order of magnitude above the relation, and are thus removed from our training sample. We do not show error bars as they are negligible and detract from the observed–intrinsic correction. This figure highlights the excess of Compton-thick predictions in our spectral fitting, which have now been removed. It also shows the difficulty in distinguishing between Compton-thick and unobscured sources, as there are unobscured sources in the $N_H = 10^{24} \text{ cm}^{-2}$ region.

removed a lot of immediately bad fits from our sample, we note that it reveals the difficulty our model has in distinguishing between unobscured and Compton-thick sources. This can be seen by the unobscured sources within the $N_H = 10^{24} \text{ cm}^{-2}$ track as their obscuration has likely been underestimated by our spectral fitting.

We show in Table 2 the number of sources within each label following the results of spectral fitting, C statistic validation, and the MIR relation. Upon recalculating the average values for N_H and Γ , we find that $\log \langle N_H \rangle = 22.7 \pm 2 \text{ cm}^{-2}$ and $\langle \Gamma \rangle = 1.87 \pm 9$. While we have lowered our values to be more in line with other observations, we are still higher by an order of magnitude. This is likely due to the small size of our hard-cut sample, as an average will be biased by our most obscured sources.

We note here that while we have validated our spectral fits using the reduced C statistic and the MIR-XRAY luminosity relation, we have not necessarily determined the most accurate fits. We have made a significant amount of assumptions in this section: our physical model within XSPEC, using the reduced C statistic to validate fits, and the calculation of our intrinsic

luminosities. To bring more certainty to our results it would be beneficial to refit our spectra using a different physical model, perhaps with a different fitting algorithm such as BXA (Buchner et al. 2014), and compare the results of both methods. Despite these assumptions, the spectral fitting appears to be visually correct (see Figure 6 and Appendix B), and so we leave these corrections to future work. In the next section, we will take the remaining 384 objects and use them to train a K-Nearest Neighbours algorithm to generalize our spectral fitting of high-count objects onto our low-count objects.

6 Generalizing Results to Low-Count Data Using a KNN Algorithm

Having performed spectral fits on the hard-cut selection, we now wanted to generalize the obscuration calculations of our high-count data onto our low-count data. We did this using a K-Nearest Neighbours (KNN) machine learning algorithm, combining the results of our spectral fitting with our 2-dimensional hardness ratios.

6.1 KNN Algorithms

K-Nearest Neighbours is a supervised learning algorithm that examines the distance between points and can be used for classification or regression, where classification is a discrete result and regression is continuous. We will be using it for classification purposes as this is far simpler than predicting N_H values. In KNN, K is a user-chosen value that designates the number of neighbours to examine. For a K value of 1, the assignment would be the closest point, while a K value of 3 means that the assignment would be based on the three nearest points. Odd values of K are typically preferred over even values to avoid ties in voting. The meaning of supervised in machine learning is that the model is trained on an initial set of data that contains features and labels, and then tested on a separate set of data that only contains features. Features are measurable properties and labels are predicted qualities.

While this algorithm is conceptually simple, it has seen a large amount of usage across disciplines due to its effectiveness. In an astrophysical context, KNN has previously been used to classify between stars, galaxies, and AGN (Li et al. 2008), and more recently to distinguish between different types of AGN (Cooper et al. 2023). It has not yet been used to estimate the obscuration of AGN. Our goal in using it within this report is to quantify the obscuration of our low-count sample using only training data from our most secure spectral fits. A successful application of a KNN algorithm in this report would strongly motivate further usage, especially as more data becomes available.

6.2 Model Evaluation

As with most other types of machine learning, KNN is strongest when used on well-separated data. We have seen previously in Figure 5 that distinct populations crop up when viewing

the obscuration as a function of hardness ratios and redshift. For this reason, we defined the features of our KNN algorithm to be the hardness ratios HR_{HM} , HR_{MS} , and the redshift z . The training sample consisted of the 386 remaining sources following the X-ray—MIR relation cut, while the test sample was made up of the 13,575 low-count objects. We set the training labels to be the class of each object, ranging from unobscured to Compton-thick. Then we assigned weights to the algorithm based on distance, where neighbours further away from a test point contributed less to the classification than a closer neighbour. All of this was done using the Python library `scikit`.

Our next goal was to determine the optimal value of K for our classifier through the usage of metrics. A simple metric is accuracy, which just measures the percentage of correct classifications. In datasets with large class imbalances, such as our own, the accuracy metric becomes a poor statistic for optimization. This is because the dominant class will skew the metric. We instead opted to use the F-measure, which is more sensitive to class imbalances. The F-measure is also biased by class imbalances, but to a lesser extent than accuracy. We will cover the basics of the F-measure below, but note that a more thorough review can be found in Christen et al. (2023).

The F-measure is dependent on two values known as the precision and recall. Precision is the “False Positive Rate” and recall is the “False Negative Rate”.

Precision captures how relevant the correct classifications are, and is given by:

$$Precision = \frac{TP}{TP + FP}$$

where TP stands for the number of true positive classifications and FP stands for the number of false positive classifications. As an example, we now consider a trio of three sources: one Compton-thick, and two unobscured. We first decide that we want to calculate the TP and FP values for the unobscured class. If both unobscured sources are classified as unobscured, then the number of true positive sources is 2. If the Compton-thick object is classified as unobscured, then the number of false positive sources is 1. We note that in this example we have specified the class of interest as unobscured, but the values of TP and FP are dependent on which class is being examined.

The other value, recall, captures how well a class is retrieved. it is given by:

$$Recall = \frac{TP}{TP + FN}$$

where FN is the number of false negative classifications. If we tweak the previous example so that one of the unobscured sources was instead classified as any other class, then our values would be $TP = 1$, $FP = 1$, and $FN = 1$.

This brings us to the F-measure, which is the harmonic mean of precision and recall:

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} = \frac{2TP}{2TP + FP + FN}$$

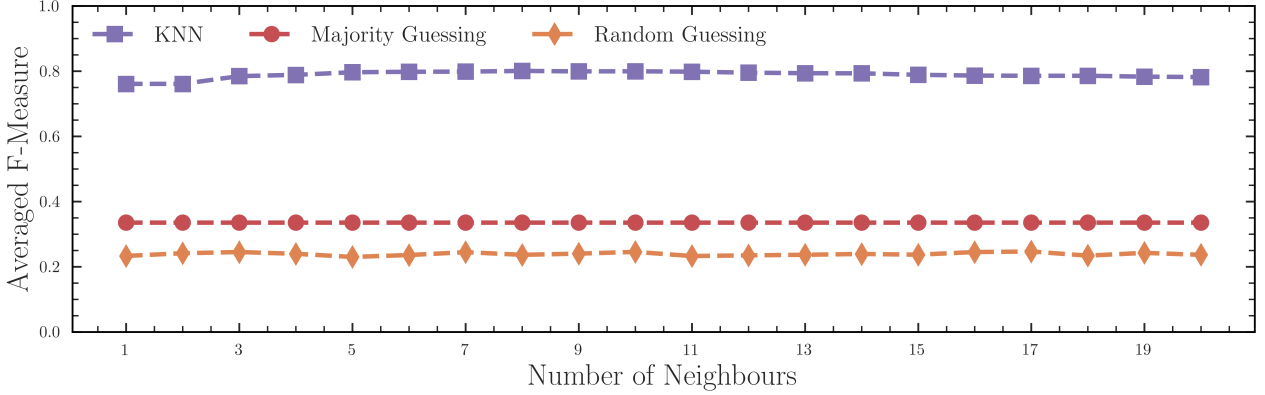


Figure 9. The averaged F-measure over 100 trials as a function of the number of neighbours. We calculated a much higher $F1$ value for our KNN classifier than for both dummy classifiers, which indicates that our classifier is better than baseline estimates, e.g. guessing randomly or choosing only the most abundant class. The $F1$ value is flat across all values of K , which indicates that the data is noisy.

To help choose a value for K , we randomly split the 378 training sample objects into smaller training and validation samples. We used 80% of the sample for training and 20% for validating, choosing these percentages due to the small size of our initial sample. We would ideally have split the data evenly, but this would not have been able to capture the data well due to our class imbalances. In even better circumstances we would have used an external set of validation data, as our internal metrics are susceptible to systemic biases. We then calculated the F-measures for each class at K values ranging from 1 to 20. The F-measures between each class were then averaged in order to produce an F-measure for the data as a whole. This average was weighted such that the more populated classes had more bearing than the less populated classes. We also included two “dummy classifiers”, which intentionally made poor classifications. One of these classifiers assigned obscurity at random, while the other would only classify objects as unobserved. We discovered that the previous F-measures were volatile, and suffered great variance depending on the distribution of training and validation sources. To account for these fluctuations, we repeated the above method 100 times and averaged the results at each K value. The results of this are shown in 9.

The results of our KNN classifier were visibly an improvement from either randomly choosing classes or classifying objects only as unobserved. There were five classes within our sample, so it was reassuring that random guessing had a typical $F1$ value around 20%. The $F1$ value of the only-unobserved classifier was below 50%, which was roughly what should be expected. The KNN algorithm plateaued around $F1 = 0.8$, but there was no definitive K value that worked best. This suggests that data is noisy due to heavily imbalanced classes and that there is negligible improvement when choosing a K greater than 5. We chose to set $K = 9$ as it would be more biased against classifying sources as Compton-thick than $K = 5$.

We attribute the difficulty in constraining K to the class imbalances and small sample size of our training data. This method might be more successful for higher energy wide-field

surveys such as *Swift* BAT, which would be able to find large populations at all levels of obscuration. It could also be easier to constrain K through the usage of other metrics, such as AUC-ROC, which may be better than the F-measure at dealing with class imbalances.

6.3 Generalisation Results

We applied the KNN algorithm to the 13,176 low-count sources which we have ignored since Section 2. The predicted classifications are shown Figure 10. We note here that we show the figure in redshift bins to improve visibility: the KNN algorithm was not trained in separate redshift bins.

From Figure 10 we see that the algorithm struggles to identify Compton-thick sources, but succeeds in identifying large populations of less obscured sources. Notably, the algorithm can classify many sources as heavily obscured despite the number of Compton-thick and heavily obscured sources being roughly equal within the training sample (6 and 13, respectively). This can be attributed to the fact that Compton-thick and unobscured AGN have hardness ratios around $HR_{HM} \approx HR_{MS} \approx 0$, while heavily obscured typically have hardness ratios around $HR_{HM} \approx 1$ and $HR_{MS} \approx -0.2$. It is therefore easier for the KNN algorithm to classify objects as heavily obscured because they are better separated from other labels than Compton-thick sources.

The figure also shows distinct populations of unobscured, lightly obscured, and obscured objects at all redshifts. Heavily obscured objects only make up a distinct group at redshifts below 1.12. As the number of low redshift objects vastly outnumbers the number of high redshift objects, the KNN algorithm does not seem to appropriately account for the evolution of hardness ratios as a function of redshift. This is particularly egregious when looking at the training sample objects at higher redshifts. The heavily obscured objects in the $0.79 < z < 1.12$ bin lie within the obscured population, and the single Compton-thick source in the $z > 1.12$ bin lies within the heavily obscured population.

Although the KNN algorithm lacks training data of more obscured sources at higher redshift, it can identify large populations at every level of obscuration except for Compton-thick. This is not necessarily problematic as 2–10 keV observations are biased against detecting Compton-thick AGN, and it is therefore desirable for our KNN algorithm to predict a small number of them. It is beneficial that the KNN algorithm can recover a large population of $\log N_H = 23$ sources given that the training sample contained only 13 heavily obscured sources, and this success can be attributed to the well-separated nature of the data. The biggest pitfall of the KNN algorithm is that it is unable to properly account for the redshift evolution of an obscured source’s hardness ratios. The majority of the data is collected at redshifts below $z < 1$, which suggests that classifications will become less accurate at higher redshifts. The number of Compton-thick sources in the Universe are thought to increase with redshift (Buchner et al. 2015; Ueda et al. 2014), so perhaps the KNN algorithm is underpredicting obscuration at higher redshifts. This could be resolved by acquiring more training data or by training multiple KNN models at different redshifts. If one were to include

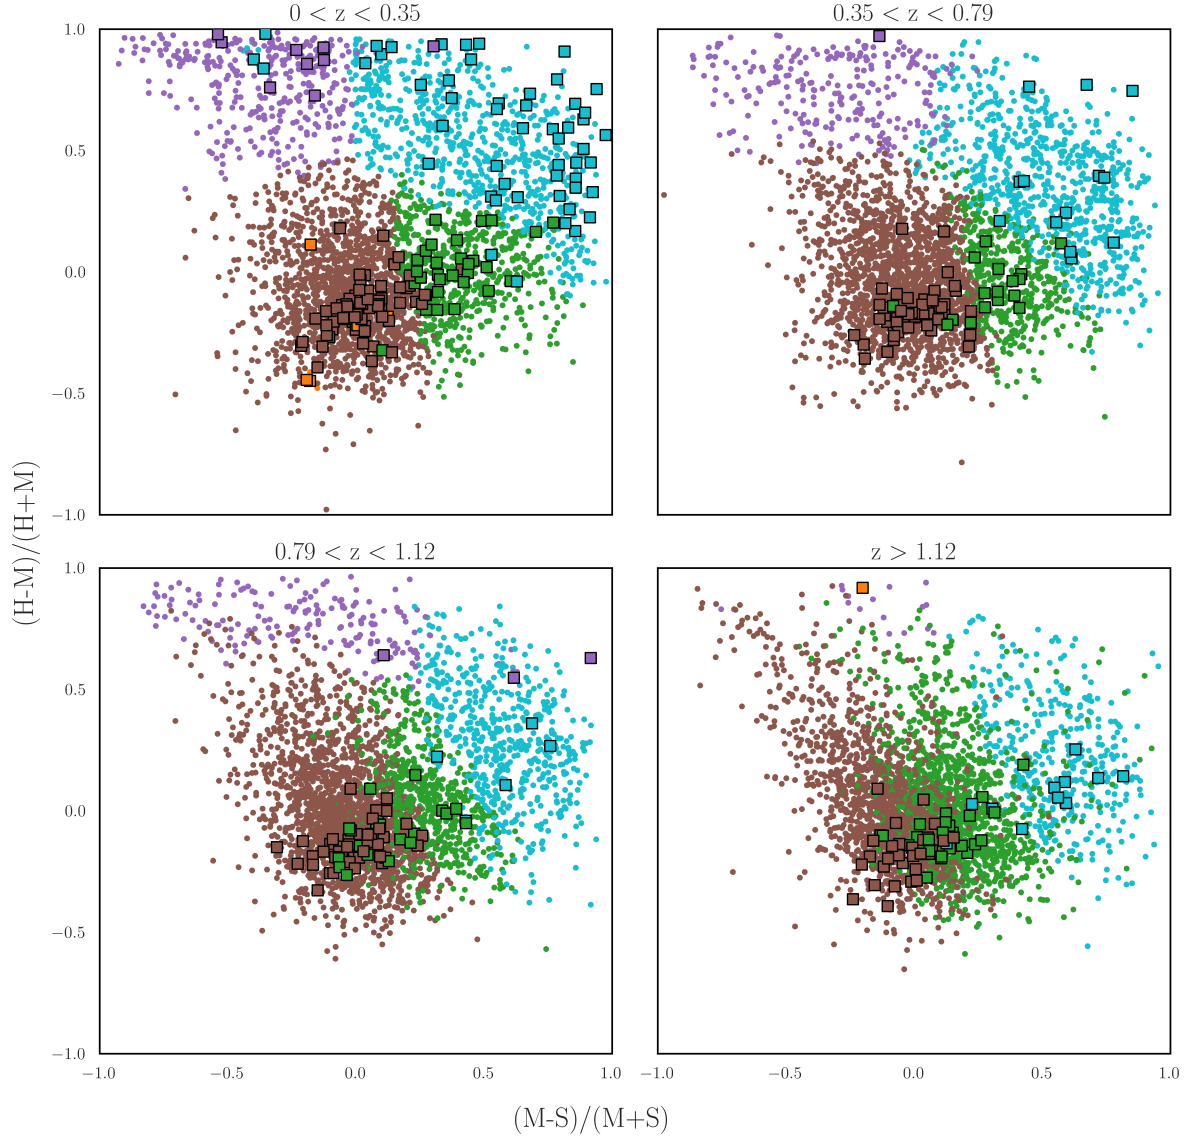


Figure 10. The results of our KNN algorithm. The training sample is shown by squares, the test sample by circles. Compton-thick objects are shown in orange, heavily obscured in purple, obscured in cyan, lightly obscured in green, and unobscured in brown. Although we show this in separate redshift bins, we stress that the KNN algorithm was not trained in this manner. We observe that Compton-thick classifications only occur at low redshift, where we have two Compton-thick training sources next to each other at $HR_{HM} \approx -0.5$, $HR_{MS} \approx -0.2$. The abundance of unobscured, lightly obscured, and obscured objects within our training sample leads to large number of classifications within our test sample. Notably, we are able to recover a large population of heavily obscured sources across redshift values ($z = 0 - 1.12$) due to the well-separated nature of heavily obscured sources in a two-dimensional hardness ratio parameter space.

more features such as the X-ray and infrared luminosities, they could perform principal component analysis to maximize the variance. Training a KNN algorithm in this manner would include more physical data while keeping dimensionality low.

It is important to note that the predictions of our KNN algorithm are unchecked. These

are low-count sources, which means that fitting their spectra is likely to lead to poorly constrained and highly degenerate solutions. The number of sources is also very computationally demanding. We do, however, perform spectral fits on the three KNN-predicted Compton-thick sources, and find that one of the three was classified as Compton-thick. The other two were classified as unobscured. The spectra of these sources are shown with the others in Appendix B. We stress that these fits are very poorly constrained, and should not be taken as an indication that the KNN algorithm is a poor metric for estimating N_H . The quality of fits become so dismal in the low-count regime that hardness ratios become a better metric for quantifying obscuration. It is, however, a serious limitation of the KNN algorithm that its results cannot be verified. One way of resolving this would be to perform spectral fits on our objects with more than 80 photon counts, and then compare their column densities to our KNN prediction. Another method would be to train a KNN algorithm on already completed datasets, e.g. the *Swift* BAT 70-month catalog (Ricci et al. 2017a) or *XXL* survey (Fotopoulou et al. 2016), and see if it can predict accurately.

6.4 Comparing Obscured Fractions to Other Surveys

We show in Figure 11 the observed fraction of $\log N_H$ within our sample, as well as the fraction of sources from several other surveys. The energy ranges and surveys we compare to are: the eFEDS hard selected sample (2.3–5 keV) from Nandra et al. (2024), the *XXL* survey (2–10 keV) from Fotopoulou et al. (2016), a hard-band selected *NuSTAR* survey (8–24 keV) Zappacosta et al. (2018), and the 70-month *Swift* BAT survey (14–150 keV) by Ricci et al. (2017a). In addition to these observations, we include the intrinsic fraction stated by Ricci et al. (2017a). Henceforth we will refer to these observations as N24, F16, Z18, and R17, respectively.

Each observation within Figure 11 covers a slightly different energy range except for our work and *XMM-Newton*. It is immediately clear that incredibly soft surveys like N24 are biased against detecting heavily obscured sources. To contrast this, higher energy surveys like Z18 observe a greatly reduced fraction of unobscured sources and a much higher fraction of heavily obscured sources, but identify larger fractions of Compton-thick sources. R17 observes a higher fraction of heavily obscured and Compton-thick sources than low energy observations, but finds a much larger unobscured fraction than Z18. The intrinsic fraction produced within R17 reveals that all of these energy ranges are biased against detecting Compton-thick AGN. When we compare our KNN predictions to F16, the only survey that matches our energy range, we find broad agreement across all levels of obscuration. In general, we tend to see more moderately obscured sources than softer surveys and less heavily obscured sources than harder surveys.

When comparing to N24, which examines AGN within 0.2–5 keV, we only find agreement in the $\log N_H = 23$ bin. This could be explained by their dominating fraction of unobscured sources, which is likely due to their inability to probe harder energies. It is therefore reassuring that we find higher fractions of obscured sources because we are probing higher energies.

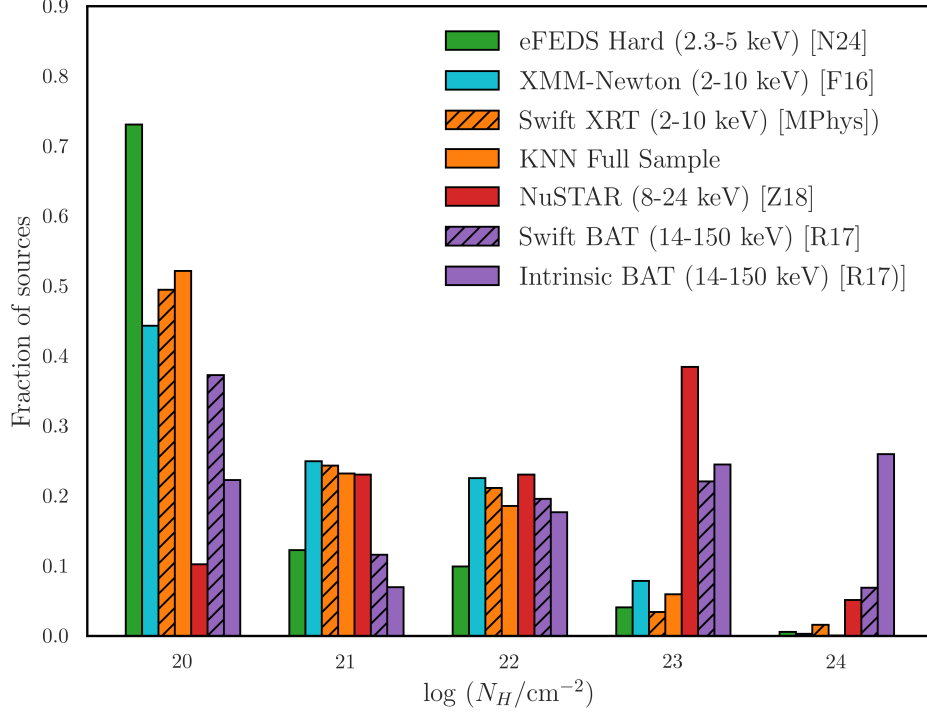


Figure 11. Observed $\log N_H$ fractions. The bar plots are sorted by their energy ranges in ascending order. The green bar indicates *eROSITA* observations, cyan indicates *XMM-Newton*, orange with dashes representing our 386 spectral fitted *Swift* XRT sources, orange without dashes being our full 13,151 sources following the application of our KNN algorithm, red being *NuSTAR* observations, purple with dashes being observed *Swift* BAT observations, and purple without dashes being the intrinsic fractions predicted by Ricci et al. (2017).

We find a slightly higher fraction of CT-AGN than F16 when looking at our hard-cut sample, but this is resolved when generalizing our KNN algorithm to the low-count sources. We estimated lower fractions of heavily obscured and Compton-thick sources than higher energy observations by Z18 and R17, which is also in line with expectations. The intrinsic fraction of obscured AGN, calculated within R17, which corrects the observed fraction of their work, shows that X-ray observations are biased against detecting Compton-thick AGN.

We have estimated N_H by classification, not regression. This means we did not predict the exact value of N_H , but instead determined obscuration using classes such as unobscured, lightly obscured, and so forth. We roughly assumed that each class corresponded to its column density (e.g. $\log N_H = 10^{20} \text{ cm}^{-2}$ for an unobscured source) in order to calculate an updated $\log \langle N_H \rangle$ value. We find that $\log \langle N_H \rangle \approx 21.9$, which is within an order of magnitude of the *XXL* survey (Fotopoulou et al. 2016). This is a rough assumption, but shows that validating our spectral fits and using our KNN algorithm has brought our results in-line with expected values.

7 Comparison to AGN Population Synthesis Models

7.1 AGN Population Synthesis Models

In the previous sections, we outlined our methods for determining the obscuration of *ExSeSS* sources. We will build off of these results by calculating the number counts of our sample, binning our objects by their obscuration, and then comparing these values to the number counts predicted by AGN population synthesis models.

AGN synthesis models compute the number density of AGN as a function of their luminosity and redshift in order to fix discrepancies between observed AGN populations and the Cosmic X-ray Background (CXB). Low energy X-ray telescopes are biased against observing Compton-thick AGN, so these observations find a fraction of observed Compton-thick AGN significantly lower than unobscured sources (Mostafa et al. 2023; Nandra et al. 2024; Signorini et al. 2023). Higher energy X-ray telescopes have worse angular resolutions and therefore have fewer sources in their surveys, but find a slightly increased fraction of Compton-thick sources relative to the soft-probing X-ray telescopes (Ricci et al. 2017a; Zappacosta et al. 2018). Both of these observations observe a lower fraction of Compton-thick AGN than is expected from analysis of the cosmic X-ray background.

The CXB, first discovered by Giacconi et al. (1962), is a nearly isotropic presence of X-ray energies peaking at around 20-30 keV (Gilli et al. 2007; Moretti et al. 2009). The main contributions to the CXB are widely thought to be from AGN as their X-ray emission dominates over the intracluster medium surrounding galaxy clusters, which is an X-ray emitting collection of ionized hydrogen (Barger et al. 2005; Szokoly et al. 2004). For the CXB to peak at 20-30 keV, there must be a significant population of Compton-thick AGN in order to keep the soft energy contribution in line with observed values while still allowing harder photons to contribute to the CXB (Fabian 1999). The goals of AGN synthesis models are therefore to match a number density of AGN to the observed X-ray number counts, the observed fraction of Compton-thick sources, and the CXB spectrum. To avoid confusion we note that number counts are not related to the number of photons from a source, but are defined as the number of sources N above a flux value S per deg^{-2} .

We will be comparing our observed number counts to the predictions of two models: Gilli et al. (2007) and Ueda et al. (2014) (hereafter G07 and U14). The G07 model is predominantly composed of deep *Chandra* and *XMM-Newton* observations but also includes shallower surveys across wider areas. It has been updated since its original publication to properly account for higher redshift sources. The U14 model, in contrast, has more observations at its disposal, including *ROSAT*, *ASCA*, *Swift* BAT, and MAXI. A more recent model which includes *NuSTAR* observations has been produced by Ananna et al. (2019), but it is not yet available for public access. There are several other models we could have used, e.g. Akylas et al. (2012) or Ballantyne et al. (2006), but we note that these models precede the G07 and U14 models by roughly a year, and are therefore made with slightly less observational data.

7.2 Calculating Number Counts

The U14 source count predictions were pre-calculated values available upon request. These assumed a redshift range of $z = 0.002-5$, X-ray luminosities of $L_X = 41-47$, and obscuration values of $\log N_H = 20-26$, $\log N_H = 20-22$, $\log N_H = 22-24$, and $\log N_H = 24-26$. We then gathered the G07 counts using their online calculator³, assuming the same luminosity range as in the U14 case, but allowing the redshift to vary depending on the maximum redshift within our observed data. The G07 model was calculated from $z = 0$ to 7 in the total and 20-22 bins, $z = 0$ to 3 in the 22-24 bin, and $z = 0$ to 1 in the 24-26 bin.

Before we calculated the number of source counts within our sample, we removed data-poor objects and binned our sample. We removed all objects with a $\log S_{2-10keV} < -14$, as the areas of these objects are too small to accurately discern. We then filtered our sample into four obscuration groups: $\log N_H = 20-26$, $\log N_H = 20-22$, $\log N_H = 22-24$, and $\log N_H = 24-26$. We partitioned these groups further by splitting them into 14 evenly spaced flux bins from $\log S_{2-10keV} = 14$ to $\log S_{2-10keV} = 10.5$. We did not require that these bins contained a minimum number of objects because our small sample of Compton-thick candidates was spread across a wide flux range. We then calculated the integrated source counts as defined by Mateos et al. (2008):

$$N(> S_j) = \sum_{i=1}^{i=M} \frac{1}{\Omega_i}$$

with Poissonian errors:

$$N(> S_j)_{err} = \frac{N(> S_j)}{\sqrt{M}}$$

where $N(> S_j)$ is the number of sources per unit sky area with flux higher than S_j , S_j is the minimum flux value allowed within a bin, and M is the total number of sources with $S_i > S_j$. Note that S_j is the minimum flux value allowed within a bin and not the minimum flux value within a bin. This allowed us to calculate the percentage of number counts contained in one obscuration group relative to the total.

7.3 Comparing to the Gilli and Ueda models

The results of our source counts are shown in Figure 12, and the residuals in Figure 13. We strongly agree with the G07 and U14 tracks at all levels of obscuration at medium fluxes ($S_{2-10keV} = 10^{-12.75} - 10^{-12}$). We are consistently below the predicted number counts at the faintest fluxes ($S_{2-10keV} < 10^{-12.75}$), however it should be noted that the *Chandra* and *XMM-Newton* surveys used to create the G07 and U14 models probe deeper than the *ExSeSS* catalogue, and are hence more sensitive to faint objects. At brighter fluxes, we observe an excess of number counts compared to both models, most extremely for our Compton-thick sources. It has been suggested by Delaney et al. (2023) that the bright excess may indicate an

³<http://www.bo.astro.it/~gilli/counts.html>

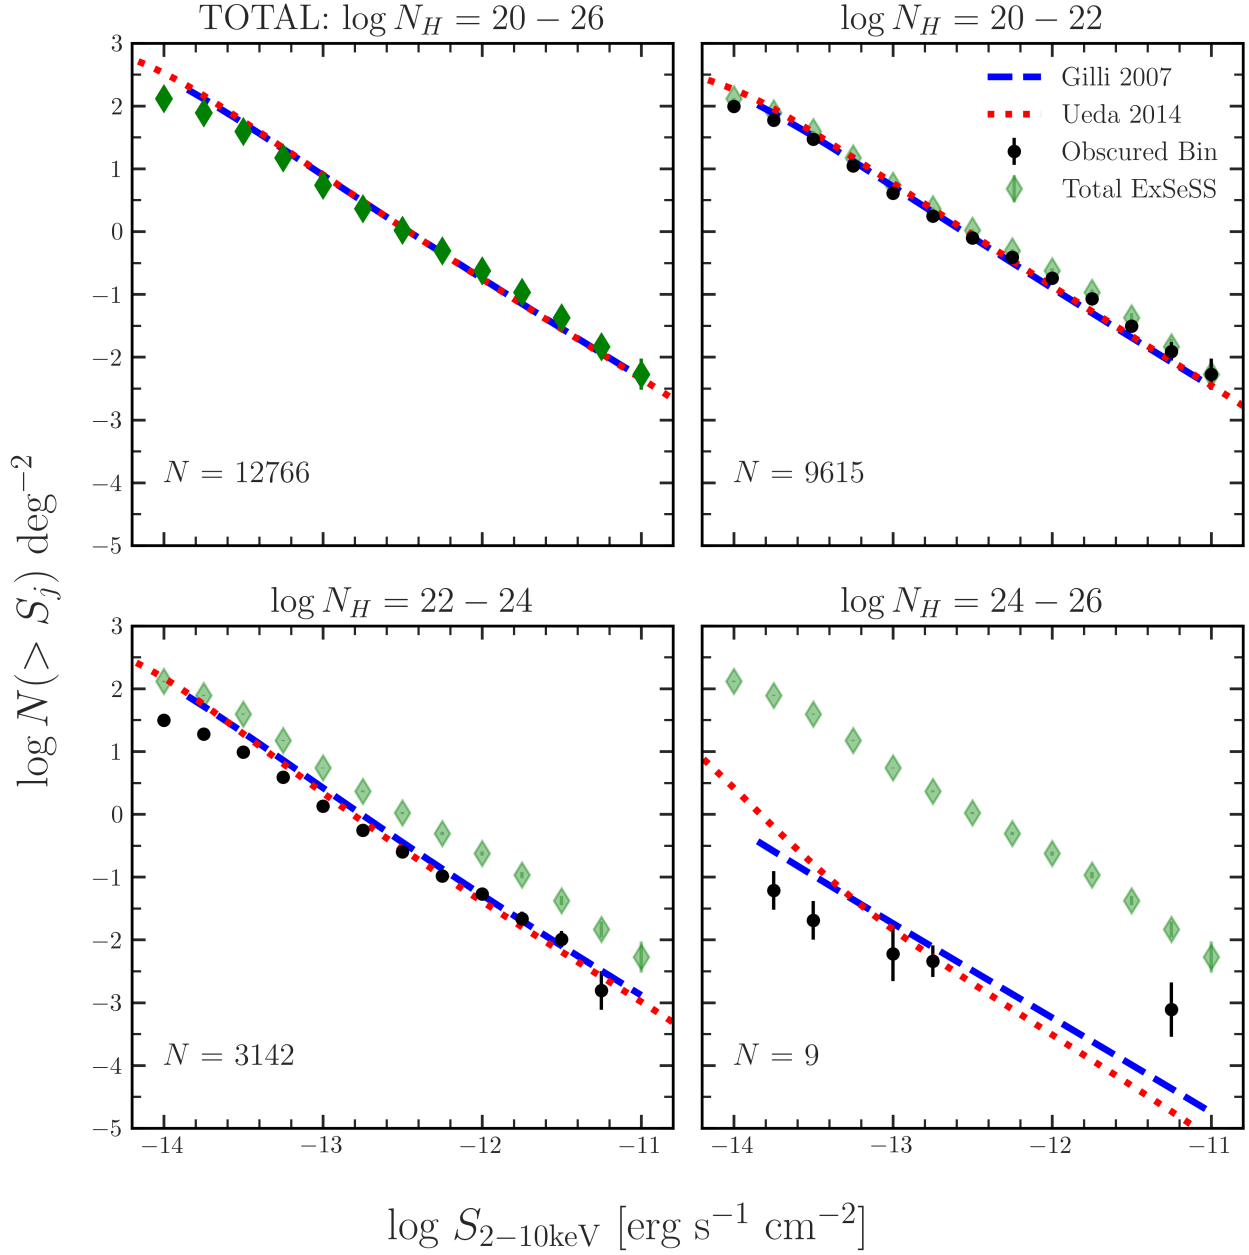


Figure 12. Integrated number counts of *ExSeSS* as a function of 2–10 keV flux ($S_2 - 10$), separated into different groups of obscuration. Top left: $\log N_H = 20 - 26$. Top right: $\log N_H = 20 - 22$. Bottom left: $\log N_H = 22 - 24$. Bottom right: $\log N_H = 24 - 26$. The blue and red lines indicate the expected number counts from two AGN synthesis models, [Gilli et al. \(2007\)](#) and [Ueda et al. \(2014\)](#). The black dots indicate the contribution from the respective obscuration group, while the green diamonds are the integrated number counts from our full sample ($\log N_H = 20 - 26$). Our results agree relatively well across all fluxes in the $\log N_H = 20 - 26$ and $\log N_H = 20 - 22$ groups, only identifying more counts at brighter fluxes and less counts at fainter fluxes. Our faintest counts deviate further from the two models within the $\log N_H = 22 - 24$ group, but still agree well at moderate fluxes. We see the most disagreement in our Compton-thick sources ($\log N_H = 24 - 26$), but note that there is strong agreement at moderate fluxes.

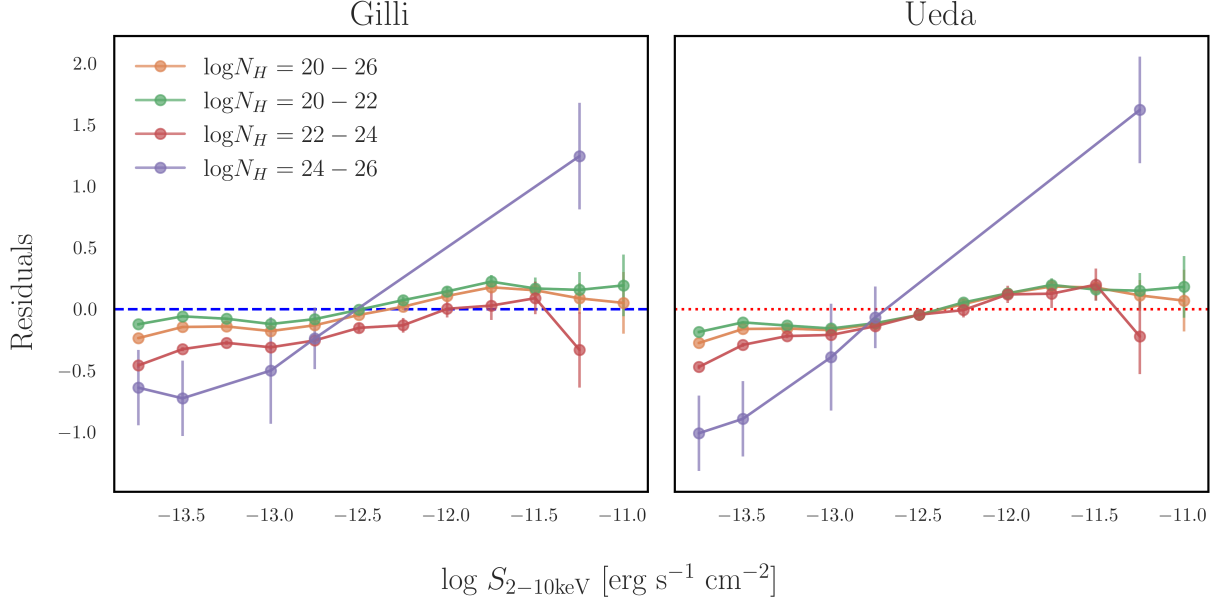


Figure 13. Residual plot of the integrated number counts across 2–10 keV flux with respect to AGN population synthesis models Left: Comparison to Gilli et al. (2007). Right: Comparison to Ueda et al. (2014) We note that there is a trend of negative residuals at fainter fluxes and positive residuals at brighter fluxes. The largest deviations are within our Compton-thick sources. Our number counts at fainter fluxes are likely a consequence of having shallower observations than *XMM-Newton*, which was used in the construction of these synthesis models. The same argument holds for our brighter fluxes, where we observe over a much greater area than *Chandra* and *XMM-Newton* at $\log S_{2-10\text{keV}} > 10^{-12}$ erg s^{−1} cm^{−2}.

AGN population which has so far been unaccounted for by current synthesis models. This is due to the wide area coverage of *ExSeSS*, which surpasses *Chandra* and *XMM-Newton* surveys at $\log S_{2-10\text{keV}} > 10^{-12}$ erg s^{−1} cm^{−2}.

The total and 20–22 groups strongly agreed with both models across all fluxes, with only slight deviations at fainter and brighter fluxes. At higher levels of obscuration, the differences become more pronounced. Within the 22–24 sources, our observations are within 1 σ of both models from $S_{2-10\text{keV}} = 10^{-12} - 10^{-11.3}$ erg s^{−1} cm^{−2}. There are only two sources within the highest flux bin, which explains the large drop in number counts around $S_{2-10\text{keV}} = 10^{-11.3}$. This drop is due to our small sample of bright objects and is likely not indicative of fewer source counts.

The largest deviation in number counts from the G07 and U14 tracks is within our 24–26 bin. Our most secure sources are located around $S_{2-10\text{keV}} = 10^{-13} - 10^{-12.75}$ and are within 1 σ of the G07 and U14 models. At fainter fluxes, we are roughly 2 σ below model expectations. This is likely due to the previously mentioned contrast between this *Swift* XRT survey and deep *XMM-Newton* / *Chandra* surveys, as well as the general difficulty in observing Compton-thick AGN.

The brightest Compton-thick candidate, *2SXPS* ID 182861, is at a redshift of $z = 0.052$. It is roughly 3 σ above the G07 and U14 predictions at $S_{2-10\text{keV}} = 10^{-11.25}$. We mentioned

previously that *ExSeSS* may suggest a larger population of bright AGN than accounted for in existing synthesis models, however, our less obscured observations agree within 1σ of the G07 and U14 models at bright fluxes. This spike in disagreement may be due to contamination by nearby bright sources, as Lansbury et al. (2017) found larger than expected fluxes within their serendipitously detected Compton-thick sources at low redshift. They described this as a “weak” relationship, as there was no evidence that the host galaxies were interacting. The excess may also stem from galaxy clusters, as there is an observed relationship between low redshift bright sources and increased clustering (Cappelluti et al. 2010; M. Koss et al. 2010). Following these arguments, we remind the reader that the *ExSeSS* catalogue is composed of serendipitous detections from *Swift* observations, and therefore might be systemically biased towards finding high number counts from low redshift and bright sources.

In addition to possible contamination, this candidate may also suffer from poor spectral fitting due to incorrect observational quantities. We discovered that the photometric redshift contained within the *ExSeSS* catalogue is $z = 0.052$, while the spectroscopic redshift determined by *BASS* DR2 is $z = 0.012$ (M. J. Koss et al. 2022). Repeating the spectral fitting for this source with the *BASS* redshift returns a calculated obscuration of $\log N_H = 20$. The intense degeneracies between Compton-thick and unobscured sources at low redshift mean that we are unable to accurately discern the obscuration of this source without follow-up from *NuSTAR*. This is one of the key limitations of our work; namely that our predictions are limited to the data at our disposal. The redshift values that we have assumed to be correct may be inaccurate, which has large implications on every aspect of this report. It is beyond the scope of this project to resolve these redshift discrepancies by cross-matching the *ExSeSS* catalogue with other surveys, but we stress that future works would have to be cognisant of this issue.

Despite the uncertainty on ID 182861, we maintain it as a Compton-thick candidate as there is a large degree of uncertainty on any of our most obscured sources, and we are unable to discern the legitimacy of any of them until observing them with a higher energy telescope. The excess in number counts could be explained both by contamination via serendipitous selection or by column density miscalculation. If it is contamination, our number counts should be lower, and thus more in line with both models. In the latter case, the number count will be a negligible change to the $\log N_H = 20 - 22$ bin. Both are suitable explanations for the discrepancy we observe.

We note that we have only been assuming Poisson uncertainties for the number counts on our sources, and have not accounted for any systemic errors such as incorrect physical quantities, our *XSPEC* model, misclassifications from our KNN algorithm, and a lack of sources fainter than $S_{2-10\text{keV}} < 10^{-13.5}$ due to the limitations of *Swift* XRT observations. We have been underestimating the uncertainty of our results, and might have significantly different results if a large population of our sources have inaccurate redshifts.

To conclude this section, our observations agree quite strongly with the predictions of the G07 and U14 models at moderate fluxes across all levels of obscuration. We disagree with the models at fainter fluxes, although our deviations may be excused by the limitations of *Swift* XRT and hence do not suggest a faint-end issue with the G07 and U14 models. At brighter

fluxes, we observe higher number counts than both models which may indicate a population of AGN that has been unaccounted for within the G07 and U14 models. This is motivated by the fact that *Swift* XRT surveys have not been used in the construction of either model, and our observations are therefore probing a unique parameter space. Alternatively, this could be due to contamination from the serendipitous nature of the *ExSeSS* catalogue. We note that the number counts at bright fluxes are exaggerated for our Compton-thick sources, and should be met with skepticism.

8 Relationship Between Obscured AGN and Galaxy Mergers

8.1 Do Obscured AGN Indicate a Galaxy Merger Event?

There has been a lot of interest in the last decade on whether Compton-thick AGN are linked to galaxy mergers. Several works have found higher fractions of Compton-thick AGN within optically identified galaxy mergers (Calabrò et al. 2018; Kocevski et al. 2015; Lansbury et al. 2017; Ricci et al. 2017b, 2021). An implication of this is that the genesis of AGN—what spurs their growth—may be fundamentally linked with galaxy mergers, and thus Compton-thick AGN could indicate a stage of evolution. The hypothesis is that galaxy mergers cause gas to flow toward the centres of the host galaxies, spurring accretion onto a black hole. Some works even suggest using mergers to detect Compton-thick AGN, as was done by Satyapal et al. (2017).

It is not as clear-cut as this, though. Other works have found no links between Compton-thick AGN and the properties of the host galaxy (Lambrides et al. 2021; Mountrichas et al. 2021), which leads to uncertainty on whether there is a link or not. Alternative ignition mechanisms for AGN include stochastic processes (Cisternas et al. 2011; Mullaney et al. 2012), or secular processes, such as barred galaxies having increased gas inflows towards the centre (Maciejewski et al. 2002).

While all three of these hypotheses could help determine how black holes begin accreting, it is not clear which one plays the dominant role. In the remainder of this report we will not attempt to thoroughly answer this quandary. Instead, we will loosely investigate the optical and infrared properties of our Compton-thick candidates in order to determine if there are any signs of galaxy mergers.

8.2 Optical and Infrared Imaging of our Compton-thick Sources

Following all the work of this report we have identified 9 Compton-thick candidates within the *ExSeSS* catalogue. We gathered optical images from Data Release 8 of the Legacy Survey and infrared images from the 4.6 μm WISE instrument. The postage stamps of our sources are shown in Figure 14.

Visual inspection of these sources indicated that there was no clear evidence of a merger event within any of our candidates. *2SXPS* ID 58583 has an irregular infrared shape, but it is not clear within the optical whether a merger is taking place. ID 4900 is located near several point source objects, as well as one larger spiral galaxy. The infrared point sources do not confirm an ongoing merger, despite their proximity to one another. We discussed ID 182861 at length in Section 7.3, and find no immediate signs of a merger event here. We discovered that ID 182861 is equivalently known as UGC 3478 and has previously been studied by Smirnova et al. (2010), who found no extended features within the host galaxy. This suggests that this source is not interacting, and thus our observations are in agreement.

We find no evidence of a relationship between Compton-thick AGN and galaxy mergers, but note that more a rigorous examination of these objects across wavelengths could indicate a previous or ongoing merger event in two of the nine objects. These are *Swift* IDs: 58583 and 4900. We conclude this section by noting that our sources are a small selection of Compton-thick candidates, and that the lack of galaxy mergers is not telling on whether Compton-thick AGN predominantly form from galaxy mergers, stochastic processes, or secular processes.

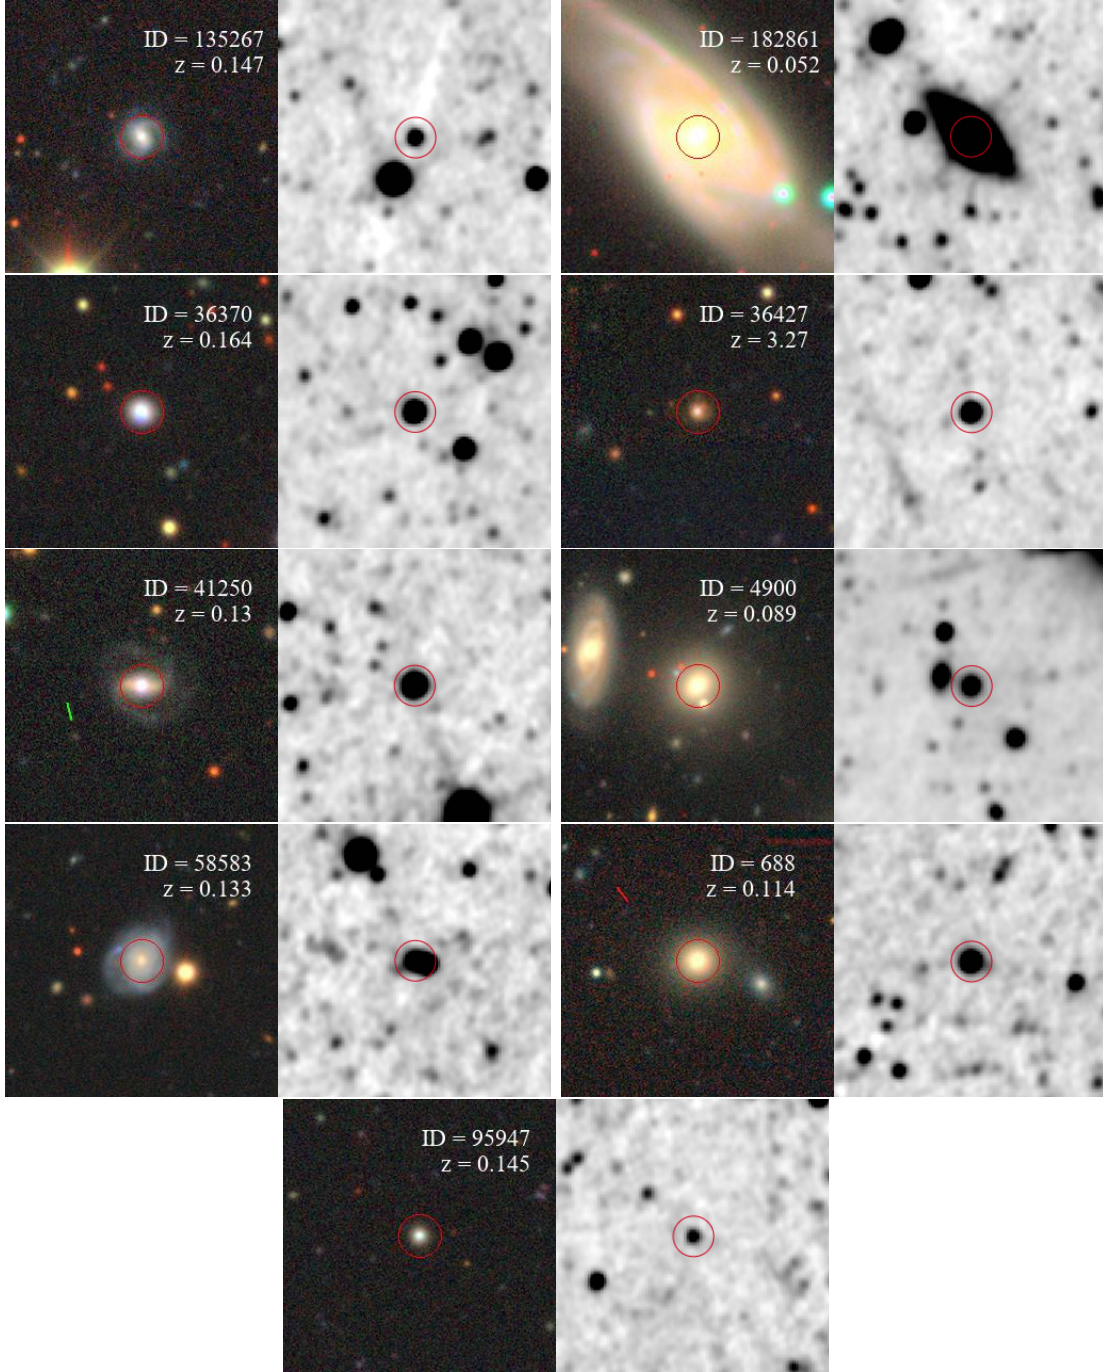


Figure 14. Optical and infrared images of the 12 Compton-thick sources following our KNN algorithm. Optical images are taken from DS8 with a $0.26''$ per pixel resolution. Infrared images are taken from WISE in the $4.6 \mu\text{m}$ waveband with with a $5.2'$ FOV.

9 Conclusion

We determined the effective column densities of a hard-selected sample of *ExSeSS* sources, and then implemented a KNN machine learning algorithm for to estimate obscuration for the low-count sources. We compared our observed obscuration fractions to other surveys, highlighting the distinctiveness of our parameter space. We also calculated the integrated number counts of our sample, identifying discrepancies with pre-existing AGN population synthesis models. We conclude by showing optical and infrared imaging of our Compton-thick candidates, finding no immediate link between Compton-thick levels of absorption and galaxy merger events. The highlights of this paper are:

- The two-dimensional hardness ratios (HR_{HM} and HR_{MS}) of AGN evolve with redshift, and can be used to identify distinct populations of obscured sources. At higher redshifts the degeneracies between Compton-thick and unobscured sources break down, and it becomes easier to identify Compton-thick sources.
- The well-separated nature of these hardness ratios allow for the usage of a KNN machine algorithm in order to calculate column densities of sources with less than 100 photon counts. We found strong agreement across $\log N_H = 20-24$ with *XMM-Newton* observations. The success of this method suggests that it could be used instead of spectral fitting for sources with sparse data.
- We find six Compton-thick candidates following the validation of our spectral fitting, and another three following generalization to our KNN algorithm. The three additional sources were found at low redshift. Two of our initial Compton-thick candidates had similar hardness ratios, while the other four were scattered across the hardness ratio space. The similarity in hardness ratios for these two sources hints at a distinct parameter space for Compton-thick sources, although it just as well may be coincidental given the lack of data.
- We find relatively strong agreement with current synthesis models at moderate fluxes, but identify an excess of integrated number counts at higher fluxes. This could be due to an AGN population that has been unseen by previous *Chandra* and *XMM-Newton* surveys, or due to an unaccounted for bias given the nature of serendipitous detection. Future work could compare the number counts from pointed and triggered *Swift* BAT observations in order to determine whether the excess of number counts is uniform across *Swift* BAT or if it is solely within the triggered observations.
- We find no definite confirmation of galaxy mergers within our Compton-thick candidates, although subsequent analysis could change this for two sources.
- The *ExSeSS* catalogue may have inaccurate redshift values. Future work could involve determining whether *ExSeSS* does indeed have inaccurate redshifts, and if so, could resolve this discrepancy by cross-matching with other surveys or SED fitting *ExSeSS* sources.

References

- Abazajian, K. N. et al. (2009). “The seventh data release of the Sloan Digital Sky Survey (SDSS)”. In: *The Astrophysical Journal Supplement Series* 182.2, p. 543. DOI: [10.1088/0067-0049/182/2/543](https://doi.org/10.1088/0067-0049/182/2/543).
- Aird, J. et al. (2015). “The evolution of the X-ray luminosity functions of unabsorbed and absorbed AGNs out to $z \sim 5$ ”. en. In: *Monthly Notices of the Royal Astronomical Society* 451.2, pp. 1892–1927. ISSN: 0035-8711. DOI: [10.1093/mnras/stv1062](https://doi.org/10.1093/mnras/stv1062).
- Akylas, A. et al. (2012). “Constraining the fraction of Compton-thick AGN in the Universe by modelling the diffuse X-ray background spectrum”. en. In: *Astronomy & Astrophysics* 546, A98. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201219387](https://doi.org/10.1051/0004-6361/201219387).
- (2016). “Compton-thick AGN in the 70-month *Swift* -BAT all-sky hard X-ray survey”. en. In: *Astronomy & Astrophysics* 594, A73. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201628711](https://doi.org/10.1051/0004-6361/201628711).
- Alexander, D. M. et al. (2008). “Reliable identification of Compton-thick quasars at $z \approx 2$ ”. en. In: *The Astrophysical Journal* 687.2, p. 835. ISSN: 0004-637X. DOI: [10.1086/591928](https://doi.org/10.1086/591928).
- Ananna, T. T. et al. (2019). “The accretion history of AGN I”. In: *The Astrophysical Journal* 871.2, p. 240. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/aafb77](https://doi.org/10.3847/1538-4357/aafb77).
- Antonucci, R. (1993). “Unified models for Active Galactic Nuclei and Quasars”. en. In: *Annual Review of Astronomy and Astrophysics* 31, pp. 473–521. ISSN: 0066-4146, 1545-4282. DOI: [10.1146/annurev.aa.31.090193.002353](https://doi.org/10.1146/annurev.aa.31.090193.002353).
- Ballantyne, D. R. et al. (2006). “Connecting Galaxy Evolution, Star Formation, and the Cosmic X-Ray Background”. In: *The Astrophysical Journal* 639.2, pp. 740–752. ISSN: 1538-4357. DOI: [10.1086/499558](https://doi.org/10.1086/499558).
- Baloković, M. et al. (2018). “New spectral model for constraining torus covering factors from broadband X-ray spectra of active galactic nuclei”. en. In: *The Astrophysical Journal* 854.1, p. 42. ISSN: 1538-4357. DOI: [10.3847/1538-4357/aaa7eb](https://doi.org/10.3847/1538-4357/aaa7eb).
- Barger, A. J. et al. (2005). “The cosmic evolution of hard X-ray selected active galactic nuclei”. In: *The Astronomical Journal* 129.2, pp. 578–609. ISSN: 0004-6256, 1538-3881. DOI: [10.1086/426915](https://doi.org/10.1086/426915).
- Bonamente, M. (2020). “Distribution of the C statistic with applications to the sample mean of Poisson data”. In: *Journal of Applied Statistics* 47.11, pp. 2044–2065. DOI: [10.1080/02664763.2019.1704703](https://doi.org/10.1080/02664763.2019.1704703).
- Buchner, J. et al. (2014). “X-ray spectral modelling of the AGN obscuring region in the CDFS”. In: *Astronomy & Astrophysics* 564, A125. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201322971](https://doi.org/10.1051/0004-6361/201322971).
- (2015). “Obscuration-dependent evolution of active galactic nuclei”. en. In: *The Astrophysical Journal* 802.2, p. 89. ISSN: 1538-4357. DOI: [10.1088/0004-637X/802/2/89](https://doi.org/10.1088/0004-637X/802/2/89).
- (2019). “X-ray spectral and eclipsing model of the clumpy obscurer in active galactic nuclei”. In: *Astronomy & Astrophysics* 629, A16. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201834771](https://doi.org/10.1051/0004-6361/201834771).
- Buchner, J. and P. Boorman (2024). “Statistical Aspects of X-ray Spectral Analysis”. In: *Handbook of X-ray and Gamma-ray Astrophysics*. Ed. by Cosimo Bambi and Andrea Santangelo. Singapore: Springer Nature Singapore, pp. 5403–5451. ISBN: 978-981-19-6960-7. DOI: [10.1007/978-981-19-6960-7_175](https://doi.org/10.1007/978-981-19-6960-7_175).

- Burlon, D. et al. (2011). “Three years Swift-BAT Survey of AGN”. In: *The Astrophysical Journal* 728.1, p. 58. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/728/1/58](https://doi.org/10.1088/0004-637X/728/1/58).
- Calabrò, A. et al. (2018). “Near-infrared emission lines in starburst galaxies at $0.5 < z < 0.9$ ”. en. In: *The Astrophysical Journal Letters* 862.2, p. L22. ISSN: 2041-8205, 2041-8213. DOI: [10.3847/2041-8213/aad33e](https://doi.org/10.3847/2041-8213/aad33e).
- Cappelluti, N. et al. (2010). “Active galactic nuclei clustering in the local Universe”. en. In: *The Astrophysical Journal* 716.2, pp. L209–L213. ISSN: 2041-8205, 2041-8213. DOI: [10.1088/2041-8205/716/2/L209](https://doi.org/10.1088/2041-8205/716/2/L209).
- Carroll, C. M. et al. (2023). “A high fraction of heavily X-ray obscured active galactic nuclei”. In: *The Astrophysical Journal* 950.2, p. 127. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/acc402](https://doi.org/10.3847/1538-4357/acc402).
- Cash, W. (1978). “Parameter estimation in astronomy through application of the likelihood ratio”. In: *The Astrophysical Journal* 228, pp. 939–947.
- Chen, C. J. et al. (2017). “The X-ray and mid-infrared luminosities in luminous type 1 quasars”. en. In: *The Astrophysical Journal* 837.2, p. 145. ISSN: 0004-637X. DOI: [10.3847/1538-4357/837/2/145](https://doi.org/10.3847/1538-4357/837/2/145).
- Christen, P. et al. (2023). “A review of the F-measure”. In: *ACM Computing Surveys* 56.3, 73:1–73:24. ISSN: 0360-0300. DOI: [10.1145/3606367](https://doi.org/10.1145/3606367).
- Cisternas, M. et al. (2011). “The bulk of the black hole growth since $z \sim 1$ occurs in a secular Universe”. In: *The Astrophysical Journal* 726.2, p. 57. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/726/2/57](https://doi.org/10.1088/0004-637X/726/2/57).
- Cooper, N. et al. (2023). “Fermi LAT AGN classification using supervised machine learning”. In: *Monthly Notices of the Royal Astronomical Society* 525.2, pp. 1731–1745. ISSN: 0035-8711. DOI: [10.1093/mnras/stad2193](https://doi.org/10.1093/mnras/stad2193).
- Crummy, J. et al. (2005). “An explanation for the soft X-ray excess in active galactic nuclei”. en. In: *Monthly Notices of the Royal Astronomical Society* 365.4, pp. 1067–1081. ISSN: 00358711, 13652966. DOI: [10.1111/j.1365-2966.2005.09844.x](https://doi.org/10.1111/j.1365-2966.2005.09844.x).
- Delaney, J. N. et al. (2023). “ExSeSS – I”. en. In: *Monthly Notices of the Royal Astronomical Society* 521.2, pp. 1620–1632. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stac3703](https://doi.org/10.1093/mnras/stac3703).
- Evans, I. N. et al. (2010). “The *Chandra* Source Catalog”. en. In: *The Astrophysical Journal Supplement Series* 189.1, pp. 37–82. ISSN: 0067-0049, 1538-4365. DOI: [10.1088/0067-0049/189/1/37](https://doi.org/10.1088/0067-0049/189/1/37).
- Evans, P. A. et al. (2020). “2SXPS”. en. In: *The Astrophysical Journal Supplement Series* 247.2, p. 54. ISSN: 0067-0049. DOI: [10.3847/1538-4365/ab7db9](https://doi.org/10.3847/1538-4365/ab7db9).
- Fabian, A. C. (1999). “Active galactic nuclei”. In: *Proceedings of the National Academy of Sciences of the United States of America* 96.9, pp. 4749–4751. ISSN: 0027-8424.
- Fiore, F. et al. (2009). “Chasing highly obscured QSOs in the COSMOS field”. en. In: *The Astrophysical Journal* 693.1, p. 447. ISSN: 0004-637X. DOI: [10.1088/0004-637X/693/1/447](https://doi.org/10.1088/0004-637X/693/1/447).
- Fotopoulou, S. et al. (2016). “The XXL Survey”. en. In: *Astronomy & Astrophysics* 592, A5. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201527402](https://doi.org/10.1051/0004-6361/201527402).
- Gandhi, P. et al. (2014). “NuSTAR unveils a Compton-thick type 2 quasar in Mrk 34”. en. In: *The Astrophysical Journal* 792.2, p. 117. ISSN: 0004-637X. DOI: [10.1088/0004-637X/792/2/117](https://doi.org/10.1088/0004-637X/792/2/117).

- Giacconi, R. et al. (1962). “Evidence for X-rays from sources outside the solar system”. en. In: *Physical Review Letters* 9.11, pp. 439–443. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.9.439](https://doi.org/10.1103/PhysRevLett.9.439).
- (2001). “First results from the X-ray and optical survey of the Chandra Deep Field South”. In: *The Astrophysical Journal* 551.2, pp. 624–634. ISSN: 0004-637X, 1538-4357. DOI: [10.1086/320222](https://doi.org/10.1086/320222).
- Giacconi, R. (2009). “History of X-ray telescopes and astronomy”. en. In: *Experimental Astronomy* 25.1, pp. 143–156. ISSN: 1572-9508. DOI: [10.1007/s10686-009-9139-8](https://doi.org/10.1007/s10686-009-9139-8).
- Gilli, R. et al. (2007). “The synthesis of the cosmic X-ray background in the Chandra and XMM-Newton era”. en. In: *Astronomy & Astrophysics* 463.1, pp. 79–96. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:20066334](https://doi.org/10.1051/0004-6361:20066334).
- Harrison, F. A. et al. (2013). “The *NuSTAR* high-energy X-ray mission”. en. In: *The Astrophysical Journal* 770.2, p. 103. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/770/2/103](https://doi.org/10.1088/0004-637X/770/2/103).
- Heckman, T. and P. Best (2014). “The co-evolution of galaxies and supermassive black holes”. In: *Annual Review of Astronomy and Astrophysics* 52.1, pp. 589–660. ISSN: 0066-4146, 1545-4282. DOI: [10.1146/annurev-astro-081913-035722](https://doi.org/10.1146/annurev-astro-081913-035722).
- Hickox, R. C. and D. M. Alexander (2018). “Obscured active galactic nuclei”. In: *Annual Review of Astronomy and Astrophysics* 56.1, pp. 625–671. ISSN: 0066-4146, 1545-4282. DOI: [10.1146/annurev-astro-081817-051803](https://doi.org/10.1146/annurev-astro-081817-051803).
- Humphrey, P. J. et al. (2009). “Chi-square and Poissonian data”. In: *The Astrophysical Journal* 693.1, pp. 822–829. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/693/1/822](https://doi.org/10.1088/0004-637X/693/1/822).
- Ishibashi, W. and T. J.-L. Courvoisier (2010). “X-ray power law spectra in active galactic nuclei”. en. In: *Astronomy and Astrophysics* 512, A58. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/200913587](https://doi.org/10.1051/0004-6361/200913587).
- Jansen, F. et al. (2001). “XMM-Newton observatory. I.” In: *Astronomy & Astrophysics* 365, pp. L1–L6. DOI: [10.1051/0004-6361:20000036](https://doi.org/10.1051/0004-6361:20000036).
- Kaastra, J. S. (2017). “On the use of C-stat in testing models for X-ray spectra”. In: *Astronomy & Astrophysics* 605, A51. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201629319](https://doi.org/10.1051/0004-6361/201629319).
- Kocevski, D. D. et al. (2015). “Are Compton-thick AGN the missing link between mergers and black hole growth?”. In: *The Astrophysical Journal* 814.2, p. 104. ISSN: 1538-4357. DOI: [10.1088/0004-637X/814/2/104](https://doi.org/10.1088/0004-637X/814/2/104).
- Koss, M. et al. (2010). “Merging and clustering of the Swift BAT AGN sample”. en. In: *The Astrophysical Journal* 716.2, pp. L125–L130. ISSN: 2041-8205, 2041-8213. DOI: [10.1088/2041-8205/716/2/L125](https://doi.org/10.1088/2041-8205/716/2/L125).
- Koss, M. J. et al. (2022). “BASS. XXII. The BASS DR2 AGN catalog and data”. en. In: *The Astrophysical Journal Supplement Series* 261.1, p. 2. ISSN: 0067-0049, 1538-4365. DOI: [10.3847/1538-4365/ac6c05](https://doi.org/10.3847/1538-4365/ac6c05).
- Lambrides, E. L. et al. (2021). “Lower-luminosity obscured AGN host galaxies are not predominantly in major-merging systems at cosmic noon”. en. In: *The Astrophysical Journal* 919.2, p. 129. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/ac12c8](https://doi.org/10.3847/1538-4357/ac12c8).
- Lansbury, G. B. et al. (2017). “The *NuSTAR* serendipitous survey”. en. In: *The Astrophysical Journal* 846.1, p. 20. ISSN: 0004-637X. DOI: [10.3847/1538-4357/aa8176](https://doi.org/10.3847/1538-4357/aa8176).

- Li, L. et al. (2008). “K-Nearest Neighbors for automated classification of celestial objects”. en. In: *Science in China Series G: Physics, Mechanics and Astronomy* 51.7, pp. 916–922. ISSN: 1862-2844. DOI: [10.1007/s11433-008-0088-4](https://doi.org/10.1007/s11433-008-0088-4).
- Lutz, D. et al. (2004). “The relation between AGN hard X-ray emission and mid-infrared continuum from ISO spectra”. In: *Astronomy & Astrophysics* 418.2, pp. 465–473. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:20035838](https://doi.org/10.1051/0004-6361:20035838).
- Maciejewski, W. et al. (2002). “Gas inflow in barred galaxies — effects of secondary bars”. In: *Monthly Notices of the Royal Astronomical Society* 329.3, pp. 502–512. ISSN: 0035-8711. DOI: [10.1046/j.1365-8711.2002.04957.x](https://doi.org/10.1046/j.1365-8711.2002.04957.x).
- Mateos, S. et al. (2008). “High precision X-ray log N – log S distributions”. en. In: *Astronomy & Astrophysics* 492.1, pp. 51–69. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:200810004](https://doi.org/10.1051/0004-6361:200810004).
- Moretti, A. et al. (2009). “A new measurement of the cosmic X-ray background”. en. In: *Astronomy & Astrophysics* 493.2, pp. 501–509. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:200811197](https://doi.org/10.1051/0004-6361:200811197).
- Mostafa, R. et al. (2023). “Selection of Compton-thick AGN from a hard photometric sample using *XMM-Newton* observations”. en. In: *Monthly Notices of the Royal Astronomical Society* 525.4, pp. 5080–5104. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stad2425](https://doi.org/10.1093/mnras/stad2425).
- Mountrichas, G. et al. (2021). “The role of AGN and obscuration in the position of the host galaxy relative to the main sequence”. en. In: *Astronomy & Astrophysics* 653, A74. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202140630](https://doi.org/10.1051/0004-6361/202140630).
- Mullaney, J. R. et al. (2012). “The hidden AGN main sequence”. en. In: *The Astrophysical Journal* 753.2, p. L30. ISSN: 2041-8205, 2041-8213. DOI: [10.1088/2041-8205/753/2/L30](https://doi.org/10.1088/2041-8205/753/2/L30).
- Nandra, K. et al. (2024). “The eROSITA Final Equatorial Depth Survey (eFEDS)”. In: NASA (2024). *Fermi Gamma-ray space telescope*. URL: <https://fermi.gsfc.nasa.gov/science/etev/agn/> (visited on 01/18/2024).
- Netzer, H. (2015). “Revisiting the unified model of active galactic nuclei”. In: *Annual Review of Astronomy and Astrophysics* 53.1, pp. 365–408. DOI: [10.1146/annurev-astro-082214-122302](https://doi.org/10.1146/annurev-astro-082214-122302).
- Nousek, J. A. and D. R. Shue (1989). “ χ^2 and C statistic minimization for low count per bin data”. In: *The Astrophysical Journal* 342, p. 1207. DOI: [10.1086/167676](https://doi.org/10.1086/167676).
- Padovani, P. et al. (2017). “Active galactic nuclei”. In: *The Astronomy and Astrophysics Review* 25.1, p. 2. ISSN: 0935-4956, 1432-0754. DOI: [10.1007/s00159-017-0102-9](https://doi.org/10.1007/s00159-017-0102-9).
- Park, T. et al. (2006). “Bayesian estimation of hardness ratios”. en. In: *The Astrophysical Journal* 652.1, p. 610. ISSN: 0004-637X. DOI: [10.1086/507406](https://doi.org/10.1086/507406).
- Predehl, P. et al. (2021). “The eROSITA X-ray telescope on SRG”. en. In: *Astronomy & Astrophysics* 647, A1. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202039313](https://doi.org/10.1051/0004-6361/202039313).
- Ricci, C. et al. (2015). “Compton-thick accretion in the local Universe”. In: *The Astrophysical Journal* 815.1, p. L13. ISSN: 2041-8213. DOI: [10.1088/2041-8205/815/1/L13](https://doi.org/10.1088/2041-8205/815/1/L13).
- (2017a). “BAT AGN Spectroscopic Survey. V”. en. In: *The Astrophysical Journal Supplement Series* 233.2, p. 17. ISSN: 1538-4365. DOI: [10.3847/1538-4365/aa96ad](https://doi.org/10.3847/1538-4365/aa96ad).
 - (2017b). “Growing supermassive black holes in the late stages of galaxy mergers are heavily obscured”. In: *Monthly Notices of the Royal Astronomical Society*, stx173. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stx173](https://doi.org/10.1093/mnras/stx173).

- Ricci, C. et al. (2021). “A hard X-ray view of luminous and ultra-luminous infrared galaxies in GOALS”. en. In: *Monthly Notices of the Royal Astronomical Society* 506.4, pp. 5935–5950. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stab2052](https://doi.org/10.1093/mnras/stab2052).
- Risaliti, G. and M. Elvis (2004). “A panchromatic view of AGN”. In: vol. 308, pp. 187–224. DOI: [10.1007/978-1-4020-2471-9_6](https://doi.org/10.1007/978-1-4020-2471-9_6).
- Satyapal, S. et al. (2017). “Buried AGNs in advanced mergers”. en. In: *The Astrophysical Journal* 848.2, p. 126. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/aa88ca](https://doi.org/10.3847/1538-4357/aa88ca).
- Sengupta, D. et al. (2023). “Compton-thick AGN in the *NuSTAR* Era X”. en. In: *Astronomy & Astrophysics* 676, A103. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202245646](https://doi.org/10.1051/0004-6361/202245646).
- Signorini, M. et al. (2023). “X-ray properties and obscured fraction of AGN in the J1030 Chandra field”. en. In: *Astronomy & Astrophysics* 676, A49. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202346364](https://doi.org/10.1051/0004-6361/202346364).
- Silverman, J. D. et al. (2008). “The evolution of AGN host galaxies”. en. In: *The Astrophysical Journal* 675.2, pp. 1025–1040. ISSN: 0004-637X, 1538-4357. DOI: [10.1086/527283](https://doi.org/10.1086/527283).
- Smirnova, A. A. et al. (2010). “Seyfert galaxies that are undergoing merging but appear non-interacting”. In: *Monthly Notices of the Royal Astronomical Society* 408.1, pp. 400–406. ISSN: 0035-8711. DOI: [10.1111/j.1365-2966.2010.17121.x](https://doi.org/10.1111/j.1365-2966.2010.17121.x).
- Stern, D. (2015). “The X-ray to mid-infrared relation of AGNs at high luminosity”. en. In: *The Astrophysical Journal* 807.2, p. 129. ISSN: 0004-637X. DOI: [10.1088/0004-637X/807/2/129](https://doi.org/10.1088/0004-637X/807/2/129).
- Szokoly, G. P. et al. (2004). “The Chandra Deep Field South”. In: *The Astrophysical Journal Supplement Series* 155.2, pp. 271–349. ISSN: 0067-0049, 1538-4365. DOI: [10.1086/424707](https://doi.org/10.1086/424707).
- Ueda, Y. et al. (2014). “Toward the standard population synthesis model of the X-ray background”. en. In: *The Astrophysical Journal* 786.2, p. 104. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/786/2/104](https://doi.org/10.1088/0004-637X/786/2/104).
- Webb, N. A. et al. (2020). “The *XMM-Newton* serendipitous survey”. en. In: *Astronomy & Astrophysics* 641, A136. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201937353](https://doi.org/10.1051/0004-6361/201937353).
- Willingale, R. et al. (2013). “Calibration of X-ray absorption in our galaxy”. en. In: *Monthly Notices of the Royal Astronomical Society* 431.1, pp. 394–404. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stt175](https://doi.org/10.1093/mnras/stt175).
- Wright, E. L. et al. (2010). “The wide-field infrared survey explorer (WISE)”. In: *The Astronomical Journal* 140.6, p. 1868. DOI: [10.1088/0004-6256/140/6/1868](https://doi.org/10.1088/0004-6256/140/6/1868).
- Zappacosta, L. et al. (2018). “The *NuSTAR* extragalactic surveys”. en. In: *The Astrophysical Journal* 854.1, p. 33. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/aaa550](https://doi.org/10.3847/1538-4357/aaa550).

Appendices

Note — Appendices are provided for completeness only and any content included in them will be disregarded for the purposes of assessment.

A Intrinsic Mid-Infrared Luminosity Derivation

Magnitudes are defined as:

$$m = -2.5 \log \frac{f_v(v_{obs})}{F_v(v_{obs})}$$

where m is the apparent magnitude, f_v is the flux observed at a given frequency in units Jansky, and F_v is the zero-point. In an AB magnitude system, the zero-points are the same for all frequencies, so the equation reduces to:

$$m_1 - m_2 = -2.5 \log \frac{f_{v_1}(v_{1,obs})}{F_v(v_{obs})} \frac{F_v(v_{obs})}{f_{v_2}(v_{2,obs})} = -2.5 \log \frac{f_{v_1}(v_{1,obs})}{f_{v_2}(v_{2,obs})}$$

Assuming a power-law spectrum between the two filters such that $f_v \propto v^\alpha$ and re-arranging for the α term, we find that:

$$\alpha = \frac{m_2 - m_1}{2.5 \log \left(\frac{v_1}{v_2} \right)}$$

Now that we have determined a value for α , we need to correct the observed f_v values. The relationship between luminosity and flux accounting for redshift is given by:

$$L_v(v_{obs}[1+z]) = \frac{4\pi D_L^2}{1+z} f_v(v_{obs})$$

where z is the redshift and D_L is the luminosity distance. Then, multiplying the right side by 1 (introducing the luminosity terms at the desired frequency):

$$L_v(v_{obs}[1+z]) = \frac{4\pi D_L^2}{1+z} f_v(v_{obs}) \frac{L_v(v_{corr})}{L_v(v_{corr})}$$

Re-arranging the equation and solving for $L_v(v_{corr})$, we then get:

$$L_v(v_{corr}) = \frac{4\pi D_L^2}{1+z} f_v(v_{obs}) \frac{L_v(v_{corr})}{L_v(v_{obs}[1+z])}$$

Since we assumed that $f_v \propto v^\alpha$, it follows that $L_v \propto v^\alpha$. The equation then becomes:

$$L_v(v_{corr}) = \frac{4\pi D_L^2}{1+z} f_v(v_{obs}) \left(\frac{v_{corr}}{v_{obs}[1+z]} \right)^\alpha$$

Expressing the equation as wavelength instead of frequency and solving for the desired 6 μm using the observed 4.6 μm WISE 2 instrument, it takes the following final form:

$$L_v(\lambda_{6\mu m}) = 4\pi D_L^2 f_\lambda(\lambda_{6\mu m}) \left(\frac{\lambda_{4.6\mu m}}{\lambda_{6\mu m}} \right)^\alpha (1+z)^{-(1+\alpha)}$$

B Additional Spectra

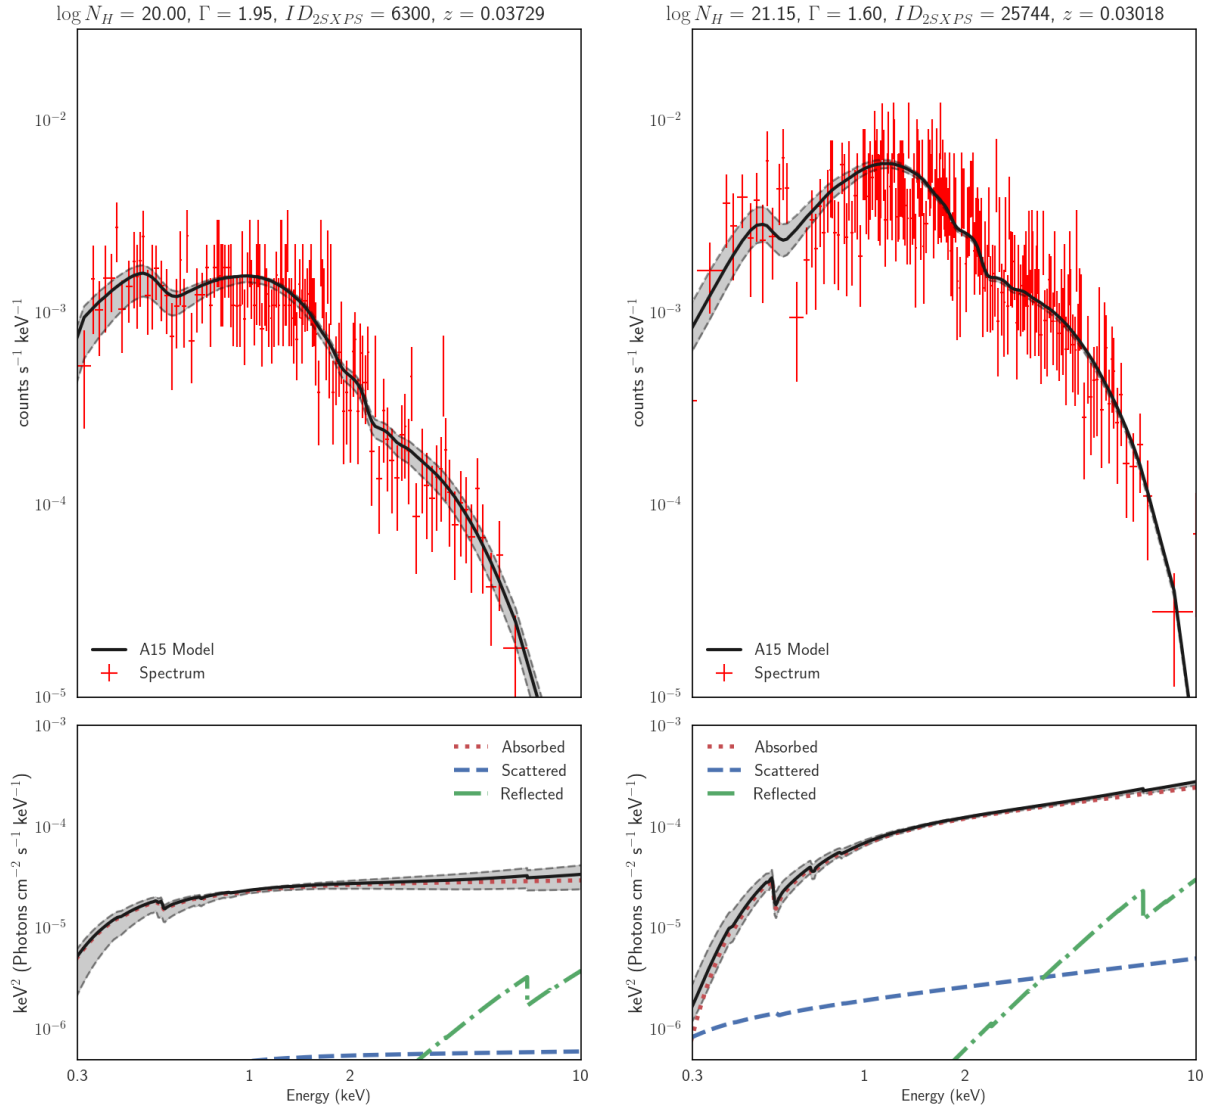


Figure 15. Spectral fitting at low redshifts. Left is an unobscured source, right is a lightly obscured source.

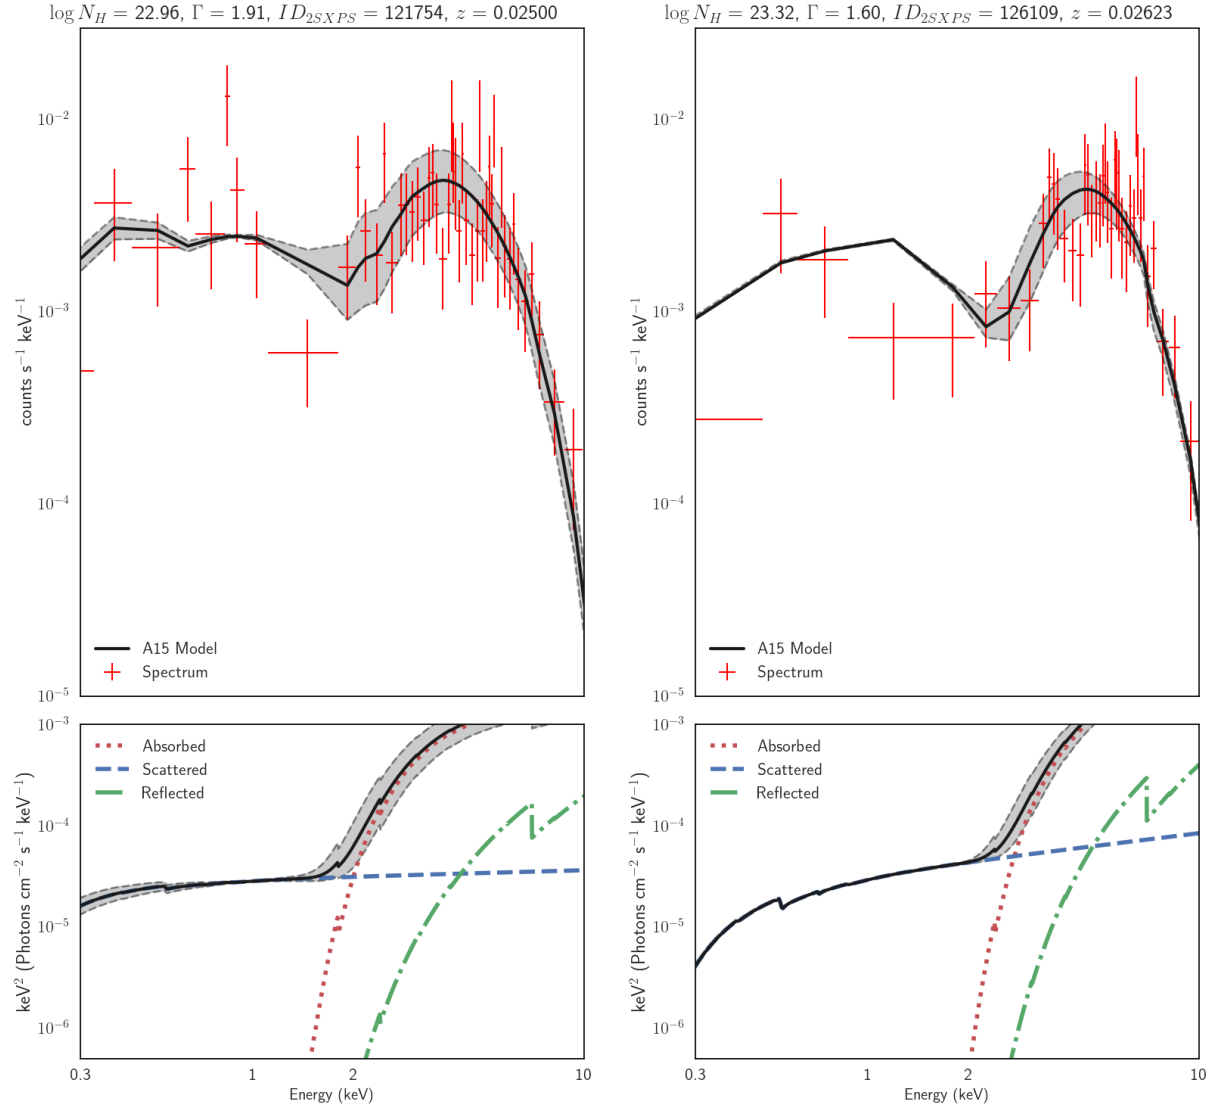


Figure 16. Spectral fitting at low redshifts. Left is a moderately obscured source, right is a heavily obscured source

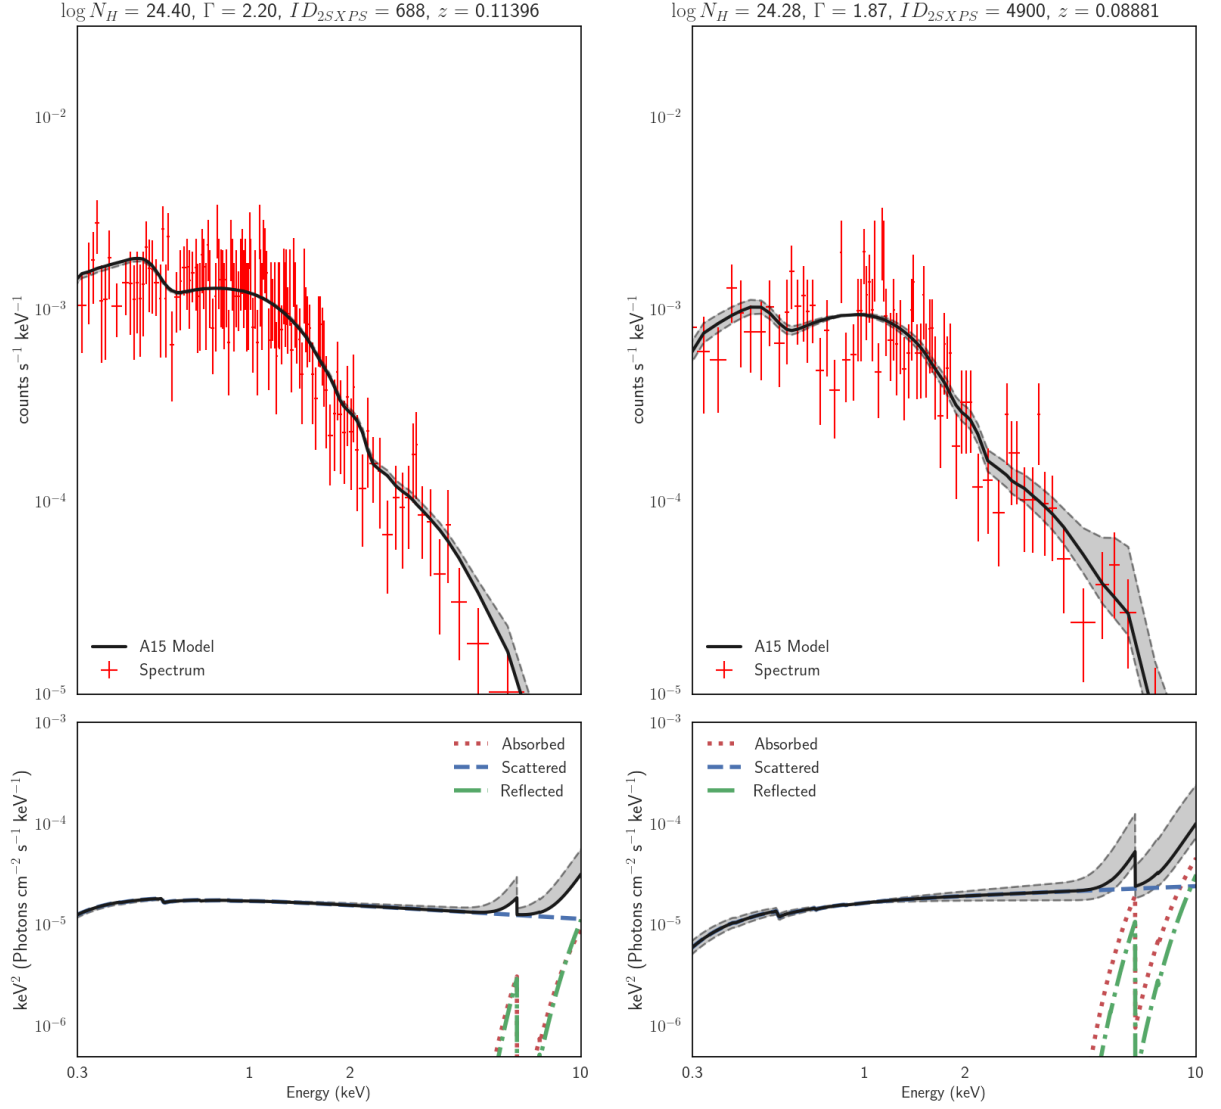


Figure 17. Spectral fitting at low redshifts for Compton-thick sources. Note that the component is entirely controlled by the scattered component—the redshift is too low for the reflected component to contribute, and the absorbed component is fully absorbed by the torus.

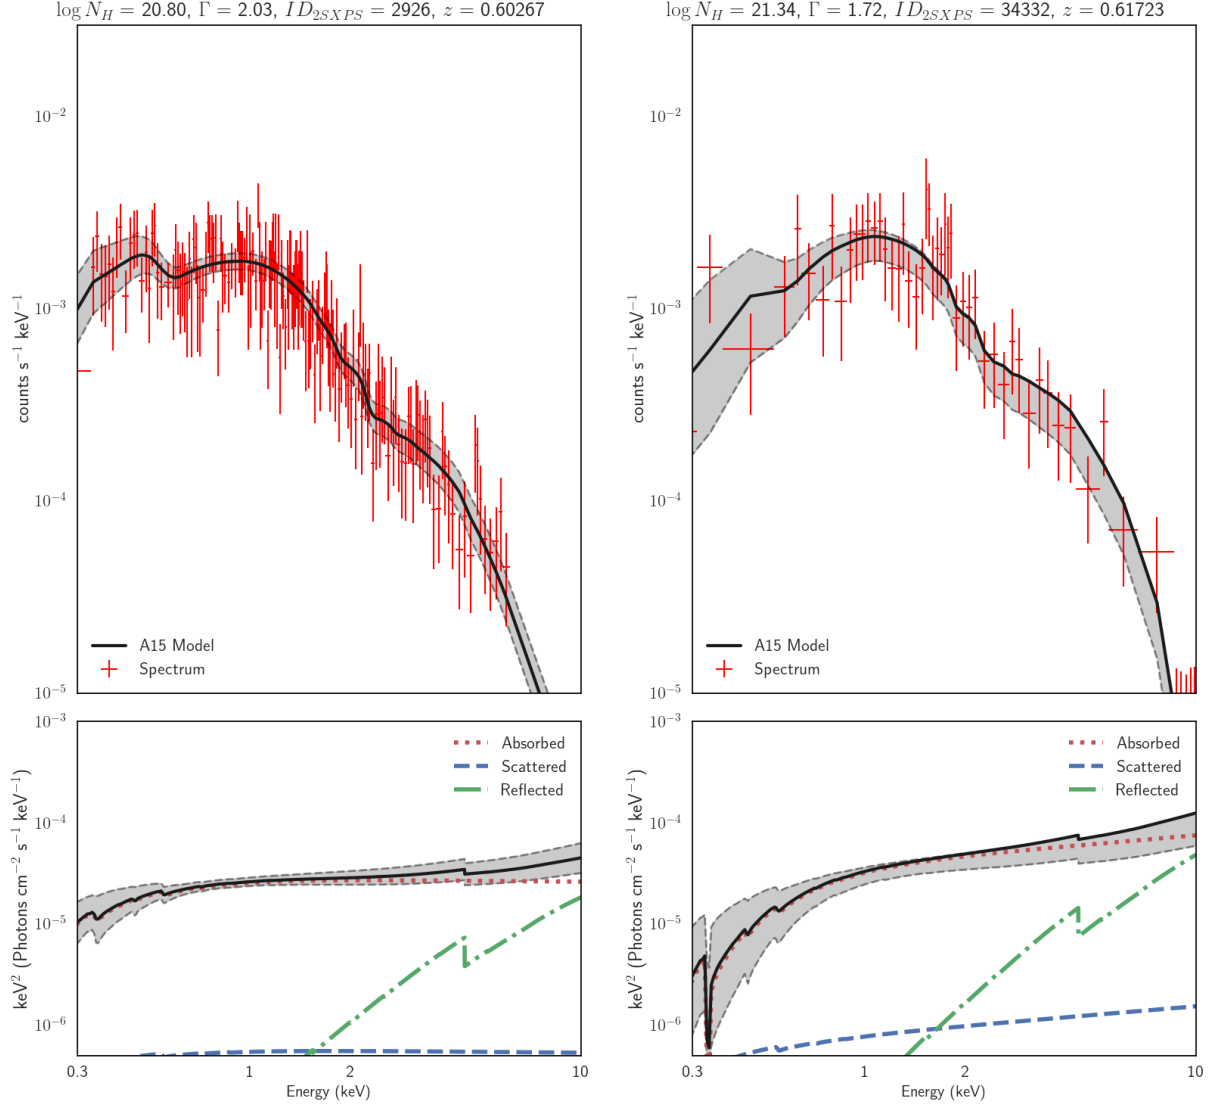


Figure 18. Spectral fitting at moderate redshifts. Left is an unobscured source, right is a lightly obscured source.

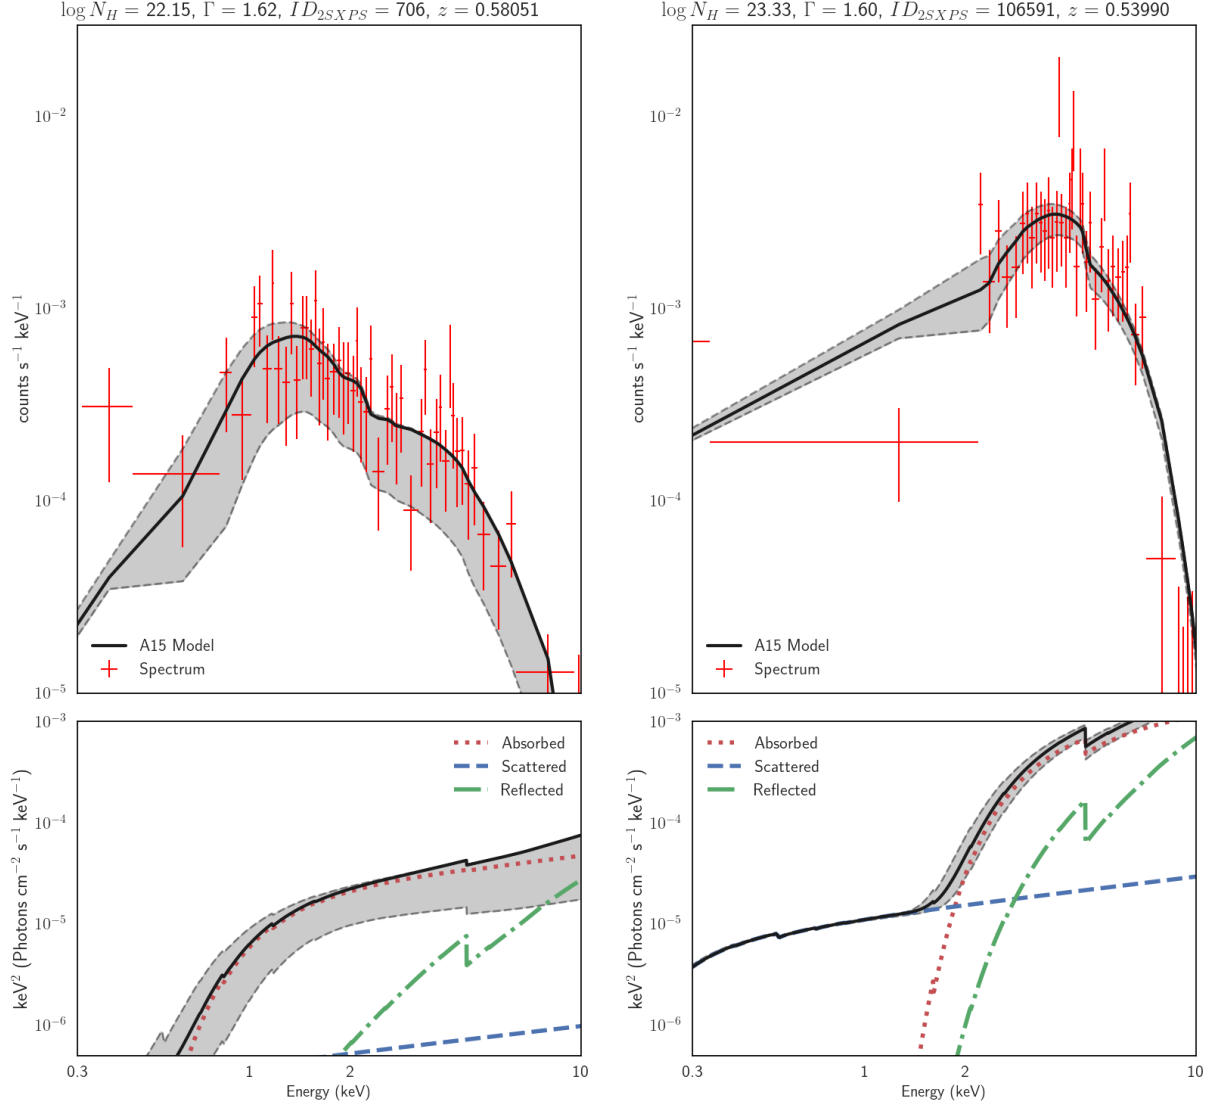


Figure 19. Spectral fitting at moderate redshifts. Left is a moderately obscured source, right is a heavily obscured source.

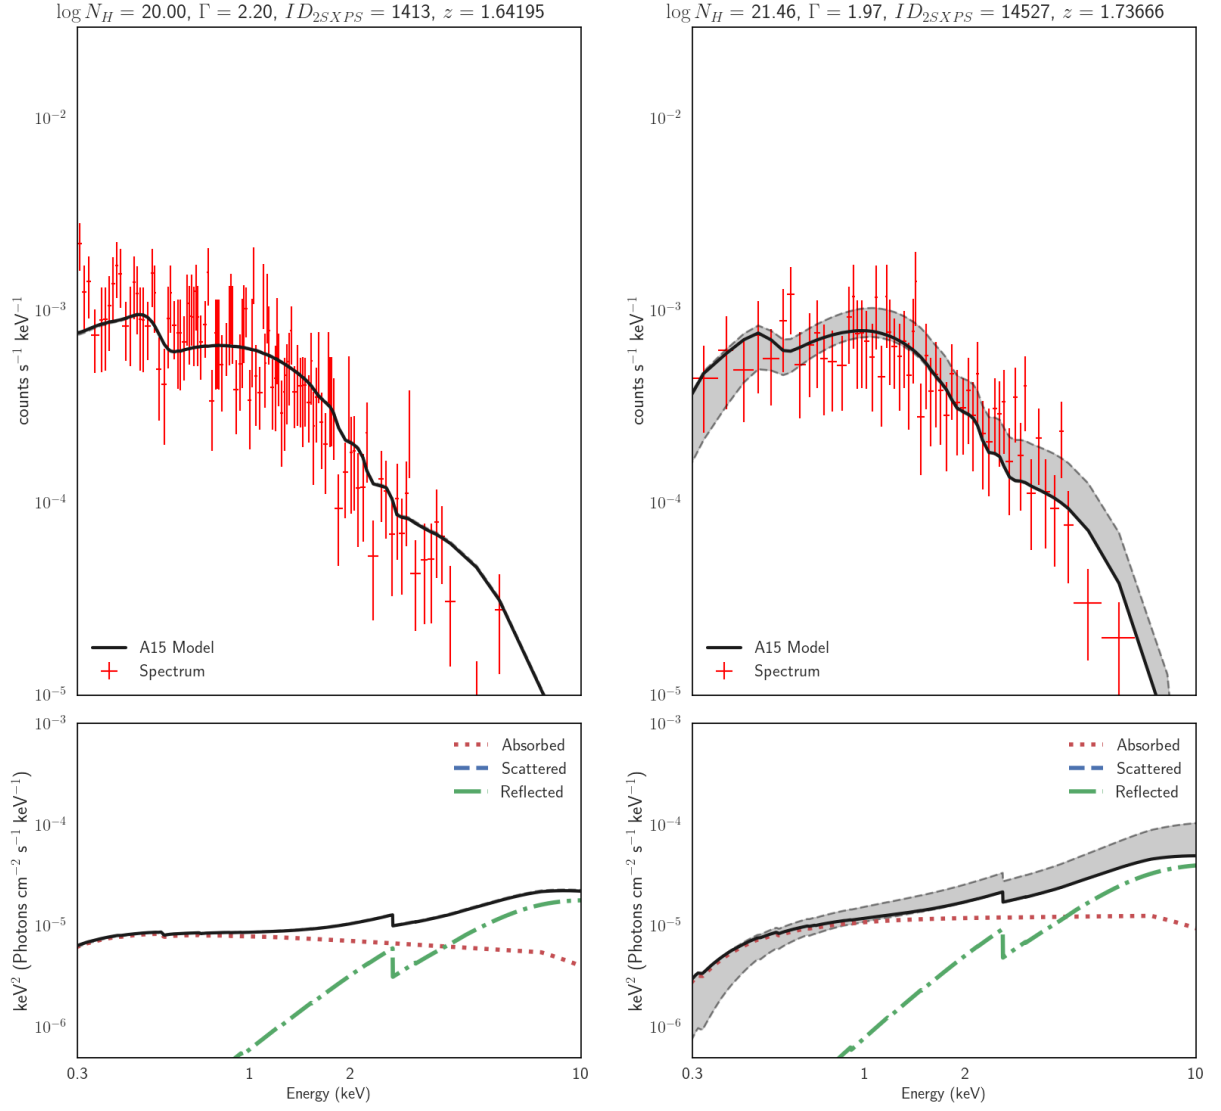


Figure 20. High redshift sources. Unobscured source on the left, lightly obscured source on the right.

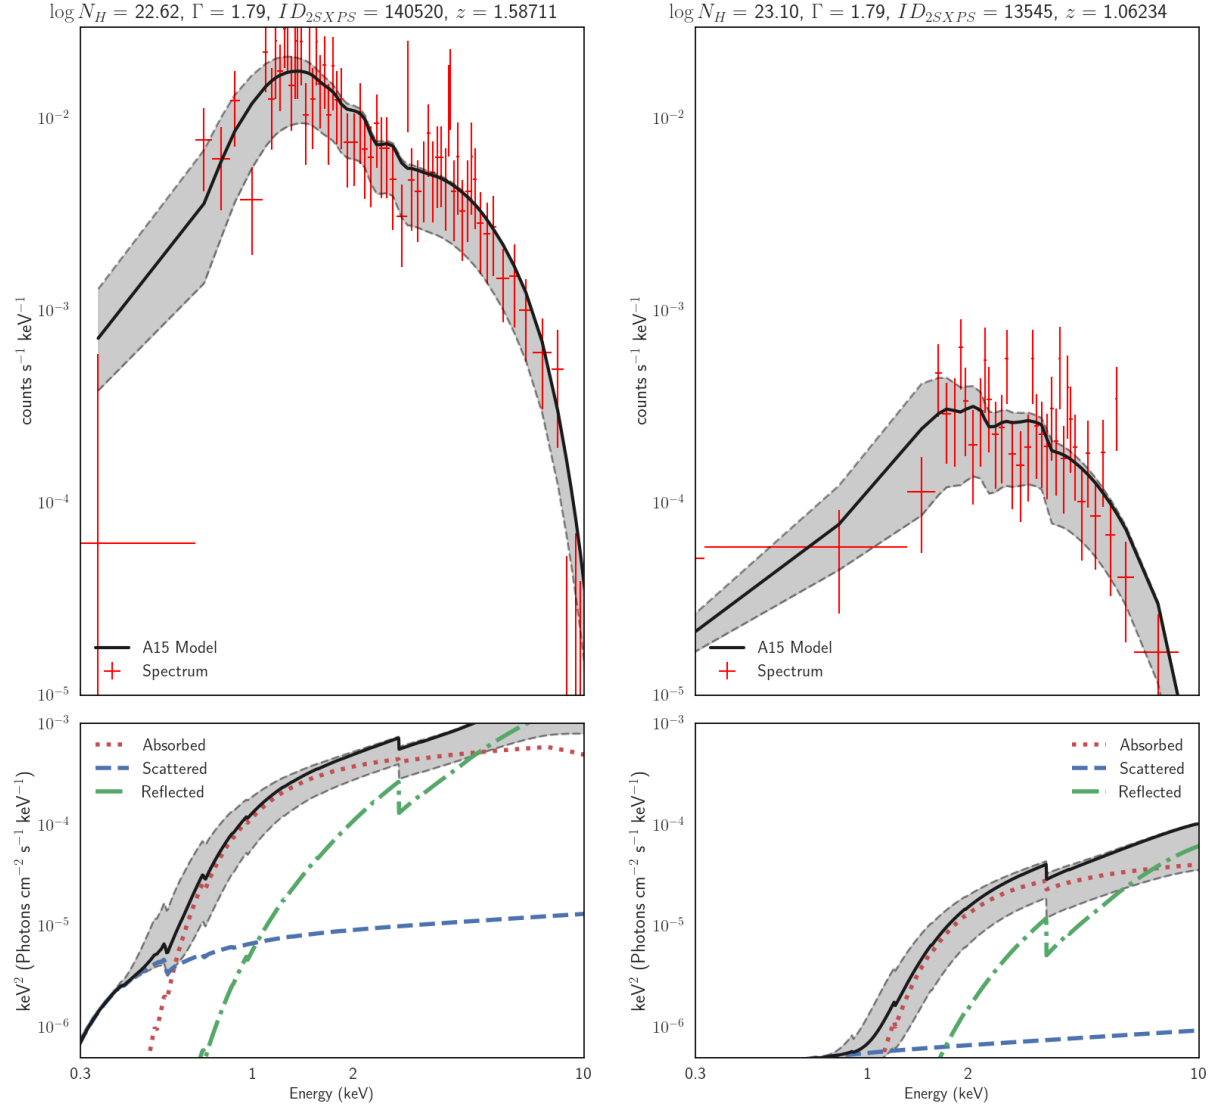


Figure 21. High redshift sources at moderate (left) and heavy (right) levels of obscuration.

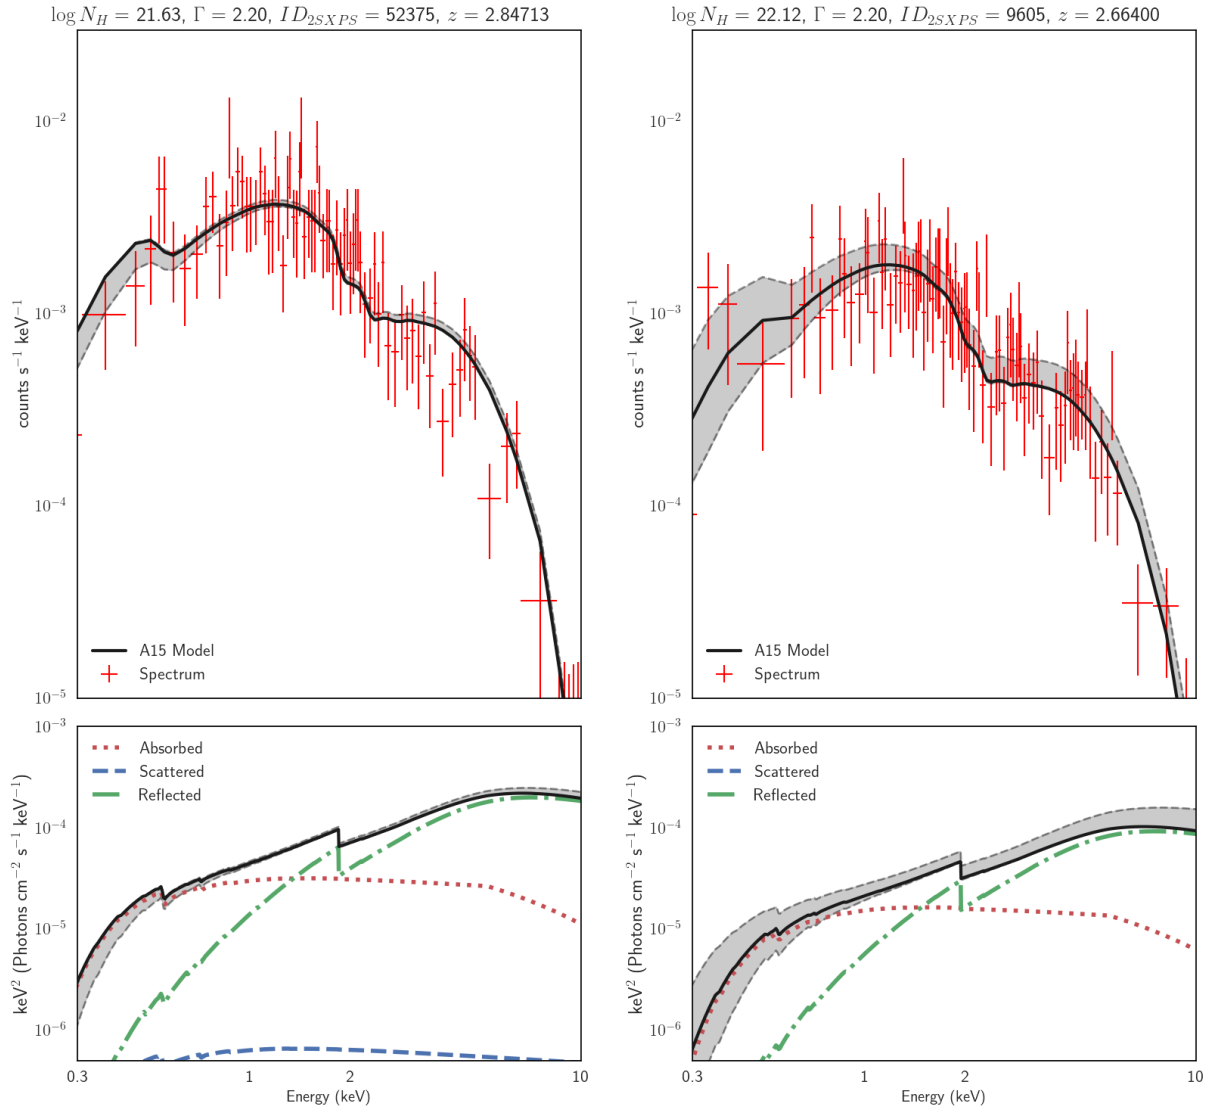


Figure 22. Very high redshift sources at light levels of obscuration (left) and moderate levels of obscuration (right).

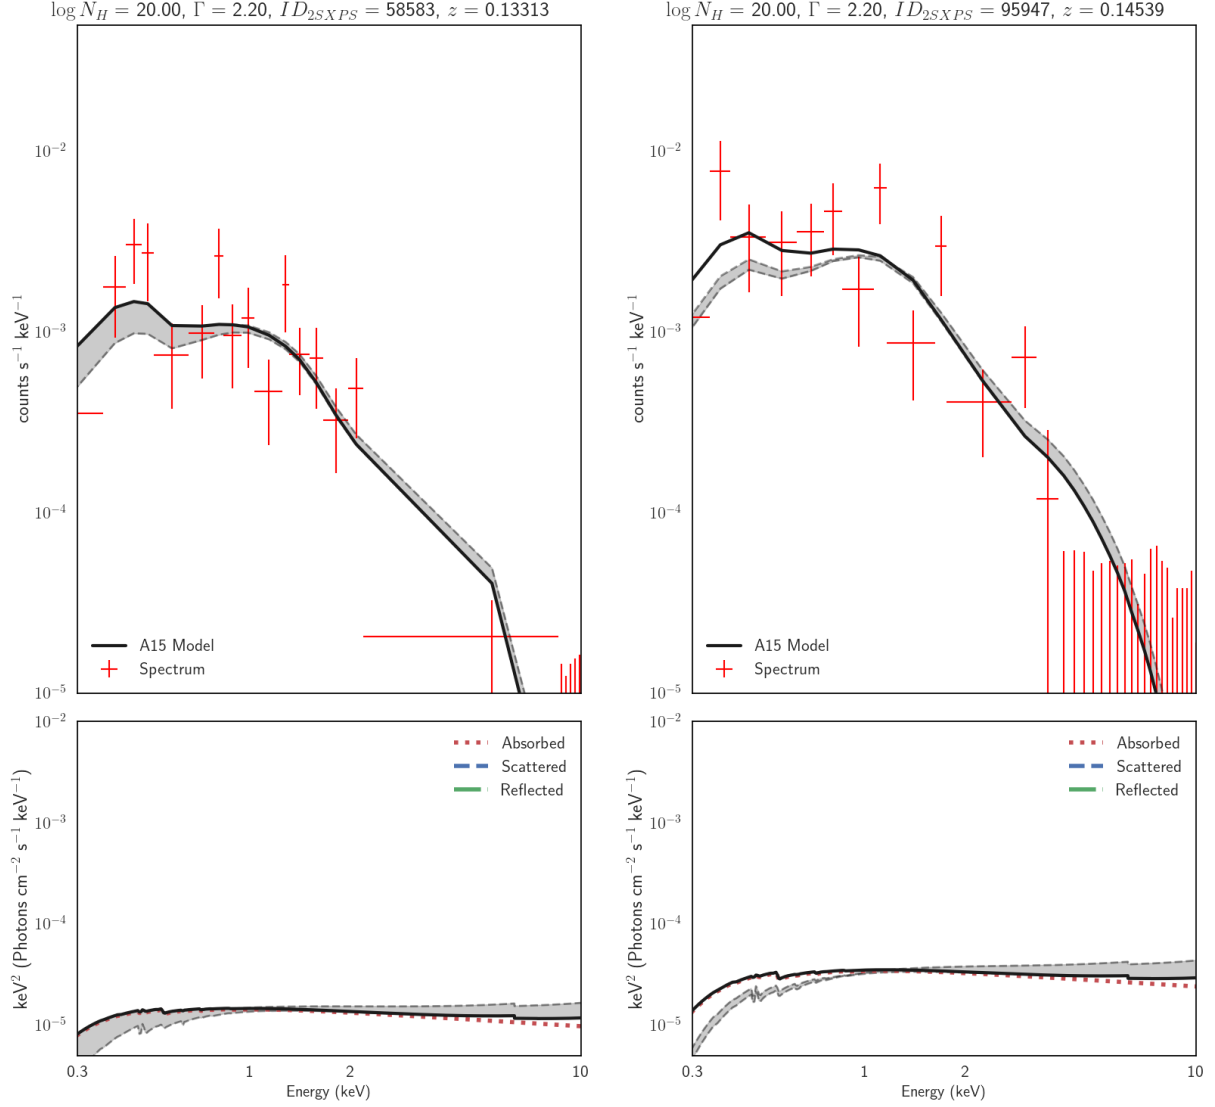


Figure 23. Sources classified as Compton-thick by the KNN algorithm. Spectral fitting reveals that they are unobscured, but this may not be accurate as these sources have less than 100 photon counts.

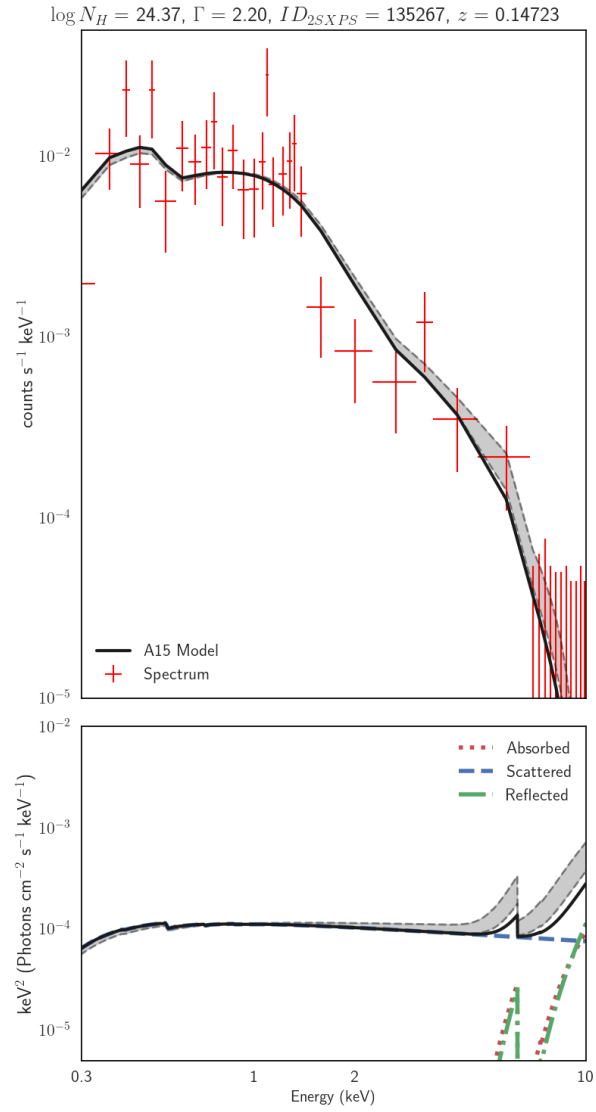


Figure 24. Source classified as Compton-thick by both the KNN algorithm and spectral fitting. The results of this may not be accurate as this source has less than 100 photon counts.