



Universiteit
Leiden
The Netherlands

X-raying extragalactic gas: warm-hot gas in the EAGLE simulations

Wijers, N.A.

Citation

Wijers, N. A. (2022, March 16). *X-raying extragalactic gas: warm-hot gas in the EAGLE simulations*. Retrieved from <https://hdl.handle.net/1887/3279147>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3279147>

Note: To cite this publication please use the final published version (if applicable).

X-raying Extragalactic Gas: warm-hot gas in the EAGLE simulations

Proefschrift

ter verkrijging van
de graad van doctor aan de Universiteit Leiden,
op gezag van rector magnificus prof.dr.ir. H. Bijl,
volgens besluit van het college voor promoties
te verdedigen op woensdag 16 maart 2022
klokke 16:15 uur

door

Nastasha Anna Wijers
geboren te Princeton, New Jersey, Verenigde Staten
in 1993

Promotores:

Prof.dr. J. Schaye

Prof.dr. J. Kaastra

Promotiecommissie:

Prof.dr. H.J.A. Röttgering

Prof.dr. J.F. Hennawi (UL and University of California, Santa Barbara)

Dr. A. Simionescu (SRON)

Prof.dr. S. Mathur (The Ohio State University)

Dr. C.-A. Faucher-Giguère (Northwestern University)

ISBN: 978-94-6419-441-8

Front cover:

X-ray emission from oxygen predicted for a part of the universe simulated in EAGLE, with strips showing X rays blocked by oxygen. The coloured parts are potentially detectable with the future space telescope Athena.

Specifically, this is the O VIII $K\alpha$ emission line surface brightness in a region of the EAGLE 100^3 cMpc³ volume at redshift 0.1, measuring 13.65×19.5 cMpc², with a depth of 6.25 cMpc. This region contains three galaxy groups (halo mass $M_{200c} = 10^{13} - 10^{14} M_{\odot}$, large blobs in the lower left quarter), one local-group-mass halo (halo mass $M_{200c} = 10^{12.5} - 10^{13} M_{\odot}$, top left corner), two haloes with masses similar to that of the Milky Way (halo mass $M_{200c} = 10^{12} - 10^{12.5} M_{\odot}$, bottom centre, each to the right of one of the galaxy groups), and a number of lower-mass haloes. The transition between colour and grayscale roughly marks the expected 5σ detection limit for this O VIII $K\alpha$ line with the Athena X-IFU, assuming extreme exposure times and spatial binning (10^7 arcmin² s). Horizontal strips show the O VIII ion column density. The transition from colour to grayscale indicates the detection limit for a planned survey with the Athena X-IFU.

Back cover:

The location of the hot (red) and cool/warm (blue) gas in the image on the front cover. Specifically, the hot gas has temperatures $> 10^{5.5}$ K, and the cool and warm gas has temperatures $< 10^{5.5}$ K. The brightness indicates the gas surface density. The stellar surface density is superimposed in transparent white.

Contents

1 Introduction	1
1.1 Motivation	1
1.2 Detecting the warm-hot gas	5
1.3 Methods	13
1.3.1 EAGLE	13
1.3.2 Post-processing	17
1.4 This thesis	22
1.4.1 Chapter 2: soft X-ray absorption lines in blind surveys	22
1.4.2 Chapter 3: soft X-ray and FUV absorption lines around galaxies	22
1.4.3 Chapter 4: soft X-ray emission lines around galaxies	23
1.4.4 Chapter 5: comparing EAGLE predictions to observed soft X-ray absorption	24
1.5 Outlook	25
2 The abundance and physical properties of O VII and O VIII X-ray absorption systems in the EAGLE simulations	27
2.1 Introduction	28
2.2 Methods	30
2.2.1 The EAGLE simulations	31
2.2.2 Column density calculation	32
2.2.3 Spectra and equivalent widths	33
2.3 Results	35
2.3.1 Column density distributions	35
2.3.2 Absorption spectra	38
2.3.3 Equivalent widths	40
2.3.4 The physical origin of the break	44
2.3.5 The effect of AGN feedback	46
2.3.6 Physical properties of the absorbers	48
2.3.7 Correlations between O VIII, O VII, and Ne VIII, O VI, and H I absorption	52
2.3.8 Correlation between absorber properties	54
2.4 Discussion	56
2.4.1 Detection prospects	56
2.4.2 Caveats	59
2.5 Conclusions	60
Appendices	61
2.A CDDF convergence	61
2.B Technical choices for the EW distribution	65
2.C Mock survey CDDFs	66

3 The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray and UV line absorption	69
3.1 Introduction	71
3.2 Methods	73
3.2.1 EAGLE	74
3.2.2 Galaxies and haloes in the EAGLE simulation	74
3.2.3 CGM definitions	76
3.2.4 The ions considered in this work	76
3.2.5 Column densities from the simulated data	77
3.2.6 Equivalent widths from the simulated data	78
3.2.7 Absorption profiles	80
3.3 Results	81
3.3.1 Ion properties	82
3.3.2 The baryonic content of haloes	82
3.3.3 Column density distributions and equivalent widths	88
3.3.4 Column density profiles	91
3.3.5 Halo gas as a function of radius	95
3.4 Detection prospects	99
3.4.1 Detection limits for different instruments	100
3.4.2 Halo detection rates	101
3.4.3 Extent of detectable absorption	101
3.5 Discussion and comparison to previous works	103
3.6 Conclusions	104
Appendices	106
3.A Measuring column densities and equivalent widths	106
3.B How to split the CDDFs	108
3.C Redshift evolution	110
4 The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray line emission	113
4.1 Introduction	114
4.2 Methods	116
4.2.1 EAGLE	116
4.2.2 The emission lines	117
4.2.3 Surface brightnesses	122
4.2.4 Galaxies and haloes	124
4.2.5 CGM definitions	125
4.3 Detectability	125
4.3.1 Overview of instruments	127
4.4 Results	131
4.4.1 Line emission in relation to haloes	131
4.4.2 Surface brightness profiles	134
4.4.3 Origin of the emission	138
4.5 Discussion	142
4.5.1 The EAGLE simulations	142
4.5.2 Literature comparison	144
4.6 Conclusions	145
Appendices	147

4.A Gas directly heated by feedback	147
4.B Numerical convergence	149
5 Comparing EAGLE simulation predictions to X-ray absorption line observations	151
5.1 Introduction	153
5.2 Methods	155
5.2.1 EAGLE	155
5.2.2 Column density calculations	156
5.3 A blind detection of two O VII absorbers?	157
5.3.1 The initial detections	157
5.3.2 A further investigation of the absorber environments	160
5.4 Soft X-ray counterparts to FUV absorbers	162
5.4.1 System 1: 3C 273, redshift 0.09	164
5.4.2 System 2: Ton S 180, redshift 0.046	167
5.5 Discussion	171
5.6 Summary	172
Bibliography	172
Nederlandse samenvatting	183
English summary	193
Publications	201
Curriculum Vitae	205
Acknowledgements	207

1

Introduction

In this thesis, I will discuss the hot, diffuse gas around and between galaxies. Specifically, I have used numerical simulations to predict the properties of this gas, and I have used those properties to predict specific observables: soft X-ray absorption and emission lines. Measuring this gas is challenging, but if we can observe and characterise it, we can learn much about the gas flows in and out of galaxies that regulate their formation and evolution. Observations of soft X-ray lines with future X-ray telescopes, such as Athena and XRISM, will enable us to do so.

1.1 Motivation

In the very early universe, 13.8 billion years ago (e.g., [Planck Collaboration et al. 2020](#)), matter was distributed almost perfectly smoothly. With time, as the universe expanded, the very small initial overdensities grew, driven by gravity. Dark matter and gas first collapsed along one dimension, forming sheets. Further collapse in a second dimension formed filaments, and collapse in all three dimensions formed roughly spherical haloes. The most massive haloes are nodes at the intersections of filaments. The underdense regions left between the structures form voids. [Fig. 1.1](#) illustrates the structure of this ‘cosmic web’. Halo densities are typically not much lower than ~ 100 times the average density of the universe, while filaments are typically a factor ~ 10 overdense.

The growth of these large-scale structures is driven by gravitational collapse, which operates in the same way on dark matter and gas. However, the gas does not behave quite like the dark matter. On large scales, the gas pressure (and around and before the era the cosmic microwave background was formed, radiation pressure) is the most important difference: this causes the gas to resist collapse in a way dark matter does not. However, gas can also cool radiatively. This cooling results from 2-body interactions between atoms, ions, and electrons, and is therefore more effective at higher densities. Within haloes, the denser gas can cool. It thereby loses pressure support, becomes compressed by surrounding gas, and cools further. The denser, lower-entropy gas sinks toward the halo centre, until it is supported by angular momentum (and remaining gas pressure). Eventually, stars form in this gas by further cooling and collapse, combined with self-gravity: a galaxy has formed at the centre of the halo. [Planelles et al. \(2015\)](#) give a review of structure formation in the universe.

Once stars form, their winds, radiation, and eventually for some, violent deaths as supernovae, affect star formation in their surroundings. They can drive gas back out of galaxies (e.g., the review by [Rupke 2018](#)), and out of their haloes (e.g., [Aguirre et al. 2001](#); [Mitchell](#)




Figure 1.1: The cosmic web, as created in the Millennium simulation (Springel et al. 2005), taken from <https://wwwmpa.mpa-garching.mpg.de/galform/virgo/millennium/>. The image shows dark matter, distributed in nodes, filaments, sheets (difficult to see in projection), and voids.

et al. 2020). Some of these feedback mechanisms may also trigger some star formation by compressing gas (e.g. Luisi et al. 2021). These effects of stars on star formation are known as (stellar) feedback. This feedback decreases the number of stars formed in galaxies; without feedback, galaxies would grow much more massive than they are observed to be (e.g., Cole 1991).

Accreting supermassive black holes in the centres of galaxies can similarly drive outflows from galaxies, and heat halo gas (see, e.g., Fabian 2012, for a review). This is known as AGN (active galactic nucleus) feedback. The exact mechanism for this is uncertain; the energy might be injected by jets, winds, or radiation from the black hole and its accretion disk. This feedback mode is mainly important for high-mass galaxies, with masses roughly equal to or greater than that of the Milky Way. In these galaxies, supernova feedback alone is no longer effective at suppressing star formation.

This ‘failure’ of the supernova feedback is thought to trigger the onset of AGN feedback (Bower et al. 2017): when supernova feedback becomes ineffective and gas builds up in the galaxy, this includes a build-up of gas around the black hole, which grows more massive by accreting some of this gas. Part of the energy of this accreted gas is not swallowed by the black hole, but is instead injected into the surrounding gas.

The galactic outflows produced by stars, and by the massive black holes at the centres of galaxies, also affect the gas around galaxies. They drive gas from galaxies into the halo, drive some gas out of the halo altogether, and may prevent new gas from accreting (e.g., Wright et al. 2020). The galactic outflows may also fall back onto the central galaxy in a process known as recycling (e.g., Mitchell et al. 2020). However, the root source of gas for star formation remains the halo and the cosmic web.

Therefore, understanding the gas around galaxies is crucial to understanding how galax-




Figure 1.2: A cartoon of the CGM and the baryon cycle. This is fig. 1 from [Tumlinson et al. \(2017\)](#). The CGM gas comes from inflows from the IGM (blue) and outflows from the galaxy that may escape the halo (orange) or fall back onto the galaxy (recycling, pink). The purple gas illustrates the diffuse, virialized gas close to hydrostatic equilibrium. The image does not include more isotropic accretion or gas accreted in mergers. The outflows and the bulk of the CGM are multi-phase, containing cool ($\sim 10^4$ K) and warm-hot ($\sim 10^5$ – 10^7 K) gas. Inflowing filaments may remain cool, or may be disrupted through their interaction with the warm-hot CGM (see, e.g., [Mandelker et al. 2020](#)).

ies form and evolve. The state of this halo gas is sensitive to the flows of gas into and out of haloes and galaxies ([Mitchell & Schaye 2021](#)); these gas flows are called the ‘baryon cycle’. Fig. 1.2 illustrates these different gas flows.

Depending on the mass of the system, the gas in these haloes, around a central galaxy, is called the intra-cluster medium (ICM), intra-group medium (IGrM), or circumgalactic medium (CGM). Clusters are haloes with total masses $\gtrsim 10^{14} M_{\odot}$, groups have masses $\sim 10^{12}$ – $10^{14} M_{\odot}$, and the CGM typically describes gas in lower-mass systems. The Milky Way halo is estimated to have a mass in the range $\approx 10^{12}$ – $10^{12.5} M_{\odot}$ (e.g., [Zaritsky 1999](#); [Li & White 2008](#); [Monari et al. 2018](#); [Li et al. 2020](#)), with many studies favouring $\approx 10^{12} M_{\odot}$. Note that these categories are not strictly defined, and systems are often classified based on the galaxies that form them, rather than total halo masses. For example, CGM typically refers to the gas in haloes of isolated galaxies. Throughout this thesis, I study haloes over a wide range of masses, so I will often use the term ‘CGM’ to refer to halo gas more generally, including IGrM and ICM. The gas outside haloes, in the filaments, sheets, and voids of the cosmic web, is known as the intergalactic medium (IGM). For a review of the CGM, see [Tumlinson et al. \(2017\)](#).

In haloes with masses $\gtrsim 10^{11.5}$ – $10^{12} M_{\odot}$, the halo gas is expected to include a volume-filling, warm or hot phase, at temperatures $\gtrsim 10^{5.5}$ K (e.g., [Dekel & Birnboim 2006](#); [Kereš et al. 2009](#); [Van de Voort et al. 2011](#); [Correa et al. 2018](#)). Although the (inner) CGM in Milky-Way-mass haloes may have a considerable amount of rotational support ([Oppenheimer 2018](#)), the warm-hot gas in these and more massive haloes is generally expected to be close to hydrostatic equilibrium, where thermal pressure balances gravity within the halo. The gas is heated to these temperatures by shocks as it falls from the intergalactic medium onto

haloes. Shock-heating can similarly heat gas falling into filaments.

The diffuse gas at temperatures between $\sim 10^5$ and 10^7 K is typically called the warm-hot intergalactic medium (WHIM). In this thesis, I will typically use this term to refer only to intergalactic gas, though some use it to refer to all warm-hot diffuse gas, including CGM, IGrM, and ICM. Diffuse gas hotter than 10^7 K does exist, but it is expected to mostly be part of the ICM and IGrM (e.g. [Tuominen et al. 2021](#); [Wijers et al. 2020](#)).

A typical definition of halo mass is M_{200c} . It is defined by a radius R_{200c} enclosing a sphere within which the average density is $200\rho_{\text{crit}}$. The critical density is

$$\rho_{\text{crit}} = \frac{3}{8\pi G} H(z)^2, \quad (1.1)$$

where G is Newton's constant, and $H(z)$ is the Hubble expansion rate at redshift z . This is the density of a flat universe at that redshift. The halo mass M_{200c} is the mass enclosed within R_{200c} . Other halo mass definitions typically use other values of the average enclosed density, but follow the same principle. The value of $200\rho_{\text{crit}}$ is derived from an analytical calculation for the gravitational collapse of a spherical overdensity in an Einstein-de-Sitter universe, where all the mass is in (dark) matter, with a density equal to the critical density. This yields an overdensity for the halo of $18\pi^2$.

For a halo of mass M_{200c} , with a hydrostatic halo, we can estimate the temperature of the gas at R_{200c} using the virial theorem. This means the thermal motions of the gas support it against gravitational collapse. This gives us the virial temperature

$$T_{200c} = \frac{\mu m_{\text{H}}}{3k} G M_{200c}^{2/3} (200\rho_{\text{crit}})^{1/3} \quad (1.2)$$

where k is the Boltzmann constant, and μ is the mean molecular weight in units of m_{H} , the mass of the hydrogen atom. The value $\mu = 0.59$ describes primordial, fully ionized gas. We see that the temperature of the warm-hot halo gas increases with halo mass. For an $M_{200c} = 10^{12} M_{\odot}$ halo at $z = 0$, $T_{200c} = 10^{5.7}$ K. At $z = 0.7$, $T_{200c} = 10^{5.8}$ K for a $M_{200c} = 10^{12} M_{\odot}$ halo. The difference is small enough that I generally will not mention the redshift when I give halo masses. Note that the temperature of the volume-filling halo gas tends to rise towards halo centres.

Eq. [1.2](#) provides an analytical estimate of the temperature of the warm-hot CGM gas. However, much about the (warm-hot) CGM and WHIM is still unknown. An important reason for this is that warm-hot, diffuse gas is difficult to detect. We will examine why in [§1.2](#). A consequence of this difficulty is that the (total) mass of the CGM is unknown, since it has been difficult to measure observationally. The CGM mass is typically characterized as the CGM mass fraction $(M_{\text{CGM},200c} / M_{200c}) / (\Omega_{\text{b}} / \Omega_{\text{m}})$, where $M_{\text{CGM},200c}$ is the CGM gas mass within R_{200c} , and Ω_{m} and Ω_{b} are the average densities of all matter and baryons respectively, normalized by ρ_{crit} . Theoretical predictions for the CGM mass fraction differ. For example, [Davies et al. \(2020\)](#) showed that, in the IllustrisTNG simulation (100-1; [Pillepich et al. 2018](#)), the CGM mass fraction is small around Milky-Way-like galaxies, and larger at lower and higher masses. Meanwhile, the EAGLE simulation CGM gas fractions rise monotonically with halo mass.

[Oppenheimer et al. \(2020b\)](#) predict that observations with eRosita should be able to discriminate between these two predictions, by measuring the soft X-ray emission from this gas. [Mitchell & Schaye \(2021\)](#) found that the amount of gas in the CGM is a good probe of the baryon cycle. Specifically, the difference between the EAGLE and TNG predictions at masses below that of the Milky Way is a result of different models for the supernova feedback. In EAGLE, this feedback removes gas not only from the central galaxy, but also from

the halo entirely. In TNG, the outflows leave the central galaxy, but do not tend to leave the halo, and often ‘recycle’ back into the galaxy.

Since both simulations produce broadly realistic galaxy populations (e.g., Schaye et al. 2015; Pillepich et al. 2018), characterising the CGM can provide important information on how star formation in galaxies is regulated. The WHIM can also be important here. We know some gas is ejected from haloes, because we find metals, produced in galaxies, outside galaxy haloes (e.g., Aguirre et al. 2001; Booth et al. 2012). The distribution of these metals, and in particular how much metal-enriched gas reaches how far from haloes, will also be important in understanding the baryon cycle.

1.2 Detecting the warm-hot gas

In order to learn more about the warm-hot WHIM and CGM, we must detect it. There are a number of strategies for this, which are sensitive to different gas properties. I will give a brief overview of fast radio burst (FRB) dispersion measures, the Sunyaev-Zel’dovich effect, X-ray emission, and UV and X-ray absorption lines. In the thesis, I will investigate warm-hot gas detections specifically with soft X-ray (and a few UV) absorption and emission lines.

Fast radio burst (FRB) dispersion measures are sensitive to the total electron column density between the observer and the burst. These are therefore sensitive to the WHIM and warm-hot halo gas, where hydrogen and helium is fully ionized through collisions between atoms/ions and electrons. We call this ‘collisional ionization’. However, gas can also be ionized by UV and X-ray photons. Throughout the Universe, there is a ‘background’ flux of UV and X-ray photons, known as the UV/X-ray background. This background is sufficient to ionize the cool IGM gas. This ionization mechanism is called ‘photo-ionization’. This means FRB dispersion measures are sensitive to the warm-hot CGM and WHIM, but also to the rest of the IGM. Note that only hydrogen in dense gas (compared to typical CGM, $n_{\text{H}} \gtrsim 10^{-3} - 10^{-2} \text{ cm}^{-3}$; Rahmati et al. 2013) can remain largely neutral in this radiation field, because the outer layer of a dense cloud will absorb enough of the radiation to ‘shield’ the inner parts. In the IGM, at densities n_{H} (well) below 10^{-5} cm^{-3} at $z = 0$, hydrogen therefore has neutral fractions $\ll 1$.

FRB dispersion measures are therefore interesting probes of the distribution of ionized gas. The scatter in the dispersion-measure-redshift relation can be used to constrain the distribution of ionized gas on large scales (e.g., McQuinn 2014; Walters et al. 2019), and correlations between dispersion measures and galaxy positions along lines of sight to FRBs can be used to measure the ionized gas mass in galaxy haloes (e.g., Ravi 2019).

Macquart et al. (2020) have already used a small number of dispersion measures from FRBs with known host galaxies, and therefore known redshifts, to measure the average density of ionized gas. The error bars on these measurements are still large, due to the small number of measurements, and systematic uncertainties about the electron column density contributed by the FRB host galaxy and our Galaxy. However, the measured density was consistent with the baryon density measured from the cosmic microwave background (CMB).

The Sunyaev-Zel’dovich (SZ) effect comes in two main flavours. The thermal SZ (tSZ) effect is sensitive to the electron pressure, integrated over the line of sight. It is therefore specifically sensitive to hot gas. The kinetic SZ (kSZ) effect is sensitive to the electron bulk momentum; if the velocity of a system is known, this effect can be used to infer the electron density. The tSZ effect is the easier of the two to measure. It has been used to measure clusters, and with stacks of massive galaxy pairs, it has been used to detect hot gas in massive filaments (e.g., de Graaff et al. 2019; Tanimura et al. 2019). Lim et al. (2020) have used the ki-

netic and thermal SZ effects together to constrain models of halo gas mass and temperature, over a range of halo masses extending down to $M_{500c} \approx 10^{12.3} M_{\odot}$ ($M_{200c} \approx 10^{12.4} M_{\odot}$). Measurements of gas around Milky-Way-like galaxies remain difficult with current instruments, due to their low spatial resolution. For a review of SZ measurements, see [Mroczkowski et al. \(2019\)](#).

Absorption and emission lines can carry a wealth of information. Depending on which lines can be detected, and how well lines can be resolved, a set of absorption and/or emission lines can be sensitive to the temperature, density, metallicity, turbulence, and kinematics of different, identifiable gas phases.

Temperatures can be constrained from the widths of the lines, or the relative abundances of ions. The lines are broadened due to the thermal and turbulent motions of ions. The turbulent and thermal components of the line widths can be disentangled by comparing the widths of lines from different elements in the same gas phase: thermal velocities scale as $m^{-1/2}$, where m is the mass of the ion, while turbulent velocities are shared by all ions in the same phase. The widths of the different lines can therefore be used to solve for the temperature and turbulent velocity dispersion.

The relative strength of different absorption and emission lines is sensitive to the gas temperature, because the temperature determines the rate and energy of electron-ion collisions, and therefore the (collisional) ionization and excitation rates of the ions and atoms. These rates affect the relative abundances of the different ionization states of a given element, and therefore the relative strengths of the lines the different ions produce.

Different gas phases can often be identified by the (slightly) different redshifts of different sets of absorption lines. These redshifts, compared to a galaxy redshift, provide information on the bulk motion of the gas.

The column densities of metal lines are all proportional to the product of the gas column density and metallicity. This means that there can be a degeneracy between the density and metallicity of a gas cloud. This degeneracy can be broken in a number of ways. I will describe a few methods that have been used in studies of the CGM.

Firstly, in UV spectra, the hydrogen column density can sometimes be measured. This requires measuring H I absorption lines, and modelling the fraction of hydrogen that is ionized. However, it is not always possible to measure the H I column density. Depending on how many Lyman series lines can be measured, saturation of the absorption lines may mean that only a lower limit on the H I column density can be given. The damping wings of lines can often provide an upper limit or a measurement in these cases. Additionally, for warm-hot gas phases, hydrogen and helium are often fully ionized, meaning UV data will not constrain the hydrogen column density. At the cooler end of that temperature range, the difficulty often lies in the width of the lines, rather than their column density alone: in modelling the spectrum, wide absorption features are often degenerate with the model for the unabsorbed continuum.

Secondly, in hot gas, like that found in the ICM, the (electron) density can be measured from the X-ray continuum emission instead. However, in the warm/hot gas in groups and around isolated galaxies, the halo gas emission is dominated by emission lines (e.g., [Werner & Mernier 2020](#)), and this measurement will not be possible because the continuum emission is too weak.

Thirdly, the combination of line absorption and emission can be used to break the degeneracy between density and metallicity. Emission depends on the (line-of-sight integrated) product of metallicity and *squared* density, and absorption on the (integrated) product of density and metallicity. Given the luminosity and column density of a gas cloud, it is therefore

possible to solve for the density and metallicity separately. However, this only works if the absorption and emission come from the same gas. In multi-phase media like the CGM, this is not necessarily the case.

Finally, if the ionization state of a plasma is primarily determined by the ionizing radiation field, the relative ion abundances are sensitive to the gas density as well as its temperature. In this case, determining the density does require knowing the radiation field the gas is exposed to. This radiation field is, however, somewhat uncertain. For example, different groups find different spectra for the UV/X-ray background (e.g., [Haardt & Madau 2012](#); [Khaire & Srianand 2019](#); [Faucher-Giguère 2020](#)).

In galaxy clusters, X-ray emission has been used to constrain the temperature, density, metallicity, and turbulence of the ICM (see, for example, the review by [Werner & Mernier 2020](#)). Metallicity measurements in this gas depend on measuring the continuum emission (bremsstrahlung), as the emission lines at these temperatures only come from metal lines. That is because hydrogen and helium are fully ionized in the hot ICM gas.

Some X-ray emission has also been detected from the IGrM. However, around isolated, spiral galaxies, detections of X-ray emission are mostly limited to the regions closest to the central galaxy (e.g., [Bogdán et al. 2015](#)), where gas densities are highest. However, [Das et al. \(2020\)](#) did detect emission up to ≈ 200 pkpc from such a galaxy.

Most of what we know about the CGM comes from absorption lines instead. Specifically, a number of surveys have measured far ultraviolet (FUV) absorption lines from the extragalactic CGM (e.g., [Tumlinson et al. 2011](#); [Johnson et al. 2015](#) [2017](#)). These have provided a lot of information on the gas producing these lines, such as temperatures, metallicities, and kinematics. However, these lines are produced only by cool and warm ($\approx 10^4$ – $10^{5.5}$ K) gas (e.g., [Tumlinson et al. 2017](#) fig. 6), and therefore cannot be used to characterise the hotter CGM gas ($\approx 10^{5.5}$ – 10^7 K).

The absorption and emission lines for this gas mostly fall in the soft X-ray band, at 0.3–2 keV (e.g., [Perna & Loeb 1998](#); [Hellsten et al. 1998](#); [Bertone et al. 2013](#)). The gas also produces some extreme UV (EUV) lines, but these are difficult to measure due to strong Galactic absorption at those energies. [Burchett et al. \(2019\)](#) did measure Ne VIII absorption lines, at redshifts $z \gtrsim 0.5$ where the line has redshifted into the FUV range. I will discuss Ne VIII absorption in the CGM, as well as O VI (FUV) absorption, which can be produced in warm ($\sim 10^{5.5}$ K) gas, in [Chapter 3](#)

However, my focus is on soft X-ray lines. In the Milky Way, such absorption and emission lines have been detected (e.g., [Bregman & Lloyd-Davies 2007](#); [Gupta et al. 2014](#); [Miller & Bregman 2015](#); [Das et al. 2019](#)). These have mainly been O VII and O VIII lines. These measurements, in many different lines of sight, have been used to measure the temperatures of different gas phases (e.g., [Kuntz & Snowden 2000](#); [Das et al. 2019](#)), to constrain the warm-hot gas density and metallicity (e.g., [Miller & Bregman 2015](#)), and to measure its rotation ([Hodges-Kluck et al. 2016](#)). There are, however, some uncertainties about where this gas lies relative to our Galaxy (e.g., [Bregman & Lloyd-Davies 2007](#); [Gatuzz & Churazov 2018](#)). [Miller & Bregman \(2015\)](#) did fit a radial profile to emission and absorption line measurements along different lines of sight.

These absorption line measurements have been made using instruments on ESA’s *XMM-Newton* and NASA’s *Chandra* satellites ([Fig. 1.3](#)). Space-based telescopes are needed for X-ray observations because the atmosphere blocks X rays. On *XMM-Newton*, the reflection grating spectrometers (RGSs) were used. There are two RGS instruments, behind two of the three co-aligned telescopes of *XMM-Newton*. As the name suggests, the RGS spectra are measured by directing X rays onto a reflection grating. The photon positions are then




Figure 1.3: Two satellites currently used for soft X-ray line observations. We show an artist’s impression of *XMM-Newton* on the left, taken from <https://www.cosmos.esa.int/web/xmm-newton/xmm-earth>, and a similar image of *Chandra* by NASA/CXC and James Vaughan, taken from <https://chandra.harvard.edu/about/spacecraft.html>




Figure 1.4: An example of claimed soft X-ray absorption lines, which were not found by others in observations of the same source. This is a *Chandra* LETG spectrum of PKS 2155–304, with a fitted continuum (red line). The inset shows the data in a smaller wavelength range in black and, in red, the modelled continuum and the absorption lines Fang et al. (2002b) claimed to detect in this spectrum. This is their fig. 1.

measured with CCDs. The CCD energy measurements are used to attribute the photons to the correct spectral order.

The main instrument of interest on *Chandra* is the low-energy transmission grating (LETG) spectrometer. It can be placed in the path of the focussed X rays, and the diffracted X rays are recorded by the advanced CCD imaging spectrometer or high-resolution camera. As with the *XMM-Newton* RGS, the photon position then indicates its energy.

These same instruments have been used in attempts to detect soft X-ray absorption lines from the CGM and WHIM. However, these lines have usually only been detected at low significance, or the claims have been disputed by others. For example, Fang et al. (2002b) claimed to have detected such an extragalactic absorber in the spectrum of PKS 2155–304 ‘blindly’, i.e., without prior expectations for the redshift. However, observations of the same X-ray source by Nicastro et al. (2002), Rasmussen et al. (2003), and Cagnoni et al. (2004) did not find this absorber. Fig. 1.4 shows the spectrum Fang et al. (2002b) based their claim




Figure 1.5: A Ne IX He- α absorption line (Ahoranta et al. (2020)) found at $z = 0.09017$, the redshift of an O VI absorber in the FUV spectrum of the same source. This is fig. 1 of (Ahoranta et al. (2020)). The black crosses show a stack of *XMM-Newton* RGS1 and RGS2 observations of the quasar 3C 273. The spectrum is normalized to the continuum. The orange line marks the wavelength of this line at $z = 0.09017$. The red lines show the continuum (solid) and the absorption line model (dotted), fitted to the *XMM-Newton* RGS data, as well as *Chandra* LETG and high-energy transmission grating (HETG) data. The absorber was detected at 3.9σ significance, though this included additional data from the *Chandra* LETG and HETG, and an O VIII line at the same redshift.

on. (Nicastro et al. (2005)) similarly claimed to find two absorbers in the sightline to Mrk 421, using the *Chandra* LETG. However, (Kaastra et al. (2006)) did not find these absorbers in *XMM-Newton* observations of this source, and argued that the (Nicastro et al. (2005)) absorption signals were not statistically significant because the number of redshift trials had not been accounted for.

Others have made use of prior information on the redshifts of potential soft X-ray absorbers. (Mathur et al. (2003)) searched the spectrum of H1821+643 at the redshifts of six known O VI absorbers, and found three, at $2-3\sigma$ significance. (Bonamente et al. (2016)) similarly searched for soft X-ray absorption lines at the redshifts of UV absorbers, and found one likely absorber, and (Ahoranta et al. (2020)) found a soft X-ray absorber at the redshift of an O VI line. Fig. 1.5 illustrates an absorption line found by (Ahoranta et al. (2020)). Note that for their data analysis, they fit spectra jointly, instead of analysing the stacked spectrum shown here.

(Kovács et al. (2019)) went a step further, and stacked spectra of different sources. They used the redshifts of Lyman α absorbers in the vicinity of massive galaxies to identify promising absorbers, and detected an O VII absorption line in the stacked spectra. For a review of searches for the WHIM using soft X-ray absorption lines, see (Nicastro et al. (2017)).

Recently, (Nicastro et al. (2018)) claimed the blind detection of two O VII absorbers, at redshifts 0.36 and 0.43. They used 1 Ms of observations of 1ES 1553+113, the brightest X-ray blazar. (Johnson et al. (2019)) measured the previously unknown redshift of the blazar using a novel approach, and found that one of the absorbers was intrinsic to the blazar host group. This was therefore not a ‘random’ absorber along the line of sight. Given the signal to noise ratio of the spectrum and the number of redshift trials, they found that there was a 4 per cent chance that a signal the size of the other absorber would arise from noise.

Given the source and exposure time, the (Nicastro et al. (2018)) measurements represent a fairly optimistic case for what can be measured with current instruments. Therefore, I expect

that we will not blindly detect large samples of soft X-ray absorption systems until the next generation of X-ray telescopes becomes available.

In X-ray spectroscopy, the main limiting factors are the spectral resolution and the sensitivity (effective area) of the instruments. The effective area is a factor that describes how many ‘counts’ (putative photon detections) result from a given flux, per unit exposure time. Its units are therefore those of area. A major difficulty in detections of X-ray absorption lines is that, with the effective areas of current instruments, it is difficult to collect enough photons to detect absorption and emission lines robustly.

The spectral resolution of each *XMM-Newton* RGS is 100–500, decreasing over the energy range 0.3–2.1 keV (den Herder et al. 2001), which translates to a velocity resolution of 600–3000 km s⁻¹. The RGS instruments have a maximum effective area of area of 140 cm², at 0.83 keV (den Herder et al. 2001).

The *Chandra* LETG has a spectral resolution of 120–800, decreasing over 0.3–2 keV, which translates to 370–2500 km s⁻¹ (Chandra X-ray Center et al. 2020). Its effective area is \approx 5–25 cm² at 0.3–2 keV. It therefore has a higher spectral resolution than *XMM-Newton*, but a lower sensitivity.

Typical (intrinsic) line widths for soft X-ray (absorption) lines are expected to be \sim 50–200 km s⁻¹ (e.g., Wijers et al. 2020). This is well below the resolutions of the currently available instruments, which has two main consequences. Firstly, if an absorber is detected, its line shape will be unresolved. This means absorber line widths are poorly constrained, and that temperatures and column densities are more difficult to constrain. Secondly, the absorption signal is ‘diluted’ by the nearby, absorption-free regions of the spectrum, which makes it more difficult to detect. Figs. 1.4 and 1.5 illustrate how this affects measurements of X-ray lines with current instruments. The shape of the line is not resolved, and only relatively strong absorption signals can be detected due to the noise, which is largely Poisson noise due to limited photon counts.

Future X-ray missions aim to improve the detectability of soft X-ray absorption and emission. In this thesis, I will discuss Athena and Lynx for absorption and emission, Arcus for absorption, and XRISM for emission lines. I will discuss the instruments further in the relevant chapters, but I give a brief overview below. Fig 1.6 illustrates the performance improvements of some of these instruments for soft X-ray line absorption. I estimate detection limits for emission lines in Chapter 4.

XRISM (X-ray Imaging and Spectroscopy Mission) is a JAXA and NASA mission, with contributions from ESA. It is due to be launched in 2022 or early 2023. The goal is to recover some of the capabilities promised by the Hitomi mission, which was unfortunately unsuccessful. Its science goals include studying X-ray emission from cluster outskirts, the IGrM, and the CGM of isolated (spiral and elliptical) galaxies. The instrument of interest for soft X-ray lines is the spectrometer, Resolve, which is similar to Hitomi’s soft X-ray spectrometer (SXS). It will likely have an effective area of 250 cm² at 1 keV, and a spectral resolution of 5 eV, based on the values achieved for Hitomi (XRISM Science Team 2020). The effective area is better than that of the *Chandra* LETG and *XMM-Newton* RGSs, but the spectral resolution is a bit worse (750–5000 km s⁻¹, decreasing from 0.3 to 2 keV). The XRISM Resolve instrument is primarily interesting for emission lines, though it should be more sensitive to absorption lines than the *XMM-Newton* and *Chandra* spectrographs, especially in the 1–2 keV range (Fig. 1.6).

Athena (Advanced telescope for high-energy astrophysics) is an L-class ESA mission, planned for launch in 2031. For absorption and emission, the instrument of interest is the X-IFU (X-ray integral field unit). An integral field unit is an instrument that produces combined

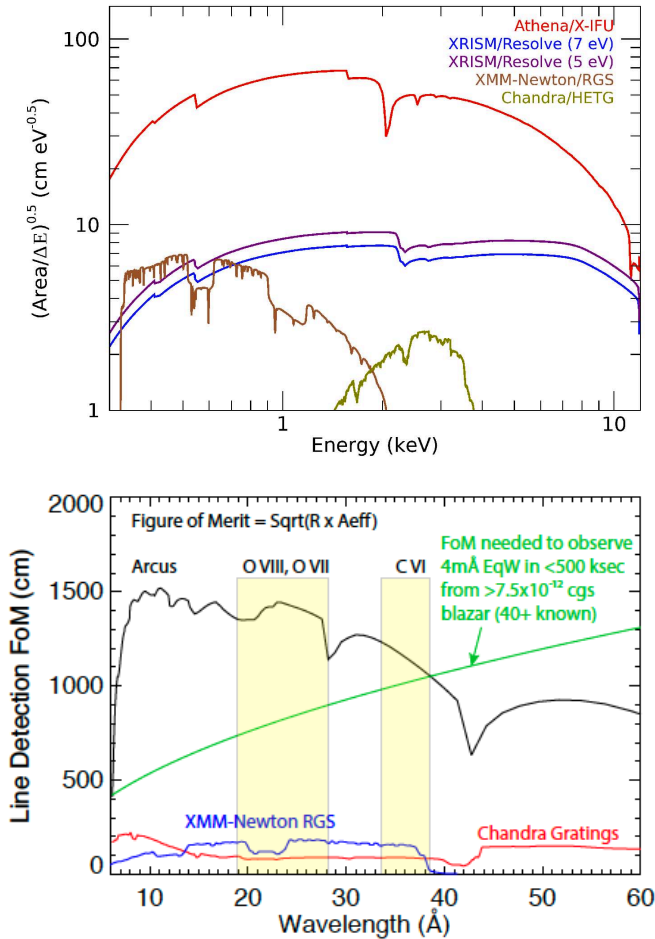


Figure 1.6: The expected performance of future instruments in soft X-ray line absorption. The top panel shows the figure of merit for weak absorption line detections for XRISM (Resolve), Athena (X-IFU), *XMM-Newton* (RGS) and *Chandra* (HETG). The 7 eV XRISM resolution is the requirement, while 5 eV is the Hitomi value. This is from Fig. 2 of [Guainazzi & Tashiro \(2018\)](#). XRISM may somewhat improve on the weak-line sensitivity of the *XMM-Newton* RGS, and the Athena X-IFU will have a much greater sensitivity than currently available instruments. The bottom panel shows a similar figure of merit (but in different units) for Arcus, compared to the *XMM-Newton* and *Chandra* spectrographs. This is fig. 1 of [Smith et al. \(2016\)](#). The bands indicate observable absorption lines at $z < 0.2$, and the green line indicates the limit for 5σ blind detections of lines with an equivalent width of 4 mÅ with the absorption line survey envisioned for Arcus.

images and spectra, with a spectrum at every image pixel. The X-IFU uses a transition edge sensor (TES) array to measure the position of each photon, and to measure its energy with high precision. In X-ray measurements, it is typical for individual photons to be recorded. The X-IFU will have a maximum effective area of $\approx 10^4 \text{ cm}^2$, at a photon energy of $\approx 1 \text{ keV}$, and a spectral resolution of $\Delta E = 2.5 \text{ eV}$ at energies $< 7 \text{ keV}$ (Barret et al. 2018). This effective area is almost two orders of magnitude larger than that of *XMM-Newton*, and the spectral resolution is similar to that of the *Chandra* LETG. The detection of O VII and O VIII absorption lines from the WHIM in the spectra of quasars and gamma-ray burst afterglows is one of the science goals of the mission.

Arcus was proposed as a NASA MIDEX mission. Its science goals include the characterisation of the CGM, IGrM, and ICM out to about R_{200c} and beyond, and measurements of the O VII and O VIII column density distributions through blind surveys. Its instrument is a spectrograph, using a reflection grating, and CCDs to detect the reflected light. The mission is designed specifically for high-resolution X-ray spectroscopy. It would have an effective area of 250 cm^2 and a spectral resolution of 2000 (2500) at wavelengths below (above) 21.60 \AA (Smith et al. 2016). This corresponds to a velocity resolution of $36\text{--}300 \text{ km s}^{-1}$ at $0.3\text{--}2 \text{ keV}$, which means absorption line shapes would begin to be resolved. The effective area is lower than that of Athena (though higher than that of the *XMM-Newton* and *Chandra* spectrometers), but the spectral resolution is higher. The anticipated minimum detectable column densities for the two instruments are similar, assuming background sources and observing times matching each instrument's respective proposed observing program.

Lynx is a NASA mission concept. Its launch was proposed for after 2036. For line emission, the instrument of interest is the Lynx X-ray microcalorimeter (LXM). It would have an effective area of 2000 cm^2 at 1 keV . The main array would have a spectral resolution of 3 eV , while a sub-array (the Ultra-High Resolution Array) would have a higher spectral resolution of 0.3 eV (The Lynx Team 2018). For absorption, the Lynx X-ray Grating Spectrometer (XGS) would provide a 4000 cm^2 effective area and a higher spectral resolution: at least ≈ 5000 , corresponding to a velocity resolution of 60 km s^{-1} . For Lynx, the goal is to use soft X-ray emission to characterise the CGM, and absorption to extend the CGM measurements out to R_{200c} , and to measure the WHIM in cosmic filaments.

For absorption lines, future instruments are designed to be more sensitive through increases in both the effective area and spectral resolution of the instruments. Other considerations in the instrument design are the field of view, spatial resolution, and instrumental backgrounds. They are not typically very important for absorption, which is usually measured against very bright point sources. However, these are important for X-ray emission. For a given total exposure time, a large field of view means more of an object can be imaged to a required depth. For surveys, the grasp (effective area times field of view) is therefore a useful figure of merit.

The spatial resolution helps imaging in two ways. The first is a straightforward improvement of the image resolution as the point spread function becomes smaller. The second is an improvement in background subtraction. A substantial part of the X-ray background consists of AGN point sources. A higher spatial resolution allows an observer to detect them and exclude these point sources from the analysis. At $z < 0.1$, the CGM will be at least marginally spatially resolved with the instruments I consider in this thesis.

Background and foreground signals are important complicating factors in the search for X-ray emission from the CGM. In general, these are a combination of an instrumental background, mainly from cosmic rays interacting with the detectors, and astrophysical backgrounds, primarily consisting of X-ray emission from the solar wind (solar wind charge

exchange), emission from our own local environment and our Galaxy (possibly including its halo), and (extragalactic) AGN. Since these signals can vary with time (such as solar wind charge exchange) and position on the sky (such as emission from the Milky Way and background AGN), X-ray data need to be fitted jointly for the target signal and backgrounds and foregrounds. This introduces systematic uncertainties, based on possible mis-modelling of the background signals, on top of the statistical noise.

Additionally, for studies of the CGM, X-ray emission from the central galaxy, such as from AGN, X-ray binaries, or a hot ISM phase, can ‘contaminate’ the CGM signal. A higher spatial resolution limits the extent of this effect. Some of these signals can also be distinguished from the CGM by their spectrum, but a hot ISM phase will produce a similar spectrum to the CGM: both are collisionally ionized plasmas.

To find absorption lines in X-ray spectra, a careful analysis of the data is required. X-ray spectra typically show many absorption lines from our own Galaxy. Lines can also originate in the host galaxy, and fast outflows from a quasar can imprint spectral features at redshifts different from the quasar itself. Besides these astrophysical signals, the instruments themselves imprint features on the spectra. The effective area of instruments depends on the energy of incoming photons, with various elements inside the instrument causing sharp jumps called absorption edges. All of these factors need to be accounted for when determining the continuum against which a line would stand out, and when determining the identity of absorption features.

1.3 Methods

1.3.1 EAGLE

The EAGLE (Evolution and Assembly of GaLaxies and their Environments) simulations are a set of cosmological, hydrodynamical simulations (Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016). The simulations follow the evolution of gas and dark matter from initial conditions at very high redshifts ($z > 100$) to the present day ($z = 0$), and include models for the formation and effects of stars and supermassive black holes. The EAGLE simulation data have been made publicly available (McAlpine et al. 2016; The EAGLE team 2017). I will discuss the EAGLE simulations specifically here. For a general review of galaxy formation simulations, see Vogelsberger et al. (2020).

The gravitational forces are calculated using the GADGET-3 TreePM scheme (Springel 2005), which uses approximations for the gravity from distant matter to speed up the calculation of those forces. Hydrodynamical forces are calculated using Smoothed Particle Hydrodynamics (SPH). This is a Lagrangian approach, where individual resolution elements (SPH particles) have a fixed mass, but change their position based on hydrodynamical (and gravitational) forces. Higher mass densities correspond to higher SPH particle densities, and the ‘size’ (smoothing length) of an SPH particle depends on the distances to its neighbours. EAGLE uses the Anarchy implementation of SPH (Schaye et al. 2015 appendix A; Schaller et al. 2015).

The simulations are cosmological in the sense that a large volume of the universe, relative to galaxies and their haloes, is simulated. The simulation we focus on is the main EAGLE simulation: a 100^3 cMpc^3 volume, with a gas mass resolution of $1.81 \times 10^6 M_\odot$ (Schaye et al. 2015). This gas mass resolution corresponds to 1504^3 SPH particles in the volume. The 100^3 cMpc^3 volume is small enough to limit the study of clusters with EAGLE: there are only 9 haloes with $M_{200c} > 10^{14} M_\odot$ at $z = 0.1$. However, there are 870 haloes with roughly the




Figure 1.7: The cosmic web, haloes, and a galaxy in the EAGLE simulation. This is fig. 1 of [Schaye et al. \(2015\)](#). The largest image shows gas in a $100 \times 100 \times 20 \text{ Mpc}^3$ volume of the EAGLE simulation at $z = 0$. The colour indicates the gas temperature, with blue being the coolest gas ($< 10^{4.5} \text{ K}$) and red the warm-hot gas ($> 10^{5.5} \text{ K}$). The intensity of the colour is set by the gas density. The insets show a smaller volume of gas, and a galaxy using stellar light in the u, g, and r SDSS bands.

mass estimated for the Milky Way halo ($M_{200c} = 10^{12} - 10^{12.5} M_{\odot}$).

In order to form (realistic) galaxies in these simulations, additional processes, which are not actually resolved in EAGLE, need to be modelled. Radiative cooling is needed for gas to condense into halo centres, where it can eventually form stars. This radiative cooling is modelled according to [Wiersma et al. \(2009a\)](#), who modelled metal line cooling by the 9 tracked metal species abundances in EAGLE. They assumed the gas was in ionization equilibrium, set by the temperature and density of the gas, and photo-ionization by a uniform (but redshift-dependent) UV/X-ray background. For this background, they used the [Haardt & Madau \(2001\)](#) model. Because the EAGLE resolution is insufficient to resolve, e.g., molecular clouds, molecular cooling processes are not included in this prescription.

Once gas has cooled, stars actually need to form in it. Star formation occurs on scales much smaller than those resolved in EAGLE, but requires dense gas. In EAGLE, star formation occurs in gas with densities $n_{\text{H}} \gtrsim 10^{-1} \text{ cm}^{-3}$ (on the $\sim \text{kpc}$ scales that can be resolved). The exact threshold depends on metallicity, where star formation can occur in lower-density gas if its metallicity is higher ([Schaye 2004](#)). This is motivated by the fact that higher-metallicity gas can cool more easily and therefore could collapse and form stars at lower densities. Metals increase gas cooling rates through their emission lines, and because they are needed to form dust, which shields gas from radiation and promotes H_2 formation.

The star formation rate is a function of the pressure of the gas. This function is designed to reproduce the Kennicutt-Schmidt relation, between star formation rate surface density and gas surface density in galaxies ([Schaye & Dalla Vecchia 2008](#)). At each time step, an

SPH particle meeting the requirements to form stars may be converted into a ‘star particle’. A star particle is a stellar mass resolution element of in EAGLE. It represents a ‘simple stellar population’: a group of stars formed at the same time, with a [Chabrier \(2003\)](#) stellar initial mass function. The initial mass of the stellar population is that of its progenitor SPH particle. The probability that an SPH particle is converted into a star particle is given by the star formation rate; at any time step, it is the ratio of the stellar mass that would form in that time step to the mass of the SPH particle.

The gas pressure in galaxies is determined by an equation of state, where the (minimum allowed) gas pressure is a function of its density alone, and cooling to temperatures below the equation of state values is not allowed. The equation of state is chosen such that the Jeans mass remains marginally resolved, which prevents an issue known as artificial fragmentation ([Schaye & Dalla Vecchia 2008](#)).

The stars return metals to their surrounding SPH particles according to [Wiersma et al. \(2009b\)](#). The model includes AGB winds and supernovae, both type Ia and core-collapse, and follows 9 individual metal species. Note that the SPH particle mass is increased when metals are added; the Lagrangian approach only means the mass is not changed by gravitational and hydrodynamical forces.

In order to form *realistic* galaxies, additional processes are required. Supernova explosions can prevent gas in galaxies from forming stars, e.g., by heating it or removing it from the galaxy. Without this ‘feedback’, gas will keep cooling and flowing into galaxies, causing them to form too many stars. Exactly how individual supernova explosions combine to drive galaxy-scale outflows is not currently known. In simulations like EAGLE, individual supernova explosions are not resolved. If the energy from individual explosions were added to neighbouring SPH particles at each time step, it would quickly be lost to radiative cooling, and not do any work on the gas. This problem is known as overcooling, and it occurs because, when the energy of a single supernova is injected into $10^6 M_{\odot}$ of gas, it only causes a small increase in temperature from the $\sim 10^4$ K of the warm ISM. At these temperatures, the radiative cooling rate is high, and the energy is quickly lost. Adding kinetic energy instead does not solve this problem, because shocks convert it to thermal energy, which is then radiated away.

Different simulations solve this problem in different ways, but all of them in some way prevent the energy injected by supernovae from being lost to radiative cooling. In EAGLE, the method of [Dalla Vecchia & Schaye \(2012\)](#) is used. For this feedback scheme, the energy produced by supernovae in a star particle over time is tracked. The energy that has been ‘saved up’ is compared to an energy budget for a single feedback event. Based on the ratio of the savings and the budget, a stochastic determination is made of whether a feedback event will occur. If it does, the entire budget is injected into a random neighbouring SPH particle as thermal energy. The large energy injection raises the temperature to a value where the radiative cooling rate is low, so the over-pressurised gas particle cools by expanding instead, and thereby does work on its neighbours. The energy of a single supernova feedback event is set to raise the gas temperature by $10^{7.5}$ K. This temperature was selected to be high enough to prevent overcooling, but low enough that feedback events do not become too rare to effectively suppress star formation.

Other stellar processes, such as winds and radiation, can also affect star formation. Modelling these processes has been found to be important in higher-resolution simulations ([Hopkins et al. 2018](#)). However, the saved-up energy from supernovae is allowed to vary in the calibration, and can be interpreted as encompassing all energy transferred from stars to gas. This energy is a free parameter, in part because it is uncertain how much of the supernova

energy remains, after, e.g., real radiative cooling, when it combines to cause the kpc-scale effects resolved in EAGLE.

For galaxies roughly as massive as the Milky Way, or more massive, supernova feedback no longer effectively suppresses star formation, even without overcooling (e.g., [Bower et al. 2017](#)). At these masses, supermassive black holes (SMBHs) become important in regulating star formation. Accreting supermassive black holes also inject energy into their surrounding gas. They can eject gas from galaxies, but they can also heat the halo gas, preventing it from cooling, sinking to the galaxy at the halo centre, and replenishing the galaxy’s ‘fuel’ for star formation.

The formation mechanisms for SMBHs are not very well known, so in EAGLE, a ‘seed’ black hole is simply added to any sufficiently massive halo that does not already have one. The black holes grow in mass by accretion and mergers, following [Rosas-Guevara et al. \(2015\)](#) and [Schaye et al. \(2015\)](#). AGN feedback faces an overcooling problem, just like the stellar feedback. In EAGLE, this is solved in a similar way: energy is ‘saved up’, based on the mass accreted by the black hole. This energy is injected stochastically into neighbouring gas particles, with a probability based on the saved-up energy and the amount required for a feedback event. An AGN feedback event in EAGLE raises the gas temperature by $10^{8.5}$ K. This value was a free parameter in the calibration.

The calibration of the EAGLE simulations is described by [Crain et al. \(2015\)](#). The idea here is that many of the subgrid processes are not fully understood, so the parameters of the models are set to values that produce realistic galaxy populations. Specifically, EAGLE was calibrated to reproduce the observed galaxy stellar mass function, reasonable galaxy sizes, and the stellar-mass-black-hole-mass relation at $z = 0.1$.

We note that in EAGLE, these same parameters reproduce the same observations less well when the resolution of the simulations is increased. This is somewhat expected, since the ‘subgrid’ (unresolved) processes are modeled at a particular scale in the calibration. When the resolution is increased, for example, some of the previously ‘subgrid’ radiative cooling in the supernova feedback scheme is now explicitly resolved, and effectively applied twice. Therefore, a separate calibration at an 8 times higher mass resolution was done, to the same observables. This simulation was run in a 25^3 cMpc³ volume. Generally, we investigate the convergence of results by comparing the main, 100^3 cMpc³ EAGLE simulation (L100 N1504) to this recalibrated, high-resolution simulation (L025N0752). This is a test of ‘weak convergence’, as described by [Schaye et al. \(2015\)](#), as opposed to ‘strong convergence’, which uses the same parameters for the feedback models at different resolutions.

Galaxy and halo data

[McAlpine et al. \(2016\)](#) describe the galaxies and haloes of the EAGLE simulation and the public release of the galaxy and halo catalogues. The haloes in EAGLE were found using a friends-of-friends (FoF) algorithm ([Davis et al. 1985](#)). For the dark matter particles, each was linked to every other particle within some fixed distance. In this case, that distance was 0.2 times the mean inter-particle spacing in the simulated volume. This results in haloes with a more or less fixed density at their edges. Stars, gas, and black holes are attributed to the same FoF halo as their nearest dark matter particle, or do not belong to any halo if the nearest dark matter particle does not.

Galaxies (‘subhaloes’) are identified using subfind ([Springel et al. 2001](#); [Dolag et al. 2009](#)). It works by identifying overdense regions within haloes, and then finding all particles that are gravitationally bound to each region. These self-bound overdensities are the subhaloes. The binding calculation includes kinetic and potential energy, as well as thermal energy for

gas particles. The most bound particle, with the lowest gravitational potential, defines the halo centre of potential. The galaxy containing this particle is the central galaxy (subhalo zero). Any particles in the halo that are not bound to any specific subhalo are attributed to the central galaxy.

Halo masses like M_{200c} , which I use throughout this thesis, are defined from the centre of potential. A sphere is grown from this point, and R_{200c} is the radius where the average density enclosed in the sphere is $200\rho_{\text{crit}}$.

For galaxy stellar masses, I typically use masses in 30 pkpc apertures. The choice of aperture has negligible effects on galaxy properties when their stellar masses are $< 10^{11} M_{\odot}$. At $M_{\star} = 10^{11} M_{\odot}$, the median effect of the aperture on the stellar mass is only 0.1 dex (Schaye et al. 2015). At larger stellar masses, the effects are larger. However, I mostly consider halo masses instead of stellar masses in this thesis. In Chapter 5 I select galaxies above a few minimum stellar masses. Two of these minimum masses are well below $M_{\star} = 10^{11} M_{\odot}$; the third minimum is $M_{\star} = 10^{11.2} M_{\odot}$, close to this limit. When I consider gas around galaxies of different stellar masses in Chapter 3 galaxies with $M_{\star} > 10^{11} M_{\odot}$ are a single category. Therefore, the effect of the aperture choice is small for the work in this thesis.

1.3.2 Post-processing

The quantities I calculated from EAGLE were primarily column densities, spectra, and surface brightnesses. From the spectra, the quantity of interest is the equivalent width of absorption lines. I will first give a brief introduction on absorption lines and surface brightnesses, then I will describe how I calculate them.

The equivalent width (EW) is the size of an absorption line. It is defined as

$$EW = \int \left(1 - \frac{F(\lambda)}{F_0(\lambda)} \right) d\lambda, \quad (1.3)$$

where λ is the wavelength, F is the measured flux, and F_0 is the unabsorbed, continuum flux. Since the spectra ‘stretch’ as they are redshifted, identical absorbers (in physical density units) at different redshifts will have different equivalent widths. The rest-frame equivalent width EW_{rest} , which a local observer would measure, is related to the equivalent width EW_{meas} that a redshift zero observer would measure as $EW_{\text{meas}} = (1+z)EW_{\text{rest}}$. Equivalent widths can also be measured in energy units, if the fluxes are integrated over spectral energy instead. In this space, the measured equivalent width scales as $(1+z)^{-1}$.

The column density is an integral of a particle density along the line of sight. This means it measures how much of an atom or ion there is in a given system, such as a filament, halo, or gas clump, along the line of sight to the background source. Column densities are always measured in particles per unit *physical* surface area. They are an important measure of absorption lines, related to their EWs.

Absorption can be measured against extended sources, such as galaxies, but in this thesis, I will focus on absorption against sources that are effectively point-like (typically quasars). In this case, the particle density in a thin column along the line of sight is an approximation of the ‘pencil-beam’ column density that would be measured in a spectrum.

The column density and equivalent width are related via the optical depth τ . It is proportional to the column density N , and $F(\lambda)/F_0(\lambda) = e^{-\tau}$. The constant of proportionality between N and τ depends on the atomic physics of absorption line in question, and the velocity distribution of the absorbing particles. In the ‘optically thin’ limit $\tau \ll 1$, $e^{-\tau} \approx 1 - \tau$, and $EW \propto N$.

For a single, uniform gas cloud, with a uniform particle density, temperature, and (e.g., turbulent) velocity dispersion, column density and equivalent width are related as follows. A cloud with column density N causes an absorption line $\tau(\lambda)$ as

$$\tau(\lambda) \propto \frac{N}{b} \exp\left(-\left(\frac{(\lambda - \lambda_0)c}{\lambda_0 b}\right)^2\right) \quad (1.4)$$

Here, b is the Doppler b -parameter. It is a measure of the Gaussian line of sight velocity distribution of the absorbing particles, and encompasses both thermal and non-thermal (e.g., turbulent) line broadening. The other parameters are the line centre λ_0 and the speed of light c . The line is further broadened by Lorentz broadening, which describes the intrinsic line width, from the uncertainty principle as applied to the lifetime (transition probability) of the excited state associated with the line.

The surface brightness (SB) is the flux per unit area on the sky; its units are, e.g., $\text{erg cm}^{-2}\text{s}^{-1}\text{arcmin}^{-2}$ or $\text{photons cm}^{-2}\text{s}^{-1}\text{sr}^{-1}$, where sr are steradians. Unlike column densities, surface brightnesses decrease with redshift, as $\text{SB} \propto (1+z)^{-3}$ (in photon units) or $\text{SB} \propto (1+z)^{-4}$ (in energy units). This redshift dependence comes from the combination of the luminosity distance (for the flux) and the angular size distance (for the sky area) dependences on redshift. The two distances cancel out so that only the redshift-dependence remains.

From the EAGLE simulations, I produce column density ‘maps’, where I divide the simulated volume into ‘slices’ along the line of sight, then measure the number of ions in thin rectangular prisms along the sightline. This produces a map of column densities as a function of position on the sky, in each line of sight slice. I similarly produced emission maps, where each map pixel measures a surface brightness instead of a column density. In addition to these maps, I also created some virtual absorption spectra. In this section, I will describe how I produced these maps and spectra, which form the basis for the analyses in the rest of the thesis.

Ion fraction and emission tables

To calculate column densities and surface brightnesses, I first calculate the number of ions or line luminosity in each individual SPH particle. I will first describe the process for ion numbers; the process for luminosities is conceptually similar.

The ion mass M_{ion} in an SPH particle is calculated as

$$M_{\text{ion}} = M_{\text{gas}} Z_{\text{elt}} f_{\text{ion}}, \quad (1.5)$$

where the mass of the SPH particle M_{gas} , and the mass fraction of the ion’s parent element Z_{elt} , are simply recorded in the EAGLE data. The fraction of parent element nuclei in the ionization state of interest, f_{ion} , is not tracked explicitly in EAGLE. Instead, I calculate this from tables created by [Bertone et al. \(2010a\)](#). These tables list ion fractions for different ions as a function of temperature, density, and redshift. EAGLE does record the temperatures and densities of SPH particles, and the redshifts of the snapshots, so I calculate ion fractions for each SPH particle by interpolating the tabulated values.

In Fig. [1.8](#) I show an example of such an ionization table, at $z = 0.1$, for O VII (the helium-like oxygen ion). The tables for other ions producing soft X-ray lines are qualitatively similar. We see that O VII is abundant in two regions of phase space. The first is at higher densities, marked ‘CIE’ in the figure. Here, the ion is in ‘collisional ionization equilibrium’. Ionization equilibrium is an assumption I make. It means that ionizing processes




Figure 1.8: An example of how ion fractions depend on the temperature and density of the gas. I show the O VII ion fraction, i.e., the ratio of O VII to oxygen, at redshift 0.1, from the tables of Bertone et al. (2010a). We indicate the mean baryon density \bar{n}_H with a dashed line. Typical warm-hot CGM densities are $\sim 10^{-5}$ – 10^{-3} cm^{-3} . We also indicate two regimes for which different ionization mechanisms dominate. In collisional ionization equilibrium (CIE), the ion fraction is independent of density. In photo-ionization equilibrium (PIE), at low densities, ionization by the UV/X-ray background dominates, and an ion is abundant over a wide range of temperatures, and a narrower range of densities.

are balanced by recombinations. It is a reasonable assumption for the column densities I will focus on (e.g., Yoshikawa & Sasaki 2006). In CIE, the ionization processes are dominated by ion-electron collisions. Recombinations also depend on ion-electron encounters, so ionization and recombination rates are both proportional to n_H^2 , and the equilibrium ion fractions are independent of density.

Another ionization process is also important: photo-ionization. Though, e.g., nearby stars and AGN can also ionize gas, I only consider ionization by the UV/X-ray background. This ionization rate (per unit volume) depends only on the ion density ($\propto n_H$), and thus becomes dominant over collisional ionization processes at low densities. In photo-ionization equilibrium (PIE), ions exist over a wide range of temperatures, and a smaller range of temperatures. Shielding effects, where the outer regions of a cloud absorb ionizing radiation and ‘protect’ the inner regions from photo-ionization, do not play a role at the densities where the ions we investigate exist in PIE.

I calculate the luminosity of an emission line in a single SPH particle in a similar way to its ion content. Bertone et al. (2010a) also produced tables of line emissivity ϵ as a function of temperature density, and redshift. I interpolate these tables to find the emissivity of an SPH particle, and then calculate its luminosity L as

$$L = \epsilon n_H^2 V \frac{n_{\text{elt}} / n_H}{(n_{\text{elt}} / n_H)_\odot}. \quad (1.6)$$

This equation defines the emissivity ϵ . Here, n_{elt} / n_H is the abundance of the element producing the line, and $(n_{\text{elt}} / n_H)_\odot$ is the solar element abundance assumed when creating the tables. The factors of density n_H and volume V account for the first-order dependence of luminosity on these quantities. For some soft X-ray emission lines, these tables contain an

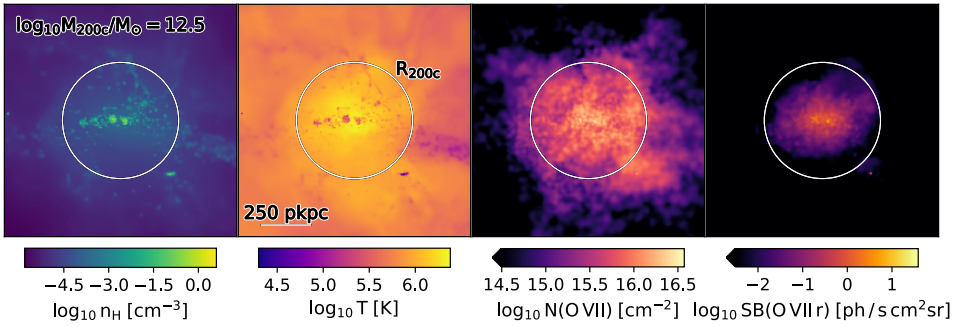


Figure 1.9: An illustration of the column density and emission maps created from EAGLE. The images show an $M_{200c} = 10^{12.5} M_{\odot}$ halo at $z = 0.1$, in a cube with side lengths of $4 R_{200c}$ (1186 kpc). The left two panels illustrate quantities used in the calculation of column densities and surface brightnesses. The leftmost panel shows the mass-weighted gas density, and the centre-left panel shows mass-weighted temperature. The centre-right panel shows O VII column density, and the rightmost panel shows the surface brightness of the O VII He- α -like recombination line.

error, so I use similar tables produced by [Ploeckinger & Schaye \(2020\)](#). I discuss this in more detail in Chapter [4](#).

The processes that ionize atoms and ions can also excite electrons, which produce emission lines when they decay back to lower-energy states. The emissivity as a function of temperature and density therefore looks qualitatively similar to the ion fraction in Fig. [1.8](#), and emission can also come from collisionally and photo-ionized gas.

[Bertone et al. \(2010a\)](#) generated their tables using CLOUDY v7.0, and used the [Haardt & Madau \(2001\)](#) model for the uniform, redshift-dependent UV/X-ray background. These assumptions match those made by [Wiersma et al. \(2009a\)](#) for the radiative cooling model used in EAGLE. The ion fraction and line luminosity predictions are therefore consistent with those assumed in the EAGLE simulations.

Some factors in these calculations are uncertain. As I mentioned previously, ionization equilibrium is an assumption, albeit a reasonable one. Aside from that, the UV/X-ray background, and therefore the photo-ionization rate, is not precisely known. Furthermore, these calculations require knowledge of various energy levels, collision cross-sections, etc., of the atoms and ions. These are fairly well known for hydrogen- and helium-like ions, with one or two electrons. However, for other ions, like Fe XVII, uncertainties remain substantial (e.g., [Gu et al. 2007](#); [de Plaa et al. 2012](#); [Bernitt et al. 2012](#); [Wu & Gao 2019](#); [Gu et al. 2019](#)).

The uncertainties in atomic physics also have more indirect effects on soft X-ray line detections. For example, the equivalent width of an O VII absorption line was recalculated by [Nicastrò \(2018\)](#) after new wavelength and oscillator strength measurements for the N II K- α triplet showed that the N II K- α complex originating from the Milky Way overlaps with the redshifted O VII line.

Column density and emission maps

In the column density maps, I approximate the column density one would measure from an absorption line in a spectrum. First, I divide the simulation volume into ‘slices’ along a given line of sight. In practice, I choose the line of sight to be parallel to the Z-axis of the simulation. This choice is convenient, and it is a random direction with respect to the structures in the

simulation, such as haloes and filaments.

Each slice defines a part of the line of sight that is considered to be a single absorption system, i.e., a group of absorption lines that come from the same structure. The standard slice depth I use is 6.25 cMpc. These slices are independent of the positions of actual structures in EAGLE, such as haloes and filaments, and ignore the effects of peculiar velocities of the gas. Therefore, these are only approximations of actual absorption systems. I tested the effect of slice depth on column density distributions and profiles in Chapters 2 and 3 and found the effects to be small.

Within each slice, I divide the volume into a grid of pixels perpendicular to the line of sight, along the X- and Y-axes. This divides the slice into long, thin columns. Typically, the pixels have an area of 3.125^2 ckpc².

I count the number of ions in each long, thin column by first calculating the number of ions in each SPH particle which lies in the slice, as described in §1.3.2. For each particle, EAGLE records a smoothing length, describing its size. I distribute the particles over the columns in the slice by assuming that the ions in each SPH particle are distributed according to a Wendland (1995) C2 kernel. The size of the kernel is the smoothing length of the particle.

I evaluated the kernel at the centre of each pixel, and normalized it according to those centre values to ensure mass conservation. The minimum smoothing length was set to half the pixel diagonal for the same reason. After this projection step, the column density is simply the number of ions in each pixel, divided by the (physical) area of the pixel.

The centre-right panel of Fig. 1.9 shows an example of a column density map created in this way, showing a small region around a single halo rather than a full $100 \times 100 \times 6.25$ cMpc³ slice. The rightmost panel shows an emission line surface brightness map, created in a similar way.

For the emission maps, the slices and pixels are defined in the same way. The luminosities of individual SPH particles are divided over the pixels in the same way as the ions in column density maps. The per-pixel luminosities are then divided by $4\pi D_L^2$, where D_L is the luminosity distance, and by the area each pixel spans on the sky, to obtain surface brightnesses. As with the column densities, the slices only approximate contiguous regions of emission that would define a line emitter.

Spectra

To make virtual absorption spectra from EAGLE, I used the SPECWIZARD code¹. I only used the ‘short spectra’ for the work in this thesis, where a spectrum is created for each individual ion. These spectra are periodic, reflecting the periodicity of the underlying simulation.

The starting point for these spectra are the same ionization tables used for the column density maps, described in §1.3.2. For these spectra, the ions in each SPH particle are assumed to follow a Gaussian distribution, with its size scaled by the smoothing length recorded by the EAGLE simulation. This is different from the Wendland (1995) C2 distribution assumed when creating the column density maps, but I found that, in practice, the exact shape of this distribution makes little difference.

For the spectra, the line of sight is initially drawn in position space, and divided into (1-dimensional) pixels. In each pixel, each SPH particle intersecting the line of sight contributes an ion column density calculated from its impact parameter and offset along the line of sight. Ion-weighted temperatures and line of sight velocities are also recorded for each pixel.

¹A version of this code can be found on github, at https://github.com/nastasha-w/specwizard_versions/tree/nastasha_quasar

In velocity space, the sightline is divided into pixels corresponding to those in position space, matched through the Hubble flow at the redshift of the simulation snapshot. Each pixel in position space produces an absorption line in velocity space according to eq. 1.4. Lorentz broadening is not included in these spectra. The line centre is determined by the position and peculiar velocity along the line of sight, and the line width is set by thermal broadening according to the ion-weighted temperature. Turbulent broadening is included through the ion-weighted peculiar velocities of the different position-space pixels; velocity dispersions within these pixels are not accounted for.

1.4 This thesis

1.4.1 Chapter 2: soft X-ray absorption lines in blind surveys

In this chapter, I focus on O VII and O VIII absorption. These two ions produce the strongest absorption lines expected in soft X rays. This chapter is focussed on blind surveys, where galaxy information is not included in the selection of sightlines. I predict column density distributions for O VII and O VIII, as well as EW distributions, and the resulting detection rates for surveys with Athena, Arcus, and Lynx. The incidence rate of absorbers decreases with column density. This decrease is particularly sharp at the highest column densities, $N \gtrsim 10^{16} \text{ cm}^{-2}$.

I also investigate the properties of the absorbers. Specifically, I study ion-weighted temperatures, densities, and metallicities as a function of column density. I use these to predict which absorbers, based on their column density, are collisionally ionized, and which are photo-ionized. The highest column density absorbers, $N \gtrsim 10^{16} \text{ cm}^{-2}$, are collisionally ionized, but at lower column densities, photo-ionization also plays a part in determining the ion fraction. These highest column densities arise in gas with overdensities $\gtrsim 10^2$, meaning this gas is most likely part of the CGM, and the transition from IGM to CGM absorption may be responsible for the ‘break’ in the column density distribution at $N \approx 10^{16} \text{ cm}^{-2}$.

The ion-weighted metallicity of the gas only has a weak dependence on column density: at column densities of $\approx 10^{14} - 10^{16.2} \text{ cm}^{-2}$, the O VII- and O VIII-weighted metallicity is mostly between 0.3 and 1 times the solar value. This is higher than typical for the IGM, which is responsible for most of the absorption at the lower end of that column density range. These high metallicities therefore result from a metallicity bias: relatively high metal ion column densities preferably arise in high-metallicity gas.

Finally, I study correlations between different ions. Correlations between column densities of O VII and O VIII and those of H I and O VI illustrate how UV absorbers can be used as signposts for soft X-ray absorption lines. Comparisons of ion-weighted temperatures and densities illustrate at which column densities ions probe similar gas, and where they do not. At high column densities, $N \gtrsim 10^{16} \text{ cm}^{-2}$, the different ions typically trace gas at different temperatures; each traces the gas at the temperature where its own ion fraction peaks in CIE.

This chapter was published as [Wijers et al. \(2019\)](#).

1.4.2 Chapter 3: soft X-ray and FUV absorption lines around galaxies

In this chapter, I predict how haloes of different masses ($M_{200c} = 10^{11} - 10^{14.5} M_{\odot}$) shape the column density distributions explored in chapter 2, and include additional ions in the

analysis: O VI, Ne VIII, Ne IX, and Fe XVII. I also examine the properties of these haloes, and how they shape column densities in the CGM. This analysis is focussed on redshift 0.1.

The highest column density absorption systems in the soft X-ray line column density distributions originate in the CGM. However, very high column density O VI and Ne VIII absorption systems can also arise from the IGM. High column density CGM absorption arises in haloes where T_{200c} is close to the temperature where the ion fraction f_{ion} peaks in CIE. These haloes also produce higher median column densities at impact parameters $\lesssim R_{200c}$ than hotter and cooler haloes.

Across halo masses, the ions tend to arise from gas close to the temperatures where their respective ion fractions peak. Exceptions are ions at the outskirts of the haloes, close to R_{200c} , and O VI. In these outskirts, and for O VI, the gas density is low enough that photo-ionization can produce the highly ionized species I study in this paper.

The ion-weighted density of these ions in the CGM is typically close to the volume-weighted density at the same distance to the central galaxy, indicating that the ions I study here trace the warm-hot, volume-filling phase of the CGM. However, the ion-weighted metallicity of the gas shows strong biases. The mass- and volume-weighted metallicities decline as a function of distance to the central galaxy, while the ion-weighted metallicities remain almost constant, at roughly solar metallicity. Overall, the CGM typically contains more metals than the interstellar medium.

The median column densities typically decline with distance to the central galaxy, and tend to drop rather sharply at $\sim R_{200c}$. In high-mass haloes, where T_{200c} is greater than the temperature where a particular ion fraction peaks, that ion's column density often shows a different behaviour within R_{200c} . Starting from the central galaxy, the median column density decreases with impact parameter, but it then peaks at $\sim R_{200c}$, or stops decreasing and forms a 'shoulder'. This behaviour is due to the fact that, in these haloes, the cooler halo outskirts produce higher ion fractions, outweighing the decline in the median density and metallicity of the gas. This results in a 'shell' of high ion density gas at relatively large radii, which causes a peak or shoulder in the median column density as a function of impact parameter.

With the planned X-ray mission Athena, O VII and O VIII absorption lines should be detectable out to $\sim R_{200c}$ around $M_{\star} \gtrsim 10^{10.5} M_{\odot}$ galaxies. With the proposed Arcus mission, detection rates should be higher close to R_{200c} , and with Lynx, these lines should be detectable up to R_{200c} around $M_{\star} \gtrsim 10^{10.0} M_{\odot}$ galaxies.

This chapter was published as [Wijers et al. \(2020\)](#).

1.4.3 Chapter 4: soft X-ray emission lines around galaxies

This chapter explores the surface brightnesses of many soft X-ray emission lines in the same vein as the column densities in Chapter 3. Unlike column densities or equivalent widths, minimum detectable emission line surface brightnesses are not typically stated in proposals for X-ray missions. Therefore, I estimate them myself, from the resolution, sensitivity, and backgrounds of each instrument, provided by the team proposing or developing it. I consider XRISM, Athena, and Lynx here.

With XRISM, a few bright emission lines should be detectable from groups and clusters ($M_{200c} > 10^{13} M_{\odot}$). With Athena, the inner CGM from a wider range of halo masses, $M_{200c} \gtrsim 10^{12.5} M_{\odot}$, should be detectable. At the lower end of this mass range, the O VII He- α -like triplet and the O VIII K- α doublet are the most easily detected lines. O VIII might be detectable out to R_{200c} for groups and clusters. With Lynx, the inner CGM of $M_{200c} = 10^{12.0} - 10^{12.5} M_{\odot}$

haloes may be detectable. These observations would require long exposure times (~ 1 Ms), and large spatial bins ($\sim 1\text{--}10$ arcmin²)

Soft X-ray line emission mostly comes from haloes, and surface brightnesses that may be detectable with the three missions I consider are very rare outside haloes. For the Fe L-shell lines I investigate, and the He- α -like lines, the surface brightnesses peak in haloes where T_{200c} is similar to the temperature where the emissivity of the emission line peaks. However, for the K- α lines, the emissivity decreases less strongly with temperature above its peak, and the increase of surface brightness with halo mass only slows above that halo mass, with a possible decline at $M_{200c} \gtrsim 10^{14} M_{\odot}$.

Line emission mostly traces gas at temperatures close to the line emissivity peak, regardless of the typical gas temperatures in a halo. This is similar to the ion-weighted temperature trends of Chapter 3. K- α doublets in haloes with T_{200c} above their emissivity peak form an exception: their ion-weighted temperatures are similar to the mass- and volume-weighted temperatures of the haloes. Like the ions, the luminosity-weighted metallicities show a strong bias towards high metallicities: the luminosity-weighted metallicities remain $\approx 0.3\text{--}1$ times the solar value Z_{\odot} out to at least $3 R_{200c}$ across halo masses, while the mass-weighted values decline with distance to the central galaxy, and are $\approx 0.03\text{--}0.1 Z_{\odot}$ at R_{200c} . Unlike the ions, line emission traces densities higher than those of the volume-filling CGM phase. This is because emission depends more strongly on gas density ($\propto n_{\text{H}}^2$) than the ion density does ($\propto n_{\text{H}}$), so it is more strongly biased towards high densities.

This chapter is submitted for publication in MNRAS.

1.4.4 Chapter 5: comparing EAGLE predictions to observed soft X-ray absorption

In this chapter, I compare EAGLE predictions of the column densities of soft X-ray absorbers to a few existing detections of extragalactic soft X-ray absorption lines (outside, e.g., quasar outflows). Overall, EAGLE seems to match the observations reasonably, with a few possible issues. The combinations of column densities of different ions found by [Ahoranta et al. \(2020, 2021\)](#) are plausible in EAGLE, if selection effects on the X-ray column densities are accounted for. Given the measured column densities, the temperatures and multi-phase nature of the detected absorption systems agree with EAGLE predictions.

The first issue in the comparison is that it involves a very small number of absorbers (4 in total), making it difficult to draw strong conclusions. For example, the uncertainties in the measured number density of O VII absorbers of [Nicastrò et al. \(2018\)](#) are large enough that they agree with all the numerical simulations the data were compared to. In fact, [Johnson et al. \(2019\)](#) argue that one of the absorbers cannot be used to measure a cosmological absorber number density, because they found it originates in the IGrM surrounding the blazar that was used as a background source for the absorption line measurements, and that the other absorber may be a spurious detection.

The second issue is that only the strongest absorption lines are detectable with currently available instruments, meaning that there are strong selection effects at play. The size of this effect can be difficult to quantify, due to an initial selection of promising sightlines in [Ahoranta et al. \(2020, 2021\)](#). The reasonable match between EAGLE and the absorption system of [Ahoranta et al. \(2020\)](#) becomes a more marginal agreement without the assumption that high X-ray column densities were selected.

Finally, based on the [Ahoranta et al. \(2021\)](#) absorber and the [Nicastrò et al. \(2018\)](#) absorber not associated with the blazar host group, EAGLE seems to underpredict the fraction of O VII

absorbers that lie at large distances ($> R_{200c}$) from galaxies. However, given the previous two issues, this does not represent a strong tension. Additionally, there are indications from UV data that the issue does not lie with an underestimation of the metal-enrichment of the IGM in EAGLE.

This chapter is based on [Nicastrò et al. \(2018\)](#), [Johnson et al. \(2019\)](#), [Ahoranta et al. \(2020\)](#), and [Ahoranta et al. \(2021\)](#), and in particular, on the EAGLE predictions I contributed to these works, and what we concluded from the comparisons.

1.5 Outlook

In this thesis, I have largely focussed on the power that future observatories will have to detect warm-hot, diffuse gas around and between galaxies. Many of these observatories will not be launched for a decade or more: the launch of Athena is planned for 2031, and Arcus and Lynx are still at the proposal stage. Though I have not investigated it in this thesis, HUBS (Hot Universe Baryon Surveyor; [Cui et al. 2020a](#)), a proposed CNSA mission with a launch around 2030, would also provide interesting information. Its spectral resolution would be similar to that of Athena, with a smaller effective area, but a larger field of view. It would be aimed at detecting soft X-ray emission. Like the other instruments I have discussed here, its science goals include the detection and characterization of the warm-hot CGM and WHIM.

XRISM, however, should be launched within 2 years, before April of 2023. The XRISM Resolve instrument should already allow us to detect soft X-ray line emission from the gas in galaxy groups and clusters. The eRosita telescope is currently collecting data, and [Oppenheimer et al. \(2020b\)](#) have shown that X-ray emission from the innermost CGM of Milky-Way-mass galaxies, and a larger fraction of the IGrM, should be detectable with its survey by stacking observations. It will be sensitive enough to constrain the physics of the CGM: the difference between the EAGLE and IllustrisTNG 100-1 simulation predictions for the CGM emission should be observable in this way. The difference between the predictions from these two simulations largely stems from the very different gas fractions of haloes hosting Milky-Way-mass galaxies in the two simulations. In IllustrisTNG, supernova feedback ‘recycles’ much of the gas ejected from the galaxy through the CGM, back into the halo ([Nelson et al. 2019](#)), while in EAGLE more of the ejected and swept-up gas escapes the halo altogether ([Mitchell et al. 2020](#)).

Large samples of fast radio bursts form a different avenue towards constraining the properties of the CGM. A large sample of FRBs (535) have been detected with the CHIME/FRB (Canadian Hydrogen Intensity Mapping Experiment Fast Radio Burst [The CHIME/FRB Collaboration et al. 2021](#)) project, and [Connor & Ravi \(2021\)](#) used this sample to detect a CGM/IGrM contribution to FRB dispersion measures. FRB dispersion measures should provide interesting constraints on the CGM mass, and potentially its density profile ([Ravi 2019](#)), complementary to the X-ray measurements. While X-ray absorption and emission lines are sensitive to metal-enriched gas at particular temperatures, dispersion measures are sensitive to all ionized gas, i.e., the whole warm-hot gas phase of the CGM and possibly some of the cool gas. This means that FRB dispersion measures are an excellent way to measure the mass of the ionized CGM, and its density profile. The total mass could be constrained by combining these measurements with UV absorption line measurements of the cool gas. The sensitivity of the X-ray lines to temperature and metallicity means they will be very useful in characterizing the hot phase. A set of emission or absorption lines can be used to constrain the temperature of the gas, and combinations of dispersion measures, absorption lines, and/or emission lines can constrain the gas density and metallicity.

When inferring CGM properties from those measurements, simulations will provide a useful guide. Firstly, direct comparisons between predictions from simulations employing different prescriptions for, e.g., feedback will be a good test of the different models. Secondly, even if simulations do not quite match the data, a comparison between the intrinsic properties of the CGM and a predicted observable will give a good sense of the biases at play. For example, I show in Chapters 3 and 4 that the ion-weighted metallicity is close to constant as a function of distance to the central galaxy, up to at least $3 R_{200c}$, and across three orders of magnitude in halo mass. However, the mass- and volume-weighted metallicities of the halo gas clearly decrease with distance to the central galaxy, and vary more with halo mass. This means that a metallicity trend measured by finding the metallicity of a sample of metal-line absorbers may not reflect the conditions of the typical gas at different distances from haloes of different masses.

To this end, theoretical predictions of these observables from different simulations and other models would be useful, to test the robustness and sensitivity of different observables to different processes. For example, Nelson et al. (2018) have made predictions of O VII and O VIII absorption from IllustrisTNG, and Bertone et al. (2010a) and Van de Voort & Schaye (2013) had previously made predictions of X-ray line emission in OWLS (Schaye et al. 2010a). Similar theoretical predictions of FRB dispersion measures from the CGM and IGM would help us make the best use of the information becoming available. Comparisons of simulations to specific combinations of observations would also help quantify the effects of various detection limits that affect these measurements, and would provide a more stringent test of the simulations than single observables.

2

The abundance and physical properties of O VII and O VIII X-ray absorption systems in the EAGLE simulations

Nastasha A. Wijers, Joop Schaye, Benjamin D. Oppenheimer, Robert A. Crain, & Fabrizio Nicastro.

2019, MNRAS, 488, 2947–2969

ABSTRACT

We use the EAGLE cosmological, hydrodynamical simulations to predict the column density and equivalent width distributions of intergalactic O VII ($E = 574$ eV) and O VIII ($E = 654$ eV) absorbers at low redshift. These two ions are predicted to account for 40 % of the gas-phase oxygen, which implies that they are key tracers of cosmic metals. We find that their column density distributions evolve little at observable column densities from redshift 1 to 0, and that they are sensitive to AGN feedback, which strongly reduces the number of strong (column density $N \gtrsim 10^{16}$ cm $^{-2}$) absorbers. The distributions have a break at $N \sim 10^{16}$ cm $^{-2}$, corresponding to overdensities of $\sim 10^2$, likely caused by the transition from sheet/filament to halo gas. Absorption systems with $N \gtrsim 10^{16}$ cm $^{-2}$ are dominated by collisionally ionized O VII and O VIII, while the ionization state of oxygen at lower column densities is also influenced by photoionization. At these high column densities, O VII and O VIII arising in the same structures probe systematically different gas temperatures, meaning their line ratio does not translate into a simple estimate of temperature. While O VII and O VIII column densities and covering fractions correlate poorly with the H I column density at $N_{\text{HI}} \gtrsim 10^{15}$ cm $^{-2}$, O VII and O VIII column densities are higher in this regime than at the more common, lower H I column densities. The column densities of O VI and especially Ne VIII, which have strong absorption lines in the UV, are good predictors of the strengths of O VII and O VIII absorption and can hence aid in the detection of the X-ray lines.

2.1 Introduction

Within extragalactic astronomy, the missing baryon problem is well-established. We know from cosmic microwave background measurements (e.g. Planck Collaboration et al. 2014) and big bang nucleosynthesis (see the review by Cyburt et al. 2016) how many baryons there were in the very early universe. However, at low redshift, we cannot account for all of these baryons by adding up observed populations: stars, gas observed in the interstellar, circumgalactic, and intra-cluster media, and photoionized absorbers in the intergalactic medium (IGM). According to a census presented by Shull et al. (2012), about 30% of baryons are observationally missing at low redshift.

Cosmological simulations predict that missing baryons reside in shock-heated regions of the IGM. Cen & Ostriker (1999) already found that the missing gas is mainly hot and diffuse. They called this gas the warm-hot intergalactic medium (WHIM), defined as gas with temperatures in the range 10^5 – 10^7 K. Some of this gas has already been detected. The cooler component is traced largely by O VI absorption: $T \sim 10^{5.5}$ K in collisional ionization equilibrium (CIE), which applies to high-density gas. O VI absorption has been studied observationally by many groups (e.g., Tumlinson et al. 2011; Johnson et al. 2013), often using the Cosmic Origins Spectrograph (COS) on the Hubble Space Telescope (HST). Others have investigated O VI absorption in galaxy and cosmological simulations (e.g., Cen & Fang 2006; Tepper-García et al. 2011; Cen 2012; Shull et al. 2012; Rahmati et al. 2016; Oppenheimer et al. 2016, 2018; Nelson et al. 2018). Predictions for the hotter part of the WHIM have been made with simulations (e.g., Cen & Fang 2006; Branchini et al. 2009; Bertone et al. 2010a; Cen 2012; Nelson et al. 2018), typically, but not exclusively, focussing on the O VII and O VIII X-ray lines ($T \sim 10^{5.4}$ – $10^{6.5}$ and $\sim 10^{6.1}$ – $10^{6.8}$ K in CIE, respectively). Others have made predictions for the WHIM gas using analytical models (Perna & Loeb 1998), or using a combination of analytical models and simulations (Fang et al. 2002a; Furlanetto et al. 2005).

Besides being a major baryon reservoir, the WHIM also provides an important way to understand accretion and feedback processes in galaxy formation. As gas collapses onto a galaxy, some of it forms stars or is accreted by the central supermassive black hole. This creates feedback, where these stars (through, for example, supernova explosions) and active galactic nuclei (AGN) inject energy and momentum into their surrounding gas, slowing, stopping, or preventing further gas accretion. Current-generation cosmological simulations like EAGLE (Schaye et al. 2015), IllustrisTNG (Pillepich et al. 2018), and Horizon-AGN (Dubois et al. 2014) include star formation and AGN feedback, and metal enrichment from stars. The feedback in EAGLE and IllustrisTNG, and the AGN feedback in Horizon-AGN, is calibrated to reproduce galaxy (stellar and black hole) properties, because individual stars, supernova explosions, and black hole accretion disks are too small to be resolved in these simulations. That means that these simulations require a model of the effect of this feedback on scales they can resolve.

The effect of this feedback on the circumgalactic medium (CGM) is not as well-constrained. For example, feedback in BAHAMAS (McCarthy et al. 2017) and IllustrisTNG is calibrated to halo gas fractions, but this only applies to high-mass haloes where observations are available ($M_{500c} > 10^{13} M_{\odot}$). One way to constrain the effects of feedback on the CGM, is to study how far metals, which are created in galaxies, spread outside their haloes. In other words, we can investigate what fraction of the metals ends up in the intergalactic medium, which impacts how many metal absorbers we expect to find in quasar sightlines. This is not the only possible effect, as feedback also impacts the temperature, density, and kinematics of gas.

The main way we expect to find the hot WHIM gas, is through absorption in e.g. quasar spectra (e.g., [Brenneman et al. 2016](#), [Nicastró et al. 2018](#), [Kovács et al. 2019](#)). This is because this WHIM gas typically has low densities: since emission scales with the density squared and absorption with the density along the line of sight, this makes absorption more readily detectable than emission in most of the WHIM. Observationally, absorption from these ions has been found around the Milky Way, as described by e.g. [Bregman \(2007\)](#), but claims of extragalactic O VII and O VIII are rare (e.g., [Nicastró et al. 2005](#)) and often disputed (e.g., [Kaastra et al. 2006](#)) or uncertain (e.g., [Bonamente et al. 2016](#)). [Nicastró et al. \(2017\)](#) review these WHIM searches in absorption. [Nicastró et al. \(2018\)](#) recently found two extragalactic O VII absorbers, using very long observations with the XMM-Newton RGS of the spectrum of the brightest X-ray blazar. These are consistent with current predictions, though the small number of absorbers means uncertainties on the total absorber budget are still large.

Besides these oxygen ions, other ions are also useful for studying the WHIM. For example, Ne VIII traces WHIM gas that is somewhat hotter than O VI traces ($T \sim 10^{5.8}$ K in CIE). Theoretically, [Tepper-García et al. \(2013\)](#) investigated its absorption in the OWLS cosmological, hydrodynamical simulations ([Schaye et al. 2010a](#)) and [Rahmati et al. \(2016\)](#) did so with EAGLE. Observationally, e.g., [Meiring et al. \(2013\)](#) studied Ne VIII absorption, and recently, [Burchett et al. \(2019\)](#) found Ne VIII absorption associated with the circumgalactic medium (CGM). A different tracer of the WHIM is broad Ly α absorption ($T \gtrsim 10^5$ K), which traces the cooler component of the WHIM, similar to O VI but without the need for metals (e.g., [Tepper-García et al. 2012](#), [Shull et al. 2012](#)). These ions are useful because they trace the WHIM themselves, but they can also be used to find O VII and O VIII absorbers. Recently, [Kovács et al. \(2019\)](#) did this: they found extragalactic O VII absorption by stacking quasar spectra using the known redshifts of 17 Ly α absorbers associated with massive galaxies.

[Mroczkowski et al. \(2019\)](#) give an overview of how the Sunyaev-Zeldovich (SZ) effect has been used to search for the WHIM. This effect probes line-of-sight integrated pressure (thermal SZ) or free electron bulk motion (kinetic SZ), which both have the same density dependence as absorption and are therefore less biased towards high-density gas than emission is. Attempts to detect the IGM by this method have focussed on SZ measurements of cluster pairs to detect filaments between them, and cross-correlations with other tracers of large-scale structure ([Mroczkowski et al. 2019](#) and references therein). For example, [de Graaff et al. \(2019\)](#) used stacked massive galaxy pairs (mean stellar mass $10^{11.3} M_{\odot}$, virial mass $\sim 10^{13} M_{\odot} h^{-1}$) to detect a filamentary thermal SZ signal. [Tanimura et al. \(2019\)](#) similarly stacked luminous red galaxy pairs with stellar masses $> 10^{11.3} M_{\odot}$, and found a filamentary thermal SZ signal larger than, but consistent with, predictions from the BAHAMAS simulations ([McCarthy et al. 2017](#)). Using the EAGLE simulations, [Davies et al. \(2019\)](#) made theoretical predictions for the thermal SZ effect from the CGM at different halo masses.

Another way to find hotter gas is through X-ray emission, but this scales with the density squared, and is therefore generally best for studying dense gas. [Bregman \(2007\)](#) reviews some X-ray observations of extragalactic gas, though this is limited to the denser parts of the intracluster medium (ICM) and CGM. Following earlier work on WHIM X-ray emission in simulations by e.g., [Pierre et al. \(2000\)](#) and [Yoshikawa et al. \(2003\)](#), [Bertone et al. \(2010a\)](#) studied soft X-ray line emission, and found that this emission should indeed be tracing mostly denser and metal-rich gas, i.e., ICM and CGM. [Tumlinson et al. \(2017\)](#) discuss some more recent results on X-ray line emission from the CGM. Instead of lines, [Davies et al. \(2019\)](#) investigated broad-band soft X-ray emission from the CGM in EAGLE, and found that it can be used as a proxy for the CGM gas fraction.

One way to search for the missing baryons in absorption is by doing blind surveys, where

observers look for absorbers in the lines of sight to suitable background sources. There are a number of upcoming and proposed X-ray telescopes for which there are plans to carry out such surveys, including Athena (Barret et al. 2016), Arcus (Brenneman et al. 2016; Smith et al. 2016), and Lynx (The Lynx Team 2018). The O VIII and especially O VII ions tend to be the focus of such plans, since oxygen is a relatively abundant element (Allende Prieto et al. 2001) and these ions have lines with large oscillator strengths (Verner et al. 1996), making these lines more readily detectable in the hot WHIM.

In simulations, O VII and O VIII absorption has been studied: e.g. Branchini et al. (2009) and Cen & Fang (2006) predicted the equivalent width distributions for blind surveys using an earlier generation of simulations and mock spectra generated from these. More recently, Nelson et al. (2018) studied absorption by these ions in the IllustrisTNG simulations, but they did not study equivalent widths.

Here, we will study the column density and equivalent width distributions of O VII (21.60 Å) and O VIII (18.967, 18.973 Å) in the EAGLE simulations (Schaye et al. 2015; Crain et al. 2015), to predict which absorption systems may be detected by future missions, and to establish the physical conditions these absorption systems probe. EAGLE has been used to predict column density distributions that agree reasonably well with observations at a range of redshifts, for a variety of ions (e.g., Schaye et al. 2015; Rahmati et al. 2016). However, these studies all focussed on ions with ionization energies lower than O VII (739 eV) and O VIII (871 eV) (Lide 2003).

This paper is organised as follows. In Section 4.2 we discuss our methods: the EAGLE simulations themselves (Section 4.2.1), and how we extract column densities (Section 2.2.2) and equivalent widths (Section 2.2.3) from them. We present the column density distributions we find in Section 2.3.1 and discuss our mock spectra (Section 2.3.2) and the equivalent width distributions we infer from these (Section 2.3.3). We discuss the origin of the shape of the column density distribution (Section 2.3.4) and how it probes AGN feedback (Section 2.3.5). We discuss the physical properties of O VII and O VIII absorption systems in Section 2.3.6. In Section 2.3.7 we discuss how absorption by these two ions correlates with three ions with UV lines: H I, O VI, and Ne VIII. In Section 2.3.8, we compare the gas traced by different ions along the same sightlines. In Section 4.5, we outline what our results predict for three planned or proposed missions in Section 2.4.1, and what some limitations of our work may be (Section 2.4.2). Finally, we summarise our results in Section 2.5.

2.2 Methods

We study predictions for O VII and O VIII absorption in the EAGLE simulations (Schaye et al. 2015; Crain et al. 2015). We use tabulated ion fractions as a function of temperature, density, and redshift, as well as the element abundances from the simulation, to calculate the number density of ions at different positions in the simulation. By projecting the number of ions in thick slices through these simulations onto 2-dimensional column density maps, we obtain column density distributions from the simulations. To calculate equivalent widths for some of these columns, we generated synthetic absorption spectra at sightlines through their centres.

Oxygen abundances in this paper are given in solar units. Here, the solar oxygen mass fraction is 0.00586 (Allende Prieto et al. 2001). This is simply used as a unit, and should not be updated to more recent solar abundance measurements in further work. Length units include p (‘proper’) or c (‘comoving’). The exception is the centimetre we use in (column) densities: centimetres are always proper units.

2.2.1 The EAGLE simulations

In this section, we provide a short summary of the EAGLE (Evolution and Assembly of GaLaxies and their Environments) simulations. More details can be found in Schaye et al. (2015), the paper presenting the simulations calibrated to observations, Crain et al. (2015), which describes the calibration of these simulations, and McAlpine et al. (2016), which describes the data release of the EAGLE galaxy and halo data.

The code used is a modified version of Gadget3, last described by Springel (2005), using the TREE-PM gravitational force calculation. The modifications include an implementation of smoothed particle hydrodynamics (SPH) known as ANARCHY (Schaye et al. 2015 appendix A; Schaller et al. 2015). A Λ CDM cosmogony is used, with the parameters $(\Omega_m, \Omega_\Lambda, \Omega_b, h, \sigma_8, n_s, Y) = (0.307, 0.693, 0.04825, 0.6777, 0.8288, 0.9611, 0.248)$ (Planck Collaboration et al. 2014).

Gas cooling is implemented following Wiersma et al. (2009a), using the tracked abundances of 11 elements and cooling rates for each. Collisional and photoionization equilibrium is assumed, with photoionization from the Haardt & Madau (2001) UV/X-ray background model. With those assumptions, and the element abundances as tracked in the simulation, we calculate the number of ions in a column through the simulated box. To do this, we use ion fraction tables from Bertone et al. (2010a,b), who investigated line emission with tables computed under the same assumptions as the EAGLE gas cooling. The ion fractions were computed using Cloudy, version c07.02.00¹ (Ferland et al. 1998). These tables give the ion fraction (fraction of nuclei of a given element that are part of a particular ion) as a function of \log_{10} hydrogen number density, \log_{10} temperature, and redshift. We interpolated these tables linearly, in log space for temperature and density, to obtain ion balances for each SPH particle.

These assumptions mean we ignore both local ionization sources, such as stars and AGN, and non-equilibrium ionization. This also means we ignore the effect of flickering AGN. These could boost the abundances of highly ionized species, such as those we are interested in, even if the AGN is not ‘active’ when we observe it (Oppenheimer & Schaye 2013; Segers et al. 2017; Oppenheimer et al. 2018). This boost is caused by species being ionized by the AGN radiation, and then not having time to recombine before observations are made. For typical AGN duty cycles and ion recombination times, the ions can be out of ionization equilibrium like this for a large fraction of systems. We discuss the impact of these assumptions in more detail in Section 2.4.2.

Since the resolution of the simulations is too low to resolve individual events of star formation and feedback (stellar winds and supernovae), or accretion disks in AGN, the star formation and stellar and AGN feedback on simulated scales are implemented using subgrid models, with model feedback parameters calibrated to reproduce the $z = 0.1$ galaxy stellar mass function, the relation between black hole mass and galaxy mass, and reasonable galaxy disc sizes. Star formation occurs where the local gas density is high enough, with an additional metallicity dependence following Schaye (2004). The rate itself depends on the local pressure, in a way that reproduces the Kennicutt-Schmidt relation (Schaye & Dalla Vecchia 2008). Stars lose mass to surrounding gas particles as they evolve, enriching them with metals. This is modelled according to Wiersma et al. (2009b).

A major problem in galaxy formation simulations at EAGLE-like resolution is that if reasonable amounts of stellar and AGN feedback energy are injected into gas surrounding a single-age population of stars or a black hole at each time step, the energy is radiated away

¹This older version is consistent with the EAGLE cooling rate calculations.

Table 2.1: The simulations used in this work. The names consist of three parts, in the format <name>-L<size>N<particles>. The name is the name or abbreviation for the stellar and AGN feedback model, as explained in the text. The size is the size of the simulation box in comoving Mpc, and the last part is the cube root of the number of dark matter particles (equal to the initial number of gas particles) used in the simulation. The table lists the dark matter particle mass (m_{DM}), the initial gas particle mass $m_{\text{gas, init}}$, and the Plummer-equivalent gravitational softening length at low redshift (l_{soft}).

simulation name	m_{DM} (M_{\odot})	$m_{\text{gas, init}}$ (M_{\odot})	l_{soft} (pkpc)
Ref-L100N1504	9.70×10^6	1.81×10^6	0.70
Ref-L050N0752	9.70×10^6	1.81×10^6	0.70
Ref-L025N0376	9.70×10^6	1.81×10^6	0.70
Ref-L025N0752	1.21×10^6	2.26×10^5	0.35
Recal-L025N0752	1.21×10^6	2.26×10^5	0.35
NoAGN-L050N0752	9.70×10^6	1.81×10^6	0.70

before it can do any work. This causes gas in galaxies to form too many stars. Following Booth & Schaye (2009) and Dalla Vecchia & Schaye (2012), the solution for this problem used in EAGLE is to statistically ‘save up’ the energy released by these particles, until it is enough to heat neighbouring particles by a fixed temperature increment of $10^{7.5}$ K (stellar feedback) or $10^{8.5}$ K (AGN feedback). Which particles are heated and when is determined stochastically. The expectation value depends on the local gas density and metallicity (Crain et al. 2015).

Table 2.1 gives an overview of the different EAGLE simulations we use in this work. The reference feedback model was calibrated for the standard EAGLE resolution, as used in e.g., Ref-L100N1504. Simulation Recal-L025N0752 was calibrated in the same way as the reference model, but at an eight times higher mass resolution. This is used to test the ‘weak convergence’ of the simulations, in the language of Schaye et al. (2015), compared to strong convergence tests, which use the same feedback parameters at different resolutions (appendix 2.A). The idea behind this is that the parameters for subgrid feedback will generally depend on which scale is considered subgrid, so the subgrid model parameters are expected to be resolution-dependent. Finally, NoAGN-L050N0752 is a variation of Ref-L050N0752 without AGN feedback, which was not described in Schaye et al. (2015) or Crain et al. (2015). Except for appendix 2.A where we test resolution and box size convergence, we will only use the Ref-L100N1504 (reference), Ref-L050N0752 (50 cMpc reference) and NoAGN-L050N0752 (no AGN) simulations.

2.2.2 Column density calculation

To obtain column densities from the EAGLE simulations, we calculate the number of ions in columns (elongated rectangular boxes) of finite area and fixed length. We ‘slice’ the simulation along the Z -axis, then divide each slice along the X - and Y -directions into narrow 3-dimensional columns, which, when projected, become the pixels of a (2-dimensional) column density map. We make such a column density map for each slice along the Z -direction. We use 32,000 columns along both 100 cMpc sides of the simulation box, meaning each column has an area of 3.125^2 ckpc² (comoving kpc) and each map has 1.024×10^9 pixels. We use 16 slices, each 6.25 cMpc thick. In appendix 2.A we verify that the column density statistics are converged at this pixel size, and examine the effect of slice thickness. These columns are an approximation of what is done observationally, where absorbers are defined by regions of

statistically significant absorption in spectra of nearly point-like sources (typically quasars), and column densities are obtained by fitting Voigt profiles to these regions.

To project the ions of each SPH particle onto a grid, we need to know the shape of the gas distribution that individual particles model. There is no unique function for this, since anything smaller than a single SPH particle is, by definition, unresolved. However, there is a sensible choice. In SPH, a similar assumption about the gas distribution must be made, which is used to evolve the gas particles' motions and thermodynamic properties. The function describing this is known as a kernel. We use this same function in projection: the C2-kernel² (Wendland 1995). Along the Z -axis, we simply place each particle in the slice that contains its centre. We have verified that the results are insensitive to the chosen kernel: the difference in the column density distribution function (CDDF, the main statistic we are interested in) compared to using a different kernel (the Gadget kernel) is $\lesssim 0.05$ dex for $10^{11} \text{ cm}^{-2} < N_{\text{O VII, VIII}} < 10^{16.5} \text{ cm}^{-2}$, which covers the column densities we are interested in. We discuss the kernels in more details in Appendix 2.A

An issue that arises in EAGLE is that for cold, dense gas (star-forming gas), the temperature is limited by a fixed equation of state, used to prevent artificial fragmentation during the simulation (Schaye & Dalla Vecchia 2008). This means that the temperature of this gas does not represent the thermodynamic state we would actually expect the gas to be in, and ion balance calculations using this temperature are unreliable. Following e.g. Rahmati et al. (2016), we therefore fix the temperature of all star-forming gas to 10^4 K, typical for the warm phase of the ISM. However, we found that the treatment of this gas has virtually no impact, as O VII and O VIII are high-energy ions, and therefore mainly exist in hot and/or low-density gas (see Section 2.3.6). The difference between CDDFs calculated using our standard method, using the temperatures in the simulation output, and excluding star-forming gas altogether is < 0.01 dex at $10^{11} \text{ cm}^{-2} < N < 10^{16.5} \text{ cm}^{-2}$.

The column density distribution function (CDDF) is defined as

$$f(N, z) \equiv \frac{\partial^2 n}{\partial N \partial X}, \quad (2.1)$$

where X is the dimensionless absorption length, n is the number of absorption systems, and N is the ion column density. Here,

$$dX = dz [H_0/H(z)] (1+z)^2, \quad (2.2)$$

which means that at $z = 0$, $dX = dz$ (Bahcall & Peebles 1969). At higher redshifts, using dX accounts for changes in the CDDF due to the uniform expansion of the Universe for absorption systems of fixed proper size. We calculate dz from the thickness of the spatial region we project (i.e., the length along the projection axis), through the Hubble flow, using the redshift of the snapshot and the cosmological parameters used in the simulation.

2.2.3 Spectra and equivalent widths

We calculated the equivalent widths along EAGLE sightlines using mock spectra obtained using a program called SPECWIZARD. It was developed by Joop Schaye, Craig M. Booth and Tom Theuns, and is described in Tepper-García et al. (2011, section 3.1). A given line of sight is divided into (1-dimensional) pixels in position space. First, the number of ions is calculated for each SPH particle intersecting the line of sight. Then the column density in

²In the simulation, the function depends on the 3D distance to the particle position; for the projection, we input the 2D distance (impact parameter) instead.

each pixel is calculated by integrating the particles' assumed ion distribution (defined by the SPH kernel) over the extent of the pixel. This guarantees that the column density is correct at any resolution. For the ion distribution, we assume a 3-dimensional Gaussian, since this is easy to integrate. We have verified that the difference with column densities and equivalent widths obtained with a different kernel is negligible. The ion-number-weighted temperature and peculiar velocity along the line of sight are also calculated in each pixel.

Once this real space column density spectrum has been calculated, it is used to obtain the absorption spectrum in velocity space. For each pixel in the column density spectrum, the optical depth distribution in velocity space is calculated as for a single absorption line. The absorption is centred based on the position of the pixel and the Hubble flow at the simulation output redshift, together with the peculiar velocity of the pixel. The thermal line width (b parameter) is calculated from the temperature. We have not modelled any sub-grid/unresolved turbulence, and neglect Lorentz broadening in our calculations. We also use these 'ideal' spectra directly, and do not model continuum estimation, noise, line detection, or a detecting instrument in our equivalent width calculations.

We calculate rest frame equivalent widths, EW , from these mock spectra by integrating the entire spectrum:

$$EW = \left(1 - \sum_{i=0}^{N-1} \frac{F_i}{N} \right) \lambda_{\text{rest}} \frac{H(z) l_{\text{sl,com}}}{c(1+z)}, \quad (2.3)$$

where N is the number of pixels, F_i is the flux in pixel i , normalised to the continuum, λ_{rest} is the rest-frame wavelength of the absorption line, z is the redshift, $H(z)$ is the local Hubble flow, c is the speed of light, and $l_{\text{sl,com}}$ is the comoving box size. The (observer-frame) velocity difference across the full 100 cMpc box is 7113 km s^{-1} ($\Delta z = 0.02373$) at redshift 0.1, and 12026 km s^{-1} ($\Delta z = 0.04011$) at redshift 1 with the EAGLE cosmological parameters. The spectra we generate inherit the periodic boundary conditions of the simulation box, and therefore probe velocity differences of at most half these values.

We use Verner et al. (1996) oscillator strengths and wavelengths for the 18.97 \AA O VIII doublet: $f_{\text{osc}} = 0.277, 0.139$ and $\lambda = 18.9671, 18.9725 \text{ \AA}$, respectively. For the O VII resonant line, we use values consistent with theirs: $\lambda = 21.6019 \text{ \AA}$ and $f_{\text{osc}} = 0.696$, but ours come from a data compilation by Kaastra (2018). The other O VII He-like lines (forbidden and intercombination) have wavelengths sufficiently far from the resonant line (e.g., Bertone et al. 2010a) that these should be clearly separated from each other at the resolutions achieved by instruments aboard Arcus (Smith et al. 2016; Brenneman et al. 2016), Athena (Lumb et al. 2017; Barret et al. 2018), and Lynx (The Lynx Team 2018) (see Sections 2.3.2 and 2.4.1 for further discussion of the telescopes and instruments). Therefore, we only discuss the resonant line here, and note that for lines at the detection limits of these instruments, the lines are not blended. Also, these lines have much weaker oscillator strengths, by a factor of at least ~ 6000 , so although they can be important in emission studies, the forbidden and intercombination lines will not be important in absorption.

Since we compute the equivalent width by integrating the entire spectrum, if there is more than one absorption system along the line of sight, we will only recover the total equivalent width. However, the equivalent widths we are interested in are the potentially detectable ones, i.e., the rarer, larger values, corresponding to the highest column densities. These will not generally be 'hidden' in projection by even larger EW absorption systems, though we find that multiple systems of similar strength along such lines of sight are not uncommon. We will discuss this issue further in Section 2.3.3, where we examine some spectra and the conversion between column density and equivalent width.

To determine the relation between the (projected) column density and equivalent width, we obtained mock spectra along 16384 lines of sight through the reference simulation (Ref-L100N1504) output at redshift 0.1.

Since we are mainly interested in the systems with large equivalent widths, we did not choose random lines of sight. Instead, we selected sightlines using the column density maps for O VI, O VII, and O VIII: for each ion, we selected samples randomly from evenly spaced log column density bins. We used column densities over the full 100 cMpc box depth, along the Z-direction, for this selection.

The matching of equivalent widths to projected column densities is the reason we calculate the equivalent widths along 100 cMpc sightlines, despite the possibility of combining two absorption systems this way. The peculiar velocities of the absorbers are often ~ 500 km/s, and can reach 1000 km/s or more. Compared to a Hubble flow across the redshift 0.1 box of ~ 7000 km/s, this means that matching absorption in the spectrum to one of 16 slices along the line of sight would not be a straightforward process. A mistake here could have major implications for the reconstructed equivalent width distribution, since the highest column density absorption systems are orders of magnitude more rare than more typical ones. This means that mismatches of absorption systems with large equivalent widths to lower column densities could lead to a significant overestimation of the abundance of these rare absorption systems. For 100 cMpc, there is no possibility of mismatches along the line of sight.

With this sample of column densities and equivalent widths, we make a 2-dimensional histogram of column density and equivalent width. By normalising it by the number of absorbers at each column density, we obtain a matrix from this histogram that we use to convert column density distributions to equivalent width distributions. Though we make these matrices for 100 cMpc sightlines, we will also apply them to column density distributions obtained using 6.25 cMpc slices of the simulation. That is to say, we assume that the relation between column density and equivalent width does not depend significantly on the slice thickness (sightline length).

2.3 Results

2.3.1 Column density distributions

The starting points for our column density distributions are column density maps. Fig. 2.1 shows a map of column densities calculated as we described in Section 2.2.2. These column densities are for columns through the full depth of the 100 cMpc box, using 400^2 pixels, and therefore do not produce the converged CDFs we use in the rest of the paper. They do demonstrate that O VII and O VIII trace the large-scale structure in the box, as indicated by the total gas surface density map. We will investigate the spatial distribution of these ions (column densities around haloes of different masses) in more detail in an upcoming paper.

Fig. 2.2 shows the O VII and O VIII CDFs (solid lines) we obtained from the reference simulation (Ref-L100N1504). The main feature visible here is the ‘knee’ or break in the distributions for both ions at column densities around 10^{16} cm $^{-2}$. We will investigate this feature in detail in Section 2.3.4 where we find that it marks the transition from absorption arising in extrahalo and intrahalo gas. We show in appendix 2.A that these distributions are converged with pixel size in the column density range we shown here (up to $10^{16.5}$ cm $^{-2}$), and reasonably converged with the size and resolution of the simulations. We have verified that the effects of some technical choices in the calculation of the column densities are small.

The dotted lines indicate the distribution we would get if all gas had an oxygen abundance




Figure 2.1: A map of the O VII (middle) and O VIII (bottom) column densities in the 100 cMpc EAGLE Ref-L100N1504 simulation at $z = 0.1$, at a resolution of 400^2 pixels and with a column depth of 100 cMpc. The full column density range is not shown. The top panel shows the corresponding gas surface overdensities. O VII and O VIII trace large-scale structures.

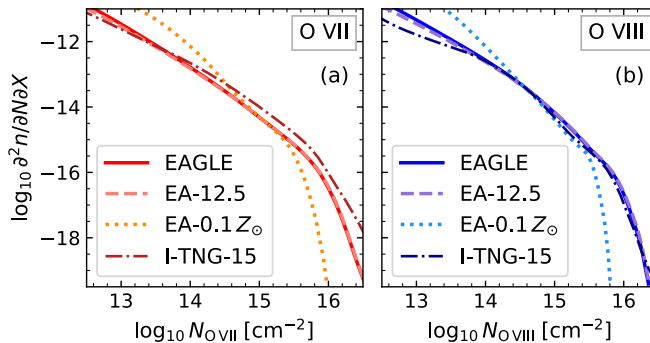


Figure 2.2: The column density distribution function (CDDF), as described in equation 2.1, of O VII (a) and O VIII (b) in the 100 cMpc EAGLE reference simulation at $z = 0.0$. The solid lines (EAGLE) show our standard CDDF: it uses the simulation oxygen abundances, and column densities are measured in 6.25 cMpc long columns. The dotted lines (EA-0.1 Z_{\odot}) use 0.1 times the solar oxygen abundance for all gas instead. All the EAGLE column densities are calculated from 32000^2 columns in each slice along the line of sight direction; each column has an area of 3.125^2 cMpc 2 . We also compare to the IllustrisTNG 100-1 CDDF (Nelson et al. 2018) I-TNG-15), shown with dot-dashed lines. Since these column densities were measured in 15 cMpc columns, we also show the CDDF we obtain using 12.5 cMpc columns for comparison (dashed lines, EA-12.5). This shows that EAGLE and IllustrisTNG-100 predict similar column density distributions for these ions, and that using realistic metallicities is crucial in determining the column density distribution.

0.1 times the solar value. This demonstrates the impact of non-uniform metal enrichment: it causes the CDDF to be much shallower, extending to larger maximum column densities. We have verified that hot or high-density absorbers are most enriched with metals, while the less dense IGM (especially the cooler part) is less enriched. This means the O VII and O VIII CDDFs are sensitive to the distribution of oxygen.

From a different perspective, 40 % of gas-phase oxygen and 27 % of total oxygen³ is in these two ions at $z = 0.1$ in the reference simulation (Ref-L100N1504). That means that absorption from these ions is also important in determining where the bulk of the metals produced in galaxies go.

We also compare our O VII and O VIII CDDFs to those of Nelson et al. (2018). They obtained column density distributions for these ions from the IllustrisTNG 100-1 simulation, using a similar method to ours. Nelson et al. (2018) use metallicity-dependent ion balances to calculate their column densities, but remark that the dependence of ion fractions on metallicity is minimal. They use 15 cMpc long columns to calculate their column densities, but we find that this makes almost no difference for the CDDF in the column density range shown here. They also use a different UV/X-ray background than adopted here.

The IllustrisTNG 100-1 simulation uses a similar cosmology, volume, and resolution as EAGLE, but a different hydrodynamics solver (moving mesh instead of SPH). While it also includes star formation and stellar and AGN feedback, it models these processes using different subgrid prescriptions.

³The total amount of oxygen here is the amount of oxygen released by stars, and does not count oxygen that remains in stellar remnants (a significant quantity) or is swallowed by black holes (a relatively insignificant quantity). Compared to gas-phase oxygen, the total also includes oxygen that was in gas, but is now in stars that formed from that gas.




Figure 2.3: The evolution of the column density distribution function (CDDF) of O VII (a, c) and O VIII (b, d) in the 100 cMpc EAGLE reference simulation from $z = 1$ to 0. The top panels show the CDDFs themselves, while the bottom panels show the CDDFs relative to the $z = 0$ CDDF. This shows that the CDDF evolves only mildly between redshifts 0 and 1.

Fig. 2.2 shows that the CDDFs from the EAGLE reference simulation (Ref-L100N1504) and IllustrisTNG-100-1 agree remarkably well; the differences are small compared to the effect of assuming a constant metallicity. Note that the EAGLE and IllustrisTNG-100-1 absorber numbers differ by a factor ≈ 2 at column densities above the break for O VII, though the differences are small compared to the dynamical range shown. A comparison at fixed metallicity (not shown) did not decrease the differences, meaning they are not dominated by different metal distributions in the two simulations.

Fig. 2.3 shows that the CDDF evolves little between redshifts 0 and 1. The shape of the CDDF does not change much, but there is some evolution: the incidence of high column densities decreases with time, while that of low column densities increases. The distribution evolves the least around its break. We have verified that these changes are not simply the result of the changing pixel area and slice thickness in physical units at the fixed comoving sizes we used, by comparing these changes to the effects of changing column dimensions described in appendix 2.A. The evolution we find in EAGLE is similar to that shown in Fig. 5 of Nelson et al. (2018).

2.3.2 Absorption spectra

We move on to an examination of a few of the spectra we will obtain equivalent widths (EWs) from. Three example spectra are shown in Fig. 2.4. We show our ideal spectra (resolution of 2 km s^{-1} , no noise), and mock spectra with spectral resolutions and effective areas corre-




Figure 2.4: Example mock spectra from the EAGLE reference simulation (Ref-L100N1504) at $z = 0.1$. On the x-axis, we show rest-frame velocity, on the y-axis, we show the flux, normalised to the continuum. The (a), (e), and (i) panels show the ideal spectra the other panels are derived from. These have a resolution of 2 km s^{-1} and are noise-free. In the other panels, we show mock spectra for the Athena X-IFU (b, f, j), Arcus (c, g, k), and Lynx XGS (d, h, l) effective areas and resolutions. These spectra have Poisson noise added to them, but no other error or uncertainty sources. We show the spectra for the O VII resonant line and the O VIII doublet. In the ideal spectra, we also show the spectrum using only the stronger of the two O VIII doublet lines to gauge the effect of the weaker one. The three panel groups (a, b, c, d), (e, f, g, h), and (i, j, k, l) represent three different sightlines, selected to illustrate some general trends. We chose sightlines with O VII column densities of $10^{16.5}$, $10^{16.0}$, and $10^{15.5} \text{ cm}^{-2}$. The column densities and EWs measured over the whole sightline, for the ideal spectra, are shown to the right of the spectra. For Athena X-IFU mock spectra, we use a (full-width half-maximum) resolution of 2.1 eV and an effective area (specified for soft X-rays) of 1.05 m^2 (Barret et al. 2018; Lumb et al. 2017). For Arcus, we used a spectral resolving power $E/\Delta E$ (where ΔE is the full-width half-maximum and E is the observed energy) of 2000 below 21.60 Å (observed wavelength), and 2500 at longer wavelengths, with an effective area (average at 16–28 Å) of 250 cm^2 (Smith et al. 2016). We used $E/\Delta E = 5000$ and an effective area (specified for 0.3–0.7 keV) of 4000 cm^2 for the Lynx XGS (The Lynx Team 2018). The background (blazar) source flux was $1 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ between 2 and 10 keV with a photon spectral index of $\Gamma = 1.8$ as specified for the Athena 5σ weak line sensitivity limit by Lumb et al. (2017), and an observing time of 100 ks. The spectra, like the simulations they come from, are periodic. We bin the instrument mock spectra to two pixels per full-width half-maximum.

sponding to the Athena X-IFU, Arcus⁴ and Lynx XGS instruments for a 100 ks observation of a blazar sightline. For the blazar, we use a source flux of 1×10^{-11} erg cm⁻²s⁻¹ between 2 and 10 keV with a photon spectral index of $\Gamma = 1.8$, as specified for the Athena 5σ weak line sensitivity limit by Lumb et al. (2017). We model the Poisson noise, but no other sources of error/uncertainty, and determine the unabsorbed number of photons per bin from the blazar spectrum and the redshifted line energy. We do not account for the slope of the blazar spectrum across our sightline. The spectra are periodic, like the simulations we derived them from, which is why we see absorption at low line of sight velocities in panel (f). We selected these spectra to be roughly representative of a larger sample we examined: ten spectra at column densities of 10^{14} – 10^{17} cm⁻², in 0.5 dex intervals, for each ion.

The ideal spectra illustrate that most absorption systems have more than one component. Single-component systems do occur, though. Athena is more sensitive than Arcus, but it cannot distinguish close components as easily. Line widths, in Athena X-IFU and even in Arcus spectra, may not provide very stringent limits on absorber temperatures, since these widths are inferred mostly from the component structure. With Lynx, it may be more feasible to separate the different components.

We note that it is not uncommon to find multiple absorption systems along a single line of sight. This should not be a major problem for our 6.25 cMpc CDDFs, but it could affect the relation between column density and EW, since the column densities and equivalent widths of the different systems simply add up, even when the stronger system’s absorption is saturated. However, we show in appendix 2.B that the EW distributions obtained using different conversions between column density and equivalent width are quite similar around the expected detection thresholds of Athena, Arcus, and Lynx, so we do not expect this to have a significant impact on our survey predictions.

We see clear differences between the spectra obtained with the simulated instruments: Arcus has much higher resolution than Athena, enabling clearer measurements of line shapes and detections of different components, while Athena’s much higher effective area enables it to detect weaker lines in the same observing time. Note that the science requirements for these two instruments specify different observing times for the proposed blind WHIM surveys. The Lynx XGS is planned to have the highest spectral resolution of the three and a large effective area, meaning it recovers these lines best. We will discuss these instruments in more detail in Section 2.4.1

Finally, the ideal spectra confirm that we need to model the O VIII doublet as two blended lines. They are intrinsically blended, but the offset between the lines is large enough that it might affect the line saturation (panel a), so modelling the two lines as one may also be inadequate. The EW of the detected line will be determined by both doublet components together.

2.3.3 Equivalent widths

As described in Section 2.2.3, we calculated the rest-frame equivalent width (EW) distribution in the reference simulation (Ref-L100N1504) by first obtaining the relation between the column density from the projection and the total EW along the same line of sight for a subset of sightlines with relatively high column densities. This relation includes the scatter in EW at fixed column density. We then used this relation to obtain the EW distribution from the column density distribution. This approximation is necessary because computing absorption spectra corresponding to the full 32000^2 grid of pixels is too computationally expensive. We

⁴The effective collecting area for Arcus is 500 cm², but the spectrometer only uses half of this to improve spectral resolution (subaperturing).




Figure 2.5: The curve of growth for the O VII resonant line (a) and the O VIII doublet (b) in the reference (Ref-L100N1504) simulation at redshift 0.1. We plot the rest frame equivalent width calculated over the full 100 cMpc sightlines against the column density along those same sightlines. The solid cyan lines are for the best-fit b parameters to the log equivalent width, using only the sightlines selected for the ion in question. The solid gray lines indicate the linear column density equivalent width relation, which applies to optically thin gas. The dotted lines roughly indicate the upper and lower range of b parameters, with the brown line indicating the bare minimum for O VII while the green lines indicate values below which absorbers are rare. This demonstrates our sampling of the column density equivalent width relation, and the range of single-absorber-equivalent widths associated with it.

did not fit any models to this relation, but simply binned the distribution of EWs as a function of column density, and used it as a conversion matrix. Below the minimum column density of 10^{13} cm^{-2} for which we selected sightlines, we used the linear curve of growth to make the conversion. We discuss the calculation of the EW distributions and the effect of including scatter in more detail in Appendix [2.B](#)

Fig. [2.5](#) shows the rest-frame equivalent width (EW) as obtained by integrating the absorption spectra, as a function of the column density computed with the same code (SPECWIZARD) used to generate the absorption spectra⁵. This shows that for the highest

⁵Here, for each ion, we only show the sightlines selected uniformly in column density for that particular ion, while we use the very similar distribution of all sightlines we have spectra for when converting column densities to EWs. If we plot the EWs against the column densities from our column density maps, there is slightly more scatter in the relation, which is clear at lower column densities. This is due to mismatches between the column densities calculated using the two different methods. These differences are generally small. At the highest column densities,

column densities and EWs ($\gtrsim 10^{15} \text{ cm}^{-2}$ for O VII, $\gtrsim 10^{15.5} \text{ cm}^{-2}$ for O VIII), the relation between the two is no longer linear. Indeed, inspection of the mock spectra (examples in Fig. 2.4) shows the lines at these column densities are typically saturated. Furthermore, the scatter in the relation at the highest column densities is large, meaning a single curve of growth cannot be used to accurately convert between column density and EW. We discuss the effect of scatter on the EW distributions in Appendix 2.B

We characterise this relation using a b -parameter dependent column density-EW relation. This b is a measure of the line width: a model Gaussian absorption line profile has a shape $1 - \exp(-\tau(\Delta v))$, with $\tau(\Delta v) \propto Nb^{-1} \exp(-(\Delta v b^{-1})^2)$. Here Δv is the distance from the line centre in rest-frame velocity units, τ is the optical depth, and N is the column density. The constant of proportionality depends on the atomic physics of the transition producing the absorption. The solid and dotted lines in Fig. 2.5 show curves of growth for different b parameters. There is no clear trend of b parameter with column density.

In Fig. 2.6 we show the EW distributions we obtain from the CDDFs measured over 6.25 cMpc and 100 cMpc sightlines, and compare them to simulation predictions from other groups and a recent measurement. First, comparing our two distributions, the differences are as expected: when we group together multiple absorption systems along a longer sightline, we measure a higher number density of high EWs, while lower-EW absorption systems are ‘swallowed up’ and therefore less common. The differences are larger for O VIII, matching what we see in the CDDF when we vary the slice thickness in appendix 2.A

We also compare our distributions to predictions made by Cen & Fang (2006) and Branchini et al. (2009), and to recent observations by Nicastro et al. (2018), including the update for the equivalent width of the more distant absorber given by Nicastro (2018). Nicastro (2018) measured two O VII resonant line absorption systems at redshifts 0.43 and 0.36 in the spectrum of a very bright blazar, and used that to calculate the EW distribution for such absorbers. These measurements are consistent with all the predictions made so far, in part because the distribution measured from two systems is still quite uncertain⁶. Still, this agreement is encouraging. Fig. 3 of Nicastro et al. (2018) shows a similar agreement. The EAGLE EW distribution shown there was obtained from the column density distribution using a fixed b -parameter, and the 100 cMpc sightline CDDF.

The Cen & Fang (2006) predictions we compare to are based on the Cen & Ostriker (2006) simulations in a 123 cMpc box with 1024^3 cells, and a dark matter particle mass of $3.9 \times 10^8 M_{\odot}/h$. They use somewhat different cosmological parameters than EAGLE. Since these simulations solve the hydrodynamics equations on a fixed Eulerian grid, galaxies and the effects of their feedback are much less well-resolved in these simulations than in EAGLE. The results we show here are from their simulations with galactic feedback (‘GSW’). The ‘G-nL’ curves are derived from simulations tracking ion abundances without assuming ionization equilibrium, while the ‘G-L’ curves were obtained by using equilibrium ion abundances instead. For the photoionization in this model, they use the $z = 0$ radiation field from the simulation, which is consistent with observations. They measured the EW distribution by making mock spectra with a resolution of 19 km s^{-1} (Fang et al. 2002a) for random sightlines. They then identified and counted absorption lines in those. Their predictions are

they are due to non-convergence of the highest column densities with projection resolution (see appendix 2.A). More generally, the spectra and 2d projections assume slightly different gas distributions for a single SPH particle, and the 2d projections deal with SPH particles smaller than the projection column size in ways that the spectra do not. We account for these small differences when we calculate the EW distributions: we use the relation between column densities from our grid projections and EWs calculated at matching sightlines.

⁶The 68% error bars in the figure account for this. Considering the overlapping error bars, we do not interpret these two data points as pointing to a steep slope in the distribution function.




Figure 2.6: Comparison between the O VII (a) and O VIII (b) rest-frame equivalent width distributions in the EAGLE reference simulation (Ref-L100N1504) at $z = 0.1$, the GSW model of Cen & Fang (2006) (CF06) at $z \sim 0$ with (G-L) and without (G-nL) assuming ionization equilibrium, the Branchini et al. (2009) (B+09) B2 model at $z \leq 1$, and observations by Nicastro et al. (2018) with the update of Nicastro (2018) (N+18). Note that the units differ from those in Fig. 2.5: the equivalent widths are in km/s (still rest-frame), and the number of absorption systems per unit redshift z is counted, rather than the number per unit absorption length X . The EAGLE distributions were computed using the 100 cMpc and 6.25 cMpc CDDFs and the 100 cMpc column density equivalent width relation calibrated to all the sightlines in our sample. This shows the rough agreement between the equivalent width distribution we predict for O VII and that predicted in other work, and an agreement with the available observations. The O VIII distributions have larger differences, and there are no extant observational constraints we are aware of.

specified for redshift 0.

In Fig. 2.6, we also compare our EW distributions to the B2 model of Branchini et al. (2009) Fig. 4). In this work, they present three models, of which they consider B2, the model shown here, to be the most realistic. It is also the most optimistic about what we can detect. The B2 model uses simulations by Borgani et al. (2004). This is an SPH simulation using 2×480^3 particles (gas + dark matter), in a $192 \text{ cMpc}/h$ box and a different set of cosmological parameters from Schaye et al. (2015) or Cen & Ostriker (2006). The mass resolution is three orders of magnitude lower than for EAGLE. It includes star formation, supernova feedback, and radiative cooling/heating (for primordial gas). Branchini et al. (2009) impose a density-metallicity relation to get the oxygen abundance for each SPH particle. In model B2, the relation from the simulations of Cen & Ostriker (1999) is imposed, including scatter. They calculate ion abundances with Cloudy, assuming ionization equilibrium. They measure the EW distribution from these sightlines by constructing lightcones, from which they obtain mock spectra at a resolution of 9 km s^{-1} , and then identify lines in these. Branchini et al. (2009) plot observed equivalent widths in their figure, for absorbers at redshifts 0–0.5. We show the rest-frame distribution obtained by assuming all their absorbers are at redshift 0.25.

For O VII, both the Branchini et al. (2009) and the Cen & Fang (2006) models agree rea-




Figure 2.7: The CDDFs for O VII (a) and O VIII (b) absorption systems with ion-weighted overdensities in the ranges indicated in the legend. These distributions are for 6.25 cMpc sightlines in the reference (Ref-L100N1504) simulation at redshift 0. This shows that the break in the distribution occurs at overdensities $\sim 10^2$ for both ions, and therefore likely derives from the transition from absorption tracing sheets and filaments to absorption tracing haloes.

sonably well with our EW distribution, especially at higher column densities. For O VIII, the differences are much larger. The Branchini et al. (2009) B2 model is roughly consistent with EAGLE in this respect, while the latter predicts many fewer absorption systems at high column densities than Cen & Fang (2006).

2.3.4 The physical origin of the break

We now focus on the main feature of the O VII and O VIII column density distributions: the break (‘knee’) at large column densities (just below 10^{16} cm^{-2} , see Fig. 2.2). We look into the temperatures, densities, and metallicities of absorption systems at different column densities to investigate this.

Figures 2.7, 2.8, and 2.9 show what type of gas dominates the CDDF at different column densities. Fig. 2.7 shows the contribution of absorption systems with different ion-weighted overdensities $\delta = \rho / \bar{\rho}_b - 1$. The shape of the total distribution (black) is different to that shown in some other figures because it shows the number of absorption systems per unit \log column density (per unit absorption length), instead of per unit column density (per unit absorption length). We show column densities up to 10^{17} cm^{-2} here, unlike in most figures, to show what happens at the highest column densities. However, since the CDDFs are not converged above $10^{16.5} \text{ cm}^{-2}$, the values here should be taken as indicative.

We see in Fig. 2.7 that higher densities tend to produce higher column densities. At the break in the CDDFs, the ion-weighted overdensities are typically $\sim 10^2$. This points towards a cause for the break: gas at $\delta \gtrsim 10^2$ is typically within dark matter haloes, which have lower covering fractions than partially collapsed intergalactic structures (filaments, sheets, and voids) found at lower overdensities. The denser gas in haloes is also more likely to produce high column densities if it has the right temperature. This picture is broadly confirmed by visual inspection of the column density maps in Fig. 2.1⁷ which show that the highest column

⁷Recall that this figure shows column densities at a much lower resolution than we use for the CDDFs, and measures column densities in 100 cMpc columns.




Figure 2.8: The CDDFs for O VII (a) and O VIII (b) decomposed into absorption systems with ion-weighted temperatures in the ranges indicated in the legend. This shows that WHIM gas dominates the distributions at all column densities likely to be observable in the foreseeable future, and that absorbers with column densities beyond the break are at temperatures associated with collisional ionization.




Figure 2.9: The CDDFs for O VII (a) and O VIII (b) decomposed into absorption systems with different ion-weighted SPH-smoothed oxygen abundances. The values indicated in the legend are ranges in the log oxygen mass fraction, relative to the solar oxygen mass fraction. This shows that most absorption systems we might detect in the foreseeable future should have oxygen-ion-weighted oxygen abundances around the solar value. (Note that mass- and volume-weighted metallicities may well be smaller.)

densities ($\gtrsim 10^{16} \text{ cm}^{-2}$) mostly occur in nodes in the cosmic web. This result is also roughly consistent with the results of e.g., Fang et al. (2002a), who compared O VII and O VIII CDDFs from their simulation to analytical models and found they could explain the CDDF at high column densities as coming from collapsed haloes.

Fig. 2.8 shows the contribution of absorption systems with different ion-weighted temperatures to the O VII and O VIII CDDFs. The black line shows the total CDDF. Gas with temperatures of 10^5 – 10^7 K is canonically referred to as the warm-hot intergalactic medium (WHIM), where e.g. Cen & Ostriker (1999) found most of the ‘missing baryons’ in their simulations. We see that absorbers tracing this gas dominate the CDDF at column densities likely to be observable with planned future missions. For O VII, we see that absorbers around the break and above tend to trace gas somewhat above $\sim 10^6$ K, while for O VIII, absorbers around the break tend to trace gas at $\sim 10^{6.5}$ K, with the temperature continuing to increase somewhat with column density. These temperature ranges are close to the temperatures where each ion attains its peak ion fraction in collisional ionization equilibrium (CIE). As we will show in more detail in Section 2.3.6 the absorption systems with column densities above the break do indeed tend to be primarily collisionally ionized while below the break, lower temperatures dominate and photoionization is important.

Finally, in Fig. 2.9 we turn to oxygen abundances. It shows that over a wide range in column densities, absorbers tend to trace gas with oxygen-ion-weighted metallicities at, or somewhat below, solar, while oxygen abundances become solar or larger at the largest column densities ($N \gtrsim 10^{16.3} \text{ cm}^{-2}$). Oxygen abundances above ≈ 3 times the solar value do occur, but they are rare. It is important to note that what we show here is the mean oxygen abundance weighted by the contribution of each gas element to the oxygen ion column density. Mass-weighted and volume-weighted mean metallicities can be much lower.

2.3.5 The effect of AGN feedback

To investigate the effect of AGN feedback on the CDDF, we compare the column density distributions we found in Section 2.3.1 to a recent EAGLE simulation not described by Schaye et al. (2015) or Crain et al. (2015): NoAGN-L050N0752. As the name suggests, this simulation does not include AGN feedback, while the rest of the (subgrid) physics is the same as in the reference model. This means it does not produce a realistic universe: AGN feedback is needed to quench star formation in high-mass galaxies and their progenitors in the EAGLE model, and to regulate the gas fractions of haloes. Fig. 2.10 also shows the CDDFs we get from the reference model in a 50 cMpc box (Ref-L050N0752) to verify that any differences are not due to the difference in box size.

The top panels of Fig. 2.10 show that, when AGN feedback is disabled, the CDDF is barely affected at the lowest column densities, intermediate column densities are slightly less common, and the highest column densities occur more frequently. The decrease at intermediate column densities is larger for O VIII than for O VII. At column densities $\gtrsim 10^{16.5} \text{ cm}^{-2}$, the CDDFs are not converged with pixel size in the column density maps (see Appendix 2.A). We show larger column densities in this plot mainly to show the difference with the no AGN simulation at high column densities more clearly, but relative differences in this range may not be reliable.

Since massive galaxies in the no AGN simulations produce too many stars, but too weak outflows, we might expect that some of this difference, especially at higher column densities probing halo gas, might be due to differences in gas metallicity. To check this, we do a similar comparison in the bottom panels of Fig. 2.10, except we assume all gas has a metallicity of




Figure 2.10: The top panels (a, b) compare the reference and no AGN models, to show the effect of AGN feedback on the $z = 0.1$ O VII (left panels) and O VIII (right panels) CDDFs. The column density distribution in the no AGN simulation (NoAGN-L050N0752, ‘noAGN-50’) is compared to the reference simulation (Ref-L100N1504, ‘Ref-100’). We also show the 50 cMpc reference simulation (Ref-L050N0752, ‘Ref-50’) results to demonstrate that differences between the reference and no AGN simulation distributions are not due to the size of the simulation domain. To gauge how much of the differences are due to differences in gas oxygen content, we show CDDFs made assuming all gas has a metallicity of $0.1 Z_{\odot}$ in the bottom panels. This shows that beyond the break in the CDDF, a boost in metal enrichment due to AGN feedback partly offsets its other effects.

0.1 times the solar value when we calculate the number of ions in our columns.

We then note that the differences below the break in the CDDF (almost) disappear when comparing constant metallicity results, while the differences above the CDDF break increase somewhat. We interpret the causes of these differences as follows. Within the haloes, we see larger column densities in the absence of AGN feedback. This could be due to higher densities, metallicities, or mass of the hot gas responsible for the absorption at these high column densities. The effect of AGN feedback at higher column densities increases, but only a bit, if we assume a fixed metallicity. This indicates that the main effect of AGN feedback is to decrease the density of the hot gas, while it also increases its metallicity. These effects partially cancel out, but the density effect dominates. The enhancement of metal ejection by AGN feedback also explains the increase in the CDDF below the break.

This picture is supported by the work of [Davies et al. \(2019\)](#). They showed that for haloes with masses $\geq 10^{12} M_{\odot}$ at $z = 0$ in EAGLE, the amount of AGN feedback injected into these halos is a good predictor of the halo gas fraction at fixed halo mass, and that this is due to AGN ejecting gas from the circumgalactic medium. In a follow-up paper, [Oppenheimer et al. \(2020a\)](#) showed that this feedback also decreases ion column densities around galaxies, though they discuss H I, C IV, and O VI.

2.3.6 Physical properties of the absorbers

To investigate the physical properties of the absorption systems, we first look at the temperature and density of gas traced by the O VII and O VIII ions. In Figs. 2.11 and 2.12, contours in different colours show the distribution of absorption system temperatures and densities in different column density ranges for O VII and O VIII, respectively. The greyscale shows the fraction of oxygen ions in the ionization state of interest, and the dotted contours show lines of constant net radiative cooling or heating time scale. The hydrogen number density displayed along the axis, used to calculate the ion balance as well as the cooling times, was calculated from the mass density assuming the primordial hydrogen mass fraction 0.752, but the difference with a solar hydrogen mass fraction conversion is negligible.

Comparison of the coloured contours and the greyscale shows that the absorbers reside in regions of the phase diagram where the fraction of oxygen ions in that state is high, as expected. These ion fractions account for photoionization by the UV/X-ray background (Haardt & Madau 2001). At high densities, the ion fraction becomes density-independent, as the influence of the ionizing radiation becomes negligible and the ion balance asymptotes to collisional ionization equilibrium (CIE).

As expected, the column density increases with the density of the absorbing gas, with the lowest column densities often coming from (almost) underdense gas. The percentages in the legend indicate the fraction of sightlines (including those with temperatures and densities outside the plotted range) in that column density range. The highest column densities are rare, accounting for less than 0.1% of sightlines. Furthermore, as the ion fraction colouring shows, most of the sightlines have their absorption coming mainly from photoionized gas: the gas at lower densities where the ion balance is density-dependent. However, at the highest column densities, especially where $N_{\text{OVII}} \gtrsim 10^{15.5} \text{ cm}^{-2}$, the absorption comes mainly from collisionally ionized gas.

For ionization models of observed absorbers, it is important to know what are reasonable assumptions when trying to establish what sort of gas an absorber traces. For both O VII and O VIII, photo-ionized gas dominates the absorption at lower column densities. For O VII, gas at the CDDF break ($N_{\text{OVII}} \sim 10^{16} \text{ cm}^{-2}$) is clearly dominated by gas in the CIE peak temperature range (Fig. 2.8). For O VIII, WHIM gas is also clearly dominant there, but CIE temperature gas is only just becoming dominant. At the CDDF break ($\sim 10^{16} \text{ cm}^{-2}$), CIE gas is important, but the absorption systems here have sufficiently low densities that photoionization still influences the ion fractions. For a given temperature, the ion fractions at $n_{\text{H}} \sim 10^{-5} \text{ cm}^{-3}$ (the value used by Nicastro et al. (2018) to model the absorbers they found) can differ from the CIE values by factors $\gtrsim 2$ within the CIE temperature range (i.e., at temperatures where, in CIE, the ion fraction is at least half the maximum value). Therefore, in modelling these absorption systems, a temperature consistent with CIE does not necessarily imply that CIE ion fractions will be accurate.

To quantify what fractions of gas are affected by this, we divide the temperature-density plane into four regions, defined by a temperature and a density cut⁸. For the temperature cut, we use the lowest temperature at which the ion fraction in CIE reaches half its maximum value: $T \approx 10^{5.5}, 10^{6.2} \text{ K}$ for O VII and O VIII, respectively. For the density cut, we focus on the same temperature ranges, where, in CIE, the ion fraction is at least half the maximum value. We consider the gas to be in CIE when it has a temperature above the cut, and a density above a minimum value. This minimum is the lowest density where the ion fractions in the

⁸Note that we use the intrinsic temperature and density of gas in EAGLE (not the ion-weighted values in projection), along with its ion content, for these calculations. This means we make cuts in the temperature-density plane of Fig. 2.13




Figure 2.11: Phase diagram of the O VII absorption system density and temperature in different column density bins in Ref-L100N1504 at $z = 0$, using 6.25 cMpc long sightlines. The rainbow-coloured contours indicate the distribution of ion-weighted average temperatures and densities for absorbers with column densities in different bins. The legend indicates the bins by their values of $\log_{10} N_{\text{OVII}} / \text{cm}^{-2}$; the percentages indicate what fraction of columns in the simulation (including those with zero column density) is in each bin. Solid and dashed contours enclose 50 and 99 % of the columns in each bin, respectively. The greyscale in the background indicates which fraction of oxygen nuclei are O VII ions. We indicate gas at fixed heating and cooling time scales using the dotted pairs of contours. They indicate constant net radiative cooling (upper contours) and heating (lower contours) time scales $t_{C/H}^{\text{net}}$ equal to the Hubble time t_H for gas with solar metallicity and with primordial abundances, using a [Haardt & Madau \(2001\)](#) background. The vertical brown line indicates the $z = 0$ average baryon density.

CIE temperature range defined above, differ from those at the highest density we have data for (10^3 cm^{-3}) by less than a factor 1.5. This means $n_{\text{H}} \approx 10^{-4.75}, 10^{-4.25} \text{ cm}^{-3}$ for O VII and O VIII, respectively. For O VII and O VIII, ≈ 32 and 8% of the ion mass is in CIE, respectively. Gas at lower densities, but with a temperature above the cut, is the gas that may be mistaken for gas in CIE based on temperature diagnostics, but where ionization modelling based on this assumption may cause errors. This is ≈ 45 and 19% of the O VII and O VIII mass, respectively. Gas below the density and the temperature cuts is not necessarily in pure PIE: the ion fractions are still influenced by collisional processes at the higher temperatures in this regime. This gas accounts for ≈ 23 and 72% of the O VII and O VIII mass, respectively. A small fraction of the O VII and O VIII mass, $\leq 0.5\%$, is at densities above the CIE threshold, but below the temperature cut. For these density thresholds and ions, there is more gas at




Figure 2.12: As Fig. 2.11 but for O VIII.

CIE temperatures that is not actually in CIE than there is CIE gas. However, note that this depends on choices we made. If we choose our density thresholds based on a maximum factor 2 difference between the ion fraction and full CIE ion fraction, both ions have more gas in CIE than in the error-prone regime.

We look into the effect of temperature further by comparing the b parameter range of Fig. 2.5 with the thermal broadening for different temperature ranges we see in figures 2.11 and 2.12. Note that those b parameters come from comparing total EW and column density, not line fitting, and do not account for instrumental broadening.

The thermal contribution to the line widths $b_{\text{th}} = \sqrt{2kTm^{-1}}$ is equal to 16–57 km/s for gas with temperatures $T = 10^{5.4}–10^{6.5}$ K. This range dominates the O VII CDDF at the column densities $\gtrsim 10^{15}$ cm $^{-2}$, where Fig. 2.5 shows we can estimate the line width from the column density and equivalent width. For O VIII, we find $b_{\text{th}} = 36–81$ km/s for temperatures around the O VIII collisional ionization peak ($T = 10^{6.1}–10^{6.8}$ K). For the gas at $T = 10^{7.0}$ K, reached by some of the highest-column-density O VIII, we find $b = 102$ km/s. For both ions, the lowest EWs we find are consistent with thermal broadening, but the typical b parameters for these absorption systems are larger than predicted by their ion-weighted temperatures. This strengthens our conclusion from visual inspection of virtual spectra that velocity structure is important for line widths, especially when the spectral resolution is too low to resolve components in absorption systems. This also means that measured line widths for these ions may not provide meaningful constraints on absorber temperatures.

We show contours indicating where the radiative net cooling (or heating) time equals the




Figure 2.13: Phase diagrams for the total gas mass (a), oxygen mass (b), and O VII (c) and O VIII (d) masses in Ref-L100N1504 at $z = 0$. The solid and dashed fuchsia contours enclose 50 and 99 % of the total mass. The green contours show the distribution of ion-weighted temperatures and densities that were computed for 6.25 cMpc long columns. This is similar to the distributions shown in figures 2.11 and 2.12 but for all non-zero ion column densities. The black, dot-dashed contours indicate a constant radiative cooling or heating time scale equal to the Hubble time for gas with solar abundances. The vertical brown lines indicate the $z = 0$ average baryon density.

Hubble time for gas with primordial and solar abundances in brown and black, respectively. We calculated these time scales using the tables of Wiersma et al. (2009a), which were also used for radiative cooling in the simulations. These include the effects of the evolving Haardt & Madau (2001) UV/X-ray background. For each metallicity, the lower-temperature contour indicates where the net heating time scale is equal to the Hubble time, with net heating within a Hubble time at lower temperatures ($C < H$). The higher-temperature contours indicate net cooling time scales, with net cooling within a Hubble time at higher densities ($C > H$). The pairs of contours converge in ‘peaks’ to the right at $\sim 10^4$ K, where radiative cooling and heating balance each other.

As column densities increase, more of the absorbing gas can cool within a Hubble time. For gas at column densities $\geq 10^{16}$ cm^{-2} , which we noted previously is dominated by halo-density gas, this means that the absorbing gas may be a part of cooling flows, where hot gas cools radiatively as it flows from the CGM to the central galaxy. As we saw in Fig. 2.9, gas with such column densities typically has roughly solar or larger oxygen abundances. However, some of it may also be gas that has recently been shock-heated in outflows.

We explore the effect of the gas phase distribution on the absorption system properties with Fig. 2.13. The panels show the distribution of gas for mass, oxygen mass, and O VII and O VIII mass. The horizontal lines for high densities at 10^4 K in the upper panels are the result of setting all star-forming gas to that temperature. The green contours (sightline distribution, for 6.25 cMpc long sightlines, as in figures 2.11 and 2.12) and fuchsia contours (mass distribution) seem to be similar, as expected, but the sightline distribution prefers lower-density




Figure 2.14: The range of counterpart column densities for O VII (a, b, c) and O VIII (d, e, f, g), given the column densities of neutral hydrogen (H I and H₂; a, d), O VI (b, e), Ne VIII (c, f) and O VII (g). We show the correlations for 6.25 cMpc sightlines in the reference simulation (Ref-L100N1504) at $z = 0.1$. The greyscale shows the distribution of sightlines in the column density of each pair of ions, while the lines show percentiles in the distribution of the ion on the y-axis at fixed column density on the x-axis. Large H I column densities ($N_{\text{HI}+\text{H}_2} \gtrsim 10^{15} \text{ cm}^{-2}$) predict a $\sim 50\%$ chance of finding a strong ($N \gtrsim 10^{15} \text{ cm}^{-2}$) O VII or O VIII counterpart. Strong O VI or Ne VIII absorption is a good predictor for strong O VII and O VIII absorption.

gas. This is due to its higher volume-filling fraction, hence higher covering fraction, relative to its mass fraction. We also show the radiative heating/cooling time contours at solar abundances as in figures 2.11 and 2.12. This demonstrates how the cooling and heating time scales shape gas conditions: most gas is either roughly at the temperature set by radiative heating over a Hubble time, or is too hot or diffuse to cool within about a Hubble time.

Comparing the ion mass distributions to the mass and oxygen distributions clearly shows the effect of ionization state fractions and metallicities. The low-density, low-temperature gas associated with the cool IGM contains few ions because it contains little oxygen. Comparing the ion mass distribution to the total oxygen mass distribution shows that the low-temperature, high-density gas (associated with haloes and galaxies) is not traced by these ions because neither collisional nor photoionization ionizes much of the oxygen to O VII or O VIII in this regime.

2.3.7 Correlations between O VIII, O VII, and Ne VIII, O VI, and H I absorption

We now look at how O VII and O VIII absorption correlate, and also how this absorption correlates with neutral hydrogen, O VI and Ne VIII. This is shown in Figs. 2.14 and 2.15. The neutral hydrogen we show is H I and H₂ together, where we count the total number of hydrogen atoms. We use the prescription of Rahmati et al. (2013) to obtain the hydrogen neutral fraction, since the ionization balance tables we use for the other ions do not account for gas self-shielding against ionizing radiation, which is important for H I fractions in cool, dense gas. For all but the highest column densities, the neutral fraction will be dominated by




Figure 2.15: As Fig. 2.14 but with the axes reversed. For a given O VII or O VIII column density, this indicates whether a counterpart agrees with what we find in EAGLE.

HI. For neutral hydrogen, we see that over a very large range of column densities ($N \sim 10^{15} - 10^{21} \text{ cm}^{-2}$), the trend between hydrogen column density and O VII and O VIII column density is flat, with a lot of scatter. However, we also see that in this range $\sim 50\%$ of the hydrogen absorbers have O VII or O VIII counterparts with column densities $\geq 10^{15} \text{ cm}^{-2}$, which is potentially observable, especially when searching at a specific redshift.

In fact, Kovács et al. (2019) have already reported extragalactic O VII absorption by stacking quasar spectra using the known redshifts of 17 Ly α absorbers. They did make an extra selection, using only absorbers associated with massive galaxies (stellar mass $\geq 10^{10} M_{\odot}$). The average equivalent width of the Ly α absorbers was 174.4 mÅ (corresponding to a column density $\sim 10^{14} \text{ cm}^{-2}$ for $b \approx 20 \text{ km s}^{-1}$), and they found an O VII column density of $1.4 \pm 0.4 \times 10^{15} \text{ cm}^{-2}$. Depending on the line width, this equivalent width may correspond to

a column density close to or in the flat part of Fig. 2.14a, and the O VII column density they measure (an average for their stacked spectra) is similar to the median we predict for that column density range.

The other ions correlate better with O VII and O VIII along 6.25 cMpc sightlines. Fig. 2.14 also shows how good O VI and Ne VIII are at predicting strong absorption in our two X-ray lines, and what sort of O VIII counterparts to expect for an O VII line. The tightest correlation is between Ne VIII and O VII, with a maximum 1–99% scatter of 1.1 dex at $N_{\text{Ne VIII}} > 10^{11} \text{ cm}^{-2}$. (The maximum scatter at fixed O VII column density, as shown in Fig. 2.15, is the same.) The tighter correlation between Ne VIII and O VII than between O VI and O VII is likely because Ne VIII is a higher-energy ion than O VI, with an ion fraction that peaks in almost the same band in the density-temperature plane as O VII, but with a narrower peak. We also notice from both figures in this section that O VII is better correlated with both the UV ions than O VIII.

We can see from Fig. 2.14 that a survey targeted at sightlines with observed high O VI or Ne VIII column densities has a greater chance of detecting O VII and O VIII systems than a blind survey. As an example, we showed in Section 2.3.3 that equivalent widths of $\sim 1 \text{ m\AA}$ require column densities of $\sim 10^{14.5} \text{ cm}^{-2}$ for O VII and $\sim 10^{15} \text{ cm}^{-2}$ for O VIII. For those threshold column densities, we predict a 50% chance of detecting an O VII (O VIII) counterpart to an O VI absorber at $N_{\text{O VI}} \approx 10^{13.4} \text{ cm}^{-2}$ ($10^{14.1}$). For Ne VIII, the corresponding column density is $N_{\text{Ne VIII}} \approx 10^{12.9} \text{ cm}^{-2}$ ($10^{13.3}$). These column densities are close to the breaks in the CDDFs $f(N, z)$ (see equation 2.1) of the FUV ions, but not beyond them. We find $f_{\text{O VI}}(10^{13.4}, z = 0.1) = 3 \times 10^{-13} \text{ cm}^2$ and $f_{\text{Ne VIII}}(10^{12.9}, z = 0.1) = 1.3 \times 10^{-12} \text{ cm}^2$ at the O VII threshold ($1.7 \times 10^{-14}, 3 \times 10^{-13} \text{ cm}^2$, respectively, for O VIII).

Fig. 2.15 is useful for checking whether an observed absorption system matches our expectations: given an O VII or O VIII column density, it shows what counterparts to expect. For neutral hydrogen, the scatter is very large at O VII or O VIII column densities $\gtrsim 10^{15} \text{ cm}^{-2}$, but for the other ions, the correlations are tighter.

2.3.8 Correlation between absorber properties

Coincident absorption from more than one ion can give powerful constraints on the temperature and/or density of the gas, particularly if the ions belong to the same element. However, this only works if the same gas is causing the different absorption lines. Since gas clouds do not have constant densities and temperatures, it is possible that different absorption lines arise in different parts of the same clouds. Since in planned blind X-ray absorption surveys with Athena, Arcus, and Lynx, O VII and O VIII are the main targets (Brenneman et al. 2016; Lumb et al. 2017; Smith et al. 2016; The Lynx Team 2018), it is useful to investigate how closely the absorption system properties for these ions match. To do this, we compare temperature and density weighted by O VII and O VIII along the same 6.25 cMpc sightlines. We also compare to O VI, which has a strong doublet in the FUV range that may provide further constraints without element abundance ratio uncertainties, and to Ne VIII, which we found in Section 2.3.7 correlates very well with O VII and O VIII. We do not compare to neutral hydrogen gas properties, since we expect neutral hydrogen to reside in cooler, denser gas than the other ions we considered in Section 2.3.7. In addition, any temperature or density constraints from hydrogen to oxygen ion column density ratios would depend on the (generally unknown) metallicity of the gas.

First, we explore density differences in Fig. 2.16. This shows the distribution of O VI,




Figure 2.16: The O VII-weighted density (vertical) of absorption systems plotted against the O VIII-weighted density (a), the Ne VIII-weighted density (b), and the O VI-weighted density (c) along the same sightline. These sightlines are our 6.25 cMpc long columns, and come from 32000² pixel projections of the reference (Ref-L100N1504) simulation at redshift 0. The greyscale shows the total sightline distribution, while the contours show the density distributions for different O VII column density bins, matching those in Fig. 2.11. These ranges are given in units of $\log_{10} \text{cm}^{-2}$. The light blue lines indicate where the densities are equal. This indicates that these ions tend to trace gas with similar densities along the same sightline, though there are small systematic differences at low column densities.

Ne VIII, O VII and O VIII-weighted densities⁹ in the 6.25 cMpc columns we use to measure column densities. Contours show the distribution of sightlines in different O VII column density ranges. The light blue lines indicate where the densities are equal. We can see that at low column densities, where the gas is photoionized ($n_{\text{H}} \lesssim 10^{-5} \text{cm}^{-3}$), O VII tends to trace slightly higher-density gas than O VIII. This makes sense, since lower-density gas in the same radiation field is more highly ionized. The differences here are largely ($\geq 98\%$ of sightlines) below 0.6 dex in all column density ranges and for all pairs of ions we show. The median differences are typically small, $\lesssim 0.1$ dex, except at $N_{\text{O VII}} > 10^{16.5} \text{cm}^{-2}$ for O VII and O VIII, where the median difference is still < 0.2 dex. The differences between other combinations of these ions are similar.

Next, we look into temperature, shown in Fig. 2.17, similar to Fig. 2.16, but for ion-weighted temperatures. Between O VII and O VI or O VIII, differences are typically larger than for density, and clearly systematic: the higher-energy O VIII ions prefer hotter gas than O VII at all (column) densities, while O VII traces hotter gas than O VI. The median differences at lower column densities ($N_{\text{O VII}} \lesssim 10^{15.5} \text{cm}^{-2}$) are $\lesssim 0.2$ dex, but at higher column densities, the median temperatures differ by up to ≈ 0.3 dex for O VII and O VI or O VIII. This occurs at column densities tracing collisionally ionized gas, where the ions both tend to trace their own CIE peak temperature gas. For Ne VIII and O VII, we see a hint of a similar CIE peak preference in the narrowing of the range of absorber temperatures as the O VII column density increases, but because these two ions have similar peak temperatures, their median differences remain $\lesssim 0.1$ dex.

We also compared temperatures and densities traced by O VI and O VIII along the same

⁹The axes show density as hydrogen number density, but the densities used for the plot are mass densities, converted to hydrogen number densities assuming a primordial hydrogen mass fraction of 0.752.




Figure 2.17: As Fig. 2.16 but for temperature. Between O VII, and O VI and O VIII, there are systematic temperature differences at all column densities, but especially the larger ones. The O VII and Ne VIII temperatures agree within 0.3 dex for $\geq 98\%$ of absorbers at all column densities shown.

6.25 cMpc sightlines. We found similar trends in the differences to those we found between O VI and O VII and between O VII and O VIII, except that both the scatter and systematic differences were larger.

The differences in gas temperature and density probed by different ions in the same structures may be large enough to be a potential issue when modelling the absorption as coming from the same gas. In future work, we intend to test this by applying ionization models to column densities measured from virtual absorption spectra. This will also enable us to test whether enforcing different lines to have identical redshifts would help reduce the systematic errors given the energy resolution expected for upcoming missions.

2.4 Discussion

2.4.1 Detection prospects

We examine the prospects for detecting O VII and O VIII line absorption with the planned mission Athena (Lumb et al. 2017; Barret et al. 2018), and the proposed missions Arcus¹⁰ (Smith et al. 2016; Brenneman et al. 2016), and Lynx (The Lynx Team 2018). Detectability depends on the background source as well as on the absorber. Here, we consider quasars/blazars like those discussed in the mission science cases. The Arcus and Athena science cases include detecting the WHIM and measuring the CDDF for WHIM absorbers (Brenneman et al. 2016; Lumb et al. 2017). For Lynx (The Lynx Team 2018), which would launch later if approved, the focus is on characterising this gas.

The Arcus X-ray Grating Spectrometer (XGS) (Smith et al. 2016; Brenneman et al. 2016) should be able to find absorbers with equivalent widths ≥ 4 mÅ at 5σ against blazars of a brightness threshold met by at least 40 known examples, at redshifts up to 0.2 with exposure times ≤ 500 ks. There are also plans to measure absorption against gamma-ray bursts. Fig. 2.6

¹⁰This was one of three (later two) proposed NASA MIDEX missions in the most recent (February 2019) funding round, but it was not chosen.

Table 2.2: Expected number of absorption systems per unit redshift above a given observer-frame EW threshold $dN(> EW_{\text{obs}})/dz$ for O VII and O VIII at different redshifts. Here, we use 6.25 cMpc slice CDDFs and 100 cMpc column density equivalent width relations (with scatter) at each redshift. The results come from the reference (Ref-L100N1504) EAGLE simulation.

EW_{obs} mÅ	O VII			O VIII		
	$z = 0.1$	$z = 0.2$	$z = 1.0$	$z = 0.1$	$z = 0.2$	$z = 1.0$
6.8	1.98	2.43	4.33	0.533	0.762	3.23
5.2	2.95	3.53	5.59	0.987	1.33	4.49
4.0	4.06	4.77	6.92	1.60	2.09	5.94
3.0	5.57	6.37	8.61	2.49	3.23	7.74
1.0	14.0	15.1	17.2	9.87	11.8	17.5

shows that this threshold is roughly at the knee of the equivalent width (EW) distribution. Fig. 2.5 shows that Arcus would mainly be probing absorption systems that do not lie on the linear part of the curve of growth with this EW regime. Smith et al. (2016) expect to find around 40 O VII absorbers at $EW > 4$ mÅ, using about 20 blazar background sources probing redshift paths that add up to $\Delta z \approx 8$.

We can make a similar prediction for the number of detections in such a survey. The results are shown in Table 2.2. At each redshift, we make predictions using EW distributions obtained from 6.25 cMpc slice CDDFs, using the (rest-frame) column density EW relation for a full sample of 100 cMpc sightlines at each redshift. We obtain the column density equivalent width relation at $z = 0.2$ and $z = 1$ in the same way as described in Sections 2.2.3 and 2.3.3 for $z = 0.1$, except that we select the higher-redshift sightlines by O VII and O VIII column density alone, ignoring O VI. This should not significantly affect our results, since the ion-selected subsamples did not yield substantially different results from the full sample at $z = 0.1$.

In this calculation, changes in the predicted number of absorbers with redshift are due to four effects. First, the CDDFs evolve, as we saw in Fig. 2.3. The evolution of the CDDFs between redshifts 0.1 and 0.2 is small, but differences with $z = 1.0$ are larger. Secondly, the relation between column density and equivalent width evolves. This evolution is weak, affecting the values in Table 2.2 by $\lesssim 4\%$. Thirdly, these CDDFs concern the absorption system distribution with respect to dX , not dz . Between redshifts 0.0 and 0.2, dX/dz increases from 1 to 1.30. Finally, the instrumental EW threshold applies to the observed EW. This means we can probe lower-EW systems at higher redshift, assuming e.g. quasar fluxes and backgrounds/foregrounds are equal.

For $\Delta z = 8$, we expect to see 32 (38) O VII absorption systems and 13 (17) O VIII absorption systems if redshift ≈ 0.1 (0.2) contributes most to the survey path length. These values agree roughly with the ≈ 40 O VII absorbers expected by Smith et al. (2016). They expect 10–15 O VIII counterparts to O VII lines along these sightlines, which is compatible with the total number of O VIII absorbers we expect. From the correlations between O VII and O VIII absorption systems, (Fig. 2.15g), and the column-density-equivalent-width relation we find (Fig. 2.5), we expect that some of the detectable O VIII absorbers do not have observable O VII counterparts: for absorbers with $N_{\text{O VIII}} \geq 10^{15.5} \text{ cm}^{-2}$, 31% (3.7%) have O VII counterparts with $N_{\text{O VII}} \leq 10^{15.5} \text{ cm}^{-2}$ ($\leq 10^{15}$).

As shown in appendix 2.C, we do not see major effects of the large-scale structure on variations in measured CDDFs in the EAGLE reference simulation (Ref-L100N1504). We expect the variations in measured EW distributions are Poisson as well. That would imply

that the expected survey-to-survey scatter in the number of absorption systems detected is similar to the maximum effect of varying the background source redshift distribution (as long as $z \lesssim 0.2$ still holds).

If the sensitivity can be increased to a minimum EW of 3 mÅ, we would expect to see about 1.3–1.6 times as many absorption systems at redshifts 0.1–0.2. If, for deep surveys, a sensitivity of 1 mÅ can be achieved, we would expect to find double to triple the number of O VII absorption systems and double to quadruple the number O VIII absorption systems per unit redshift as at 3 mÅ.

As for resolving line widths or line shapes, Fig. 2.4 shows the expected effects of line broadening and Poisson noise on observations. This figure suggests that constraining the line widths of single absorption components will be difficult, especially since what would appear to be single absorption components often are not. Comparing our approximate line widths from Fig. 2.5 to the spectral resolution expected for the instrument supports this.

The Arcus XGS resolving power was planned to be $\lambda/\Delta\lambda = 2000$, 2500 below and above 21.6 Å, respectively (Smith et al. 2016). For our broadest O VII lines (≈ 220 km/s, we have $\lambda/\Delta\lambda \approx 1.4 \times 10^3$ assuming a single line. For O VIII, $\lambda/\Delta\lambda \approx 1.0 \times 10^3$ for the broadest lines. The b parameters for O VIII in Fig. 2.5 were calculated by fitting the relation between column density and equivalent width by the two doublet lines, each with a width b . Fitting the O VIII curve of growth using a single line instead (using the sum of oscillator strengths and oscillator-strength-weighted average wavelength) only adds about 10 km/s to the best-fit b parameter. This would suggest we may be able to constrain some of the line widths with the Arcus XGS, though a good measurement of the line widths may be possible for only the broadest lines. Bear in mind that the b parameters we use are a rough measure, determined by the curve of growth rather than the actual structure of the absorption system, and that unresolved turbulence may cause lines to be broader than predicted from the velocity structure resolved in the simulation alone. Although resolving individual components will be difficult, Arcus will be able to decompose many systems into multiple components.

On Athena, the instrument of interest is the X-IFU (Barret et al. 2018). The science requirements for Athena as a whole, including this instrument, are described by Lumb et al. (2017). Weak line 5σ detections should be possible from $EW = 0.18$ eV against ‘bright sources’. The plan is to measure WHIM absorption against 100 BLLacs and 100 gamma-ray bursts, to study the WHIM at $z < 1$.

The minimum EW translates to 6.8 and 5.2 mÅ (rest-frame) for O VII and O VIII respectively, for which we show the expected number of absorption systems per unit redshift in Table 2.2. These limits are for 50 ks observations against a point source with a 0.5 mCrab flux in the 2–10 keV energy band ($F_{2-10\text{keV}} = 1 \times 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ for $\Gamma = 1.8$). These are larger EWs than expected for Arcus, though still above the knee of the EW distribution. Note that the observing time used here is ten times smaller than that of the Arcus specification, so some differences are expected. Rather than attempt to correct for this difference, we explore what Athena would see if the X-IFU is used as described.

The exact number of absorption systems expected from the survey will depend on the redshift distribution of the quasars and gamma-ray bursts used to probe the WHIM. We make predictions for Athena in the same way as for Arcus. Over the redshift range 0.1–1, Table 2.2 shows large changes in predicted absorber densities in redshift space. The expected number of absorption systems at the Athena sensitivity thresholds is less than expected in the Arcus survey, given the current factor 10 difference in planned exposure times, but if $\Delta z = 50$ –100 is achieved, we would expect the Athena survey to find a larger number of O VII absorption systems. Since the aim is to observe 100 BLLacs and 100 gamma-ray burst afterglows, such

a survey size seems reasonable.

For the Lynx XGS (The Lynx Team 2018), a goal is to use O VII and O VIII absorption to characterise the hot CGM of galaxies with halo mass $\gtrsim 10^{12} M_{\odot}$ and the hot gas in filaments with overdensities $\gtrsim 30$. For the filaments, this means looking for column densities $\gtrsim 10^{15} \text{ cm}^{-2}$ and rest-frame EWs $\sim 1 \text{ m\AA}$. Indeed, The Lynx Team (2018) expects to be able to detect lines at EW $\gtrsim 1 \text{ m\AA}$ outside halo virial radii, against bright AGN ($F_{0.5-2 \text{ keV}} \sim 1 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$) with an average exposure time of $\approx 60 \text{ ks}$. Table 2.2 shows the predicted number of detected O VII and O VIII absorbers in blind surveys with this sensitivity.

In the future we plan to improve on the rough estimates provided in this section by creating many virtual Arcus and X-IFU spectra, processing them through the instrument models, and analysing them as if they were real data.

2.4.2 Caveats

We have studied the impact of resolution and some technical choices in appendices 2.A and 2.B. There are some other approximations and uncertainties we will discuss here.

First, we assume the ion fractions are set by the temperature and hydrogen number density of the gas, assuming collisional and photoionization equilibrium with a uniform (but evolving) Haardt & Madau (2001) UV/X-ray background. Cen & Fang (2006) found that there were differences in the equivalent width distributions derived with and without equilibrium assumptions. These differences are comparable to differences between different simulations for O VII, shown in Fig. 2.6 and smaller than those differences for O VIII. We therefore consider ionization equilibrium a reasonable model simplification. Yoshikawa & Sasaki (2006) also discuss non-equilibrium effects, and plot at what temperatures and densities they expect non-equilibrium effects to be strongest. They found that WHIM detectability was not impacted much by non-equilibrium effects, but that line ratios were impacted more significantly, further complicating measurements of absorber temperatures. Oppenheimer et al. (2016) found using much higher-resolution simulations that non-equilibrium ionization was generally unimportant for O VI.

Another issue relating to ionization and equilibrium is the impact of local ionization sources. For these high-energy ions, we would expect AGN to be the most important. Oppenheimer & Schaye (2013) pointed out that if, as expected, AGN vary on time scales similar to or shorter than the recombination time scale in the IGM, then a large fraction of high-ionization metal absorbers may be more highly ionized than they would be if they were only ever illuminated by the metagalactic background radiation. This could then even affect gas near galaxies that currently do not show detectable AGN activity. Segers et al. (2017) found that in the inner parts of the halo, a variable AGN could have large effects on O VI column densities, with smaller effects out to $2 R_{\text{vir}}$. These effects persisted between periods where the AGN was ‘on’. They did not study the effect of this on O VII and O VIII, so the precise impact here is unknown. The effect could be important in regions where the temperature is too low for collisional ionization to be effective. Here, AGN radiation could add a population of cooler, high column density systems. However, it might also affect the population of strong absorbers in our current sample, if the X-ray flux is high enough to ionize their hot, dense gas even further.

A further difficulty lies in defining absorbers. We have somewhat avoided this issue when discussing column densities, by limiting our work to a reasonable proxy for absorption systems. However, in observations, observers usually fit line profiles to their spectra and gather

information on each line. Although we have synthetic spectra, we have not attempted to identify and characterise individual absorption components in these, since the finding and identification of these components will generally depend on the observing instrument and exposure time, and is time-consuming. From visual inspection of Athena and Arcus mock spectra, it seems that absorption systems would usually look more or less like single lines at the spectral resolution of Arcus, and would nearly always look this way at the Athena spectral resolution. This justifies limiting our discussion to entire absorption systems, though we do plan to revisit this question in future work.

2.5 Conclusions

We have used the EAGLE cosmological, hydrodynamical simulations to predict the rate of incidence and physical conditions of intergalactic O VII ($\lambda = 21.6019 \text{ \AA}$, $E = 574 \text{ eV}$) and O VIII ($\lambda = 18.9671, 18.9725 \text{ \AA}$, $E = 654 \text{ eV}$) absorbing gas. In the largest simulation, 40 % of gas-phase oxygen is in the form of these two ions, making them the key tracers of cosmic metals. We have extracted column density distribution functions (CDDFs) for O VII and O VIII by measuring the number of ions in a grid of long, thin columns within slices through the simulation box. Assuming a uniform metallicity instead of the predicted metallicities results in much steeper CDDFs (Fig. 2.2) because higher column densities are predicted to correspond to higher metallicities. The most notable feature in the CDDFs, which evolve only mildly between $z = 1$ and $z = 0$, looks like a power-law break, and occurs at column densities $\sim 10^{16} \text{ cm}^{-2}$ (Fig. 2.3). This break occurs where absorption systems reach overdensities $\sim 10^2$ (Fig. 2.7), which indicates it is likely caused by a transition from absorption by sheet and filament gas to absorption by denser halo gas.

AGN feedback substantially modifies the CDDF at high column densities. It impacts both the position of the break in the CDDF and its slope after the break, indicating an impact on haloes. Without AGN feedback, there are slightly fewer absorption systems at column densities of $\sim 10^{15} - 10^{16} \text{ cm}^{-2}$ and significantly more at column densities $\gtrsim 10^{16} \text{ cm}^{-2}$ (Fig. 2.10). Before the break, this is because AGN feedback increases the metallicity of the gas. Beyond the break, AGN feedback still increases the metallicity, but column densities decrease due to the decrease in the hot gas densities.

We used a large set of synthetic absorption spectra to determine the relation between equivalent width (EW) and projected column density. The best-fit b -parameters to the curve of growth are 90 and 158 km s^{-1} for O VII and O VIII, respectively, but the scatter in $b(N)$ is large (Fig. 2.5). We used the CDDF to predict the EW distribution in the EAGLE reference simulation at redshift 0.1 (Fig. 2.6), accounting for the scatter by using our sample's EW distribution at fixed column density in the conversion. We predict an O VII EW distribution consistent with the recent observations of [Nicastrro et al. \(2018\)](#).

By measuring the ion-weighted gas temperatures and densities along the same columns, we investigated the physical properties of our absorption systems. Overall, for 32 % (8 %) of O VII (O VIII) the ion fractions are within a factor 1.5 of the values corresponding to collisional ionization equilibrium, while for most of the remaining gas both collisional ionization and photoionization by the metagalactic background radiation are important for the ion fractions. Although O VII and O VIII often trace gas that is (partly) photoionized, O VII and O VIII absorption systems with column densities $\gtrsim 10^{16} \text{ cm}^{-2}$ mostly trace gas at temperatures close to those where their ion fractions peak in collisional ionization equilibrium: $\sim 10^6$ and $\sim 10^{6.5} \text{ K}$, respectively. At column densities $\sim 10^{15} - 10^{16} \text{ cm}^{-2}$, much of the gas is at temperatures consistent with collisional ionization equilibrium, but the ionization fractions

are not the same as the purely collisional values (figures 2.11 and 2.12). There are systematic differences between the temperatures traced by O VII and O VIII for the same absorbing structures (Fig. 2.17). Along with density differences, this may lead to difficulties with using column density ratios to estimate gas temperatures. Instrumental line broadening and unresolved absorption system structure mean that temperatures will also typically not be well-constrained by line widths.

On the other hand, the column densities of O VII and O VIII are correlated with each other, and with those of ions with detectable absorption lines in the far UV (Figs. 2.14, 2.15). The presence of neutral hydrogen absorption with column density $\gtrsim 10^{15} \text{ cm}^{-2}$ implies a $\sim 50\%$ chance of finding an O VII or O VIII absorber with column density $\gtrsim 10^{15} \text{ cm}^{-2}$, but the probability does not increase further for even stronger H I absorption (until ISM densities are reached). High O VI or Ne VIII column densities are good predictors of the presence of detectable O VII and O VIII column densities (Fig. 2.14). This suggests that pre-selection based on these UV lines is an efficient (though necessarily biased) search strategy for the X-ray lines.

Acknowledgements

We thank the referee for their helpful comments. We would like to thank Dylan Nelson and Enzo Branchini for sharing their data, John Helly for programs we used to access EAGLE data (READ_EAGLE), and Volker Springel for the original version of the code we use to project particles onto a grid (HSM LANDPROJECT). We would also like to thank Ali Rahmati for help testing the code we use to make projections and setting up SPECWIZARD. We also made use of WebPlotDigitizer (Rohatgi 2018) and the NUMPY (Oliphant 2006), SCIPY (Jones et al. 2001), H5PY (Collette 2013), and MATPLOTLIB (Hunter 2007) PYTHON libraries, as well as the IPYTHON (Pérez & Granger 2007) command-line interface. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. RAC is a Royal Society University Research Fellow. FN acknowledges funding from the INAF PRIN-SKA 2017 program 1.05.01.88.04.

Appendix

2.A CDDF convergence

In this appendix, we test the convergence of the CDDFs with pixel size, slice thickness, box size, and simulation resolution, and give more details on how we tested the effect of the projection kernel shape. The default projection we will compare other CDDFs to is that obtained from the reference simulation (Ref-L100N1504) at redshift 0.1, with 16 slices of 6.25 cMpc each along the z -axis, and 32000^2 pixels in the x - and y -directions. The assumed gas distribution for one SPH particle in this ‘standard’ CDDF is given by the C2-kernel (Wendland 1995) used in the simulation. The temperature of star-forming gas is set to 10^4 K for the ion fraction calculation. Any properties of CDDFs that are not otherwise mentioned here, are as described above.




Figure 2.18: Convergence of the column density distribution function (CDDF, $f(N)$) with pixel size. The left panels show O VII, the right panels show O VIII. In all panels, the solid line shows the CDDF as used throughout the paper: the column density is calculated on a grid of 32000^2 pixels, for each of 16 slices through the 100 cMpc the box along the projection axis. These CDDFs are from the reference (Ref-L100N1504) simulation at $z = 0.1$. In the top panels (a, b), this CDDF is compared to that obtained by degrading the resolution of the 3.125^2 ckpc² pixel grid to get grids of 6.25^2 (dashed), 12.5^2 (dot-dashed) and 25^2 ckpc² (dotted) pixels. The legends indicate the pixel sizes. In the bottom panels (c, d), the column density distributions extracted from two smaller regions (10×10 cMpc²) at this standard resolution (3.125^2 ckpc² per pixel) are compared with the CDDFs from the same region at nine time higher resolution. The CDDFs at different resolution are shown relative to the highest-resolution CDDF for the same region.

In Fig. 2.18, we look at the effect of changing the pixel size. Using lower resolution means underestimating the fraction of columns with large column densities, while assessments of lower column densities remain accurate. Our CDDFs appear to be converged at our standard pixel size for column densities $\lesssim 10^{16.5}$ cm⁻². At higher resolutions than our standard, we see problems arise at a similar column density.

In Fig. 2.19 we investigate the effect of using slices of different thickness to calculate the column densities. The clearest effects are at low column densities. Thinner slices result in larger values of $f(N)$, which suggests that at low column densities, absorption often comes from multiple systems along the same line of sight. An optimal slice thickness is difficult to choose; we want to use a size which is small enough that the absorption in one column is (mostly) coming from a single system, but large enough that single systems are (mostly) not split over two columns along the line of sight. We choose to use slices of 6.25 cMpc, which matches UV-ion CDDFs well below the CDDF break (Rahmati et al. 2016), but bear in mind that this choice of, essentially, how to define absorption systems affects our results somewhat. Note, however, that the differences between CDDFs of different slice thickness are small compared to the range spanned by the CDDF. This standard slice thickness cor-

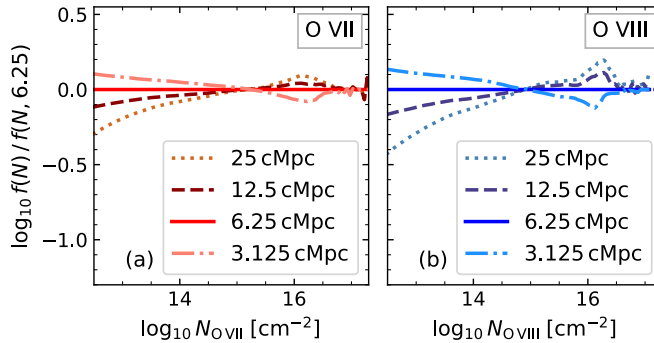


Figure 2.19: Convergence of the column density distribution function (CDDF, $f(N)$) with slice thickness. The left panel shows O VII, the right panel shows O VIII. In both panels, the solid line shows the CDDF as used throughout the paper: the column density is calculated on a grid of 32000^2 pixels, for each of 16 slices the the 100 cMpc box along the projection axis. For the other lines, the thickness of the box slices for which the column density is calculated is varied as indicated. These CDDFs are from the reference (Ref-L100N1504) simulation at $z = 0.1$. The differences between our default 6.25 cMpc slice CDDF and CDDFs using thicker or thinner slices are mostly $\lesssim 0.2$ dex, and always $\lesssim 0.2$ dex in this column density range when comparing to 3.125 and 12.5 cMpc slices. These differences are small compared to the more than 8 orders of magnitude the CDDFs span in this column density range.

responds to 404 km s^{-1} in velocity space ($z = 0.1$, rest-frame) when only accounting for the Hubble flow.

Next, we investigate whether the CDDFs are converged in the simulations themselves. Specifically, we investigate the effect of simulation volume and resolution. In Fig. 2.20 we compare different simulations to each other. Here, differences may also come from differences in initial conditions (cosmic variance). In the top panels, we examine the effect of box size. In smaller boxes, the distributions are noisier because they are less well sampled, but there is also a physical difference: smaller boxes do not sample some of the large-scale modes of the initial density variations, so smaller boxes will contain fewer high-mass, large, hot systems. The 100 cMpc box has no haloes with masses $M_{200c} > 10^{15} M_{\odot}$, and only seven with $M_{200c} > 10^{14} M_{\odot}$, so we do not expect complete convergence with box size. However, the figure indicates that, at column densities $\lesssim 10^{16.5} \text{ cm}^{-2}$, convergence with box size is sufficient. The $\lesssim 0.2$ dex difference between the CDDFs for the largest two boxes is much smaller than the range spanned by the CDDFs in the same column density range, and larger differences occur at column densities where the CDDF is not converged with projection resolution.

In the bottom panels, we investigate the effect of simulation resolution (distinct from the projection resolution we investigated in Fig. 2.18). We compare three 25 cMpc boxes, since we do not have higher-resolution simulations of larger boxes. The comparison to Ref-L025N0752 tests strong convergence and Recal-L025N0752 tests weak convergence in the terminology of Schaye et al. (2015). The Ref-L025N0752 simulation was run using the exact same parameters as Ref-L025N0376, except for the resolution, and the addition of initial density perturbations on smaller scales in the higher-resolution box. However, the stellar and AGN feedback in the EAGLE simulation is calibrated to reproduce the redshift 0.1 galaxy stellar mass function, the relation between black hole mass and galaxy mass, and reasonable galaxy sizes. This calibration depends on the simulation resolution, so Schaye et al. (2015) also recalibrated the subgrid model for feedback to produce a similarly calibrated simulation at higher resolution:




Figure 2.20: Convergence of the column density distribution function with simulation box size (top panels) and resolution (bottom panels). The left panels show O VII, the right panels show O VIII. The plots show the CDDF $f(N, z = 0.1)$, as described in equation 2.1 relative to a reference CDDF. In the top panels, the CDDFs from Ref-L025N0376 and Ref-L050N0752 are compared to the standard Ref-L100N1504 CDDF. In the bottom panels, Ref-L025N0376 (‘Ref’) is compared to the Ref-L025N0752 and Recal-L025N0752 (‘Ref-hi’, ‘Recal-hi’) CDDFs.

Recal-L025N0752. We consider this the most relevant comparison.

The weak convergence of the CDDFs is particularly good: for column densities $< 10^{16.3} \text{ cm}^{-2}$, differences are $\lesssim 0.3$ dex. They get larger at higher column densities (particularly for O VIII), but this is in a regime where the CDDFs in the 25 cMpc box are poorly converged compared to the larger boxes (top panels of Fig. 2.20), and the behaviour of the CDDFs here may therefore not be representative of what the effect of higher simulation resolution might be in a larger box.

Finally, as we mentioned in section 2.2.2, we found that the effect of the choice of SPH kernel is small. Here, we give some more details on the kernels we compared. Our default is the Wendland (1995) C2 kernel, with shape k given by

$$k_{C2}(u) = (1 - u)^4(1 + 4u), \quad (2.4)$$

where u is the distance to the particle position normalized by the smoothing length of the particle. The kernel is zero at $u > 1$. Similarly, the Gadget kernel (used by Springel (2005), from Monaghan & Lattanzio (1985)) is described at $u \leq 1$ by

$$k_{\text{Gadget}}(u) = \begin{cases} 2.5465 + 15.2789(u - 1)u^2 & \text{if } u < 0.5 \\ 5.0930(1 - u)^3 & \text{if } 0.5 \leq u \leq 1. \end{cases} \quad (2.5)$$

Both kernels are normalized to unity when integrated over the surface of a unit circle to ensure mass conservation in the projection. The C2 kernel is more centrally concentrated.




Figure 2.21: A comparison of methods to obtain the rest-frame equivalent width distribution from the CDDF. The rest frame EW distributions for O VII, $\lambda = 21.60 \text{ \AA}$ (a) and the O VIII, $\lambda = 18.97 \text{ \AA}$ doublet (b) in the Ref-L100N1504 simulation at redshift 0.1 are shown. The lines labelled ‘COG- b ’ are for EWs calculated from the CDDF using a single curve of growth with the best-fit b parameters to the relation between \log EW and the column density obtained for a set of absorption spectra. The dotted line (‘COG-lin’) is for EWs calculated using the linear curve of growth for each ion. For the lines labelled ‘matrix’, the conversion from CDDF to EW distribution was done using a matrix generated from the column density-EW relation in Fig. 2.5, which accounts for scatter. We use the CDDF for 6.25 cMpc slices.

When both kernels are normalized, their absolute difference is largest at the centre of the distribution: $k_{C2}(0) = 2.23$ and $k_{\text{Gadget}}(0) = 1.82$. The maximum relative difference is at the edge of the distribution, and diverges as $u \rightarrow 1$ due to the different slopes of the kernels near zero. The differences between the CDDFs we obtain using these kernels is $\lesssim 0.05$ dex for $10^{11} \text{ cm}^{-2} < N_{O_{VII,VIII}} < 10^{16.5} \text{ cm}^{-2}$.

2.B Technical choices for the EW distribution

Fig. 2.21 shows the rest-frame EW distribution of O VII and O VIII calculated from the CDDF using different methods. In the first two, we impose a one-to-one relation between column density and EW. We use the linear curve of growth or the curve of growth with the best-fit b parameter to the relation between \log EW and column density obtained for a sample of absorption spectra. The third method, which is our fiducial model, accounts for scatter in the relation between column density and EW by using a matrix generated from the points in the column density-EW relation from Fig. 2.5. We obtain this matrix by choosing column density and EW bins, then making a histogram of our sample of sightlines in column density and EW. After normalisation, we use this matrix to convert our column density histograms into EW histograms.

The differences between the EWs calculated using the different methods are as expected. The matrix conversion is our preferred method, because it models the scatter. Since the linear curve of growth is only valid for unsaturated absorption, it gives the maximum EW for any column density. This method overestimates the number of high EW systems ($EW \gtrsim 10 \text{ m}\text{\AA}$). The difference between the matrix and best-fit curve-of-growth methods is whether scatter is included. Including scatter increases the predicted number of absorption systems at the

highest EWs because the column density distribution is steep and declining at the highest column densities. This means that scatter at fixed column density attributes larger EWs to more lower column density absorption systems than it attributes smaller EWs to larger column density absorption systems, in absolute terms. This is why the single best-fit b -parameter curve of growth conversion underestimates the number of high EW absorption systems. The discrepancy is larger for O VII than for O VIII, because, as Fig. 2.5 shows, O VII EWs have a larger scatter at fixed column density and have significant scatter over a larger range of column densities than O VIII EWs. This means that including scatter in the modelling of the EW distribution has a larger effect for O VII.

2.C Mock survey CDDFs

We are interested in the variation in measured CDDFs for different survey sizes. In particular, we investigate whether clustering along the line of sight on scales of 6.25–50 cMpc causes deviations from Poisson statistics, which we find is not the case. We measure survey size by the absorption distance ΔX , as given by equation 2.2. We measure this at redshift 0.1, and use the Hubble flow across the box to calculate dz . We conduct mock surveys of the 100 cMpc box, of sizes $\Delta X = 0.5, 1.0, 10, 100, 1000$. To create one mock survey, we choose (x, y) positions at random, then take the column densities for all the 6.25 cMpc columns along the z -axis at that position. This mimics searching for absorption systems along longer sightlines than we use to measure the column densities. Note that due to the size (and periodicity) of the box, and the extent over which we measure a single column density, this only probes the effect of large-scale structure correlations for separations between 6.25 and 50 cMpc.

Because we are interested in the statistics of these surveys, we conduct 100 such mock surveys for O VII and O VIII for each ΔX . For comparison, we also create ‘random’ surveys of the same size, where for each random (x, y) , we select a single 6.25 cMpc slice at random. We make cumulative distributions for these surveys. Then at each threshold column density, we considered the distribution in absorption system counts among the different mock surveys. We show the 10th (blue), 50th (green, median), and 90th (red) percentiles of this distribution as the dashed lines in Fig. 2.22.

In Fig. 2.22, we do not show theoretical distributions, but the random surveys described above. To probe what range in these percentiles is consistent with variation between survey samples, we use ten different samples of 100 random surveys. We plot these in Fig. 2.22 with colors matching the same percentiles for the mock surveys.

The figure shows that the mock surveys for the different ions and survey sizes are consistent with the random surveys: in other words, large-scale structure on the scales we can probe does not appear to affect the survey statistics for measuring CDDFs. We also checked the relative differences between these percentiles and Poisson and binomial distributions: these were centred on zero, and the mock survey distributions differed from the theoretical values by amounts consistent with the random surveys. Using the CDDF instead of the cumulative distribution gives the same results.

Note that Fig. 2.22 can also be used to estimate a reasonable range of detected absorption systems above a given column density threshold for different survey sizes.




Figure 2.22: The cumulative column density distributions for mock surveys of sizes $\Delta X = 1.01$ ($\Delta z = 0.878$; a, b), $\Delta X = 10.0$ ($\Delta z = 8.66$; c, d), and $\Delta X = 100$ ($\Delta z = 86.7$; e, f), of O VII (a, c, e) and O VIII (b, d, f) in the reference (Ref-L100N1504) simulation at redshift 0.1. The column densities are calculated for 16 6.25 cMpc slices, each with 3.125^2 cMpc² pixels. The dashed lines show mock surveys, where (x, y) positions were selected at random, and then each slice along the z -axis at that position was included, mimicking the measurement of different absorption systems along full 100 cMpc sightlines. This accounts for clustering along the line of sight on scales between 6.25 and 50 cMpc (i.e. half the size of the periodic box). We obtained a sample of 100 such mock surveys. At each column density, show the 10th (blue), 50th (green, median), and 90th (red) percentiles of the distribution of absorber counts from the different mock surveys. To find the effect of large-scale structure on this distribution, we repeated this process for random surveys, show in lighter scale colors, where we select the same number of 6.25 cMpc columns, but choose them randomly. For these random surveys, we show the percentiles for ten different sets of 100 random surveys. The good agreement between the solid and dashed curves indicates that large-scale structure does not seem to affect the statistical variation in CDDFs that surveys may measure on scales of 6.25–50 cMpc.

3

The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray and UV line absorption

Nastasha A. Wijers, Joop Schaye, & Benjamin D. Oppenheimer.
2020, MNRAS, 498, 574–598

ABSTRACT

We use the EAGLE cosmological simulation to study the distribution of baryons, and FUV (O VI), EUV (Ne VIII) and X-ray (O VII, O VIII, Ne IX, and Fe XVII) line absorbers, around galaxies and haloes of mass $M_{200c} = 10^{11} - 10^{14.5} M_{\odot}$ at redshift 0.1. EAGLE predicts that the circumgalactic medium (CGM) contains more metals than the interstellar medium across halo masses. The ions we study here trace the warm-hot, volume-filling phase of the CGM, but are biased towards temperatures corresponding to the collisional ionization peak for each ion, and towards high metallicities. Gas well within the virial radius is mostly collisionally ionized, but around and beyond this radius, and for O VI, photo-ionization becomes significant. When presenting observables we work with column densities, but quantify their relation with equivalent widths by analysing virtual spectra. Virial-temperature collisional ionization equilibrium ion fractions are good predictors of column density trends with halo mass, but underestimate the diversity of ions in haloes. Halo gas dominates the highest column density absorption for X-ray lines, but lower-density gas contributes to strong UV absorption lines from O VI and Ne VIII. Of the O VII (O VIII) absorbers detectable in an Athena X-IFU blind survey, we find that 41 (56) per cent arise from haloes with $M_{200c} = 10^{12.0-13.5} M_{\odot}$. We predict that the X-IFU will detect O VII (O VIII) in 77 (46) per cent of the sightlines passing $M_{\star} = 10^{10.5-11.0} M_{\odot}$ galaxies within 100 pkpc (59 (82) per cent for $M_{\star} > 10^{11.0} M_{\odot}$). Hence, the X-IFU will probe covering fractions comparable to those detected with the Cosmic Origins Spectrograph for O VI.

3.1 Introduction

It is well-established that galaxies are surrounded by haloes of diffuse gas: the circumgalactic medium (CGM). Observationally, this gas has been studied mainly through rest-frame UV absorption by ions tracing cool ($\sim 10^4$ K) or warm-hot ($\sim 10^{5.5}$ K) gas (e.g., [Tumlinson et al. 2017](#) for a review). It has been found that the higher ions (mainly O VI) trace a different gas phase than the lower ions (e.g., H I), and that the CGM is therefore multi-phase. [Werk et al. \(2014\)](#) find that these phases and the central galaxy may add up to the cosmic baryon fraction around L_* galaxies, but the budget is highly uncertain, mainly due to uncertainties about the ionization conditions of the warm phase.

Theoretically, we expect hot, gaseous haloes to develop around $\sim L_*$ and more massive galaxies ($\log_{10} M_{200c} M_{\odot}^{-1} \gtrsim 11.5-12.0$; e.g., [Dekel & Birnboim 2006](#); [Kereš et al. 2009](#); [Van de Voort et al. 2011](#); [Correa et al. 2018](#)). The hot gas phase ($\gtrsim 10^6$ K) mainly emits and absorbs light in X-rays. For example, high-energy ions with X-ray lines dominate the haloes of simulated L_* galaxies (e.g. [Oppenheimer et al. 2016](#); [Nelson et al. 2018](#)). In observations, it is, however, still uncertain how much mass is in this hot phase of the CGM.

Similarly, there are theoretical uncertainties regarding the hot CGM. For example, we can compare the EAGLE ([Schaye et al. 2015](#)) and IllustrisTNG ([Pillepich et al. 2018](#)) cosmological simulations. They are both calibrated to produce realistic galaxies. However, they find very different (total) gas fractions in haloes with $M_{200c} \lesssim 10^{12.5} M_{\odot}$ ([Davies et al. 2020](#)), implying that the basic central galaxy properties used for these calibrations do not constrain those of the CGM sufficiently. This means that, while difficult, observations of the CGM hot phase are needed to constrain the models. The main differences here are driven by whether the feedback from star formation and black hole growth, which (self-)regulates the stellar and black hole properties in the central galaxy, ejects gas only from the central galaxy into the CGM (a galactic fountain), or ejects it from the CGM altogether, into the intergalactic medium ([Davies et al. 2020](#); [Mitchell et al. 2020](#)).

There are different ways to try to find this hot gas. The Sunyaev-Zel'dovich (SZ) effect traces the line-of-sight free electron pressure, and therefore hot, ionized gas. So far, it has been used to study clusters, and connecting filaments in stacked observations, as reviewed by [Mroczkowski et al. \(2019\)](#). Future instruments (e.g., CMB-S4, [Abazajian et al. 2016](#)) might be able to probe smaller angular scales with the SZ-effect, and thereby smaller/lower-mass systems.

Dispersion measures from fast radio bursts (FRBs) measure the total free electron column density along the line of sight, but are insensitive to the redshift of the absorption. They therefore probe ionized gas in general, but the origin of the electrons can be difficult to determine (e.g., [Prochaska & Zheng 2019](#)). [Ravi \(2019\)](#) found, using an analytical halo model, that it might be possible to constrain the ionized gas content of the CGM and intergalactic medium (IGM) using FRBs. This does require host galaxies for FRBs to be found in order to determine their redshift, uncertainties about absorption local to FRB environments to be reduced, and galaxy positions along the FRB sightline to be measured from (follow-up) surveys.

Another way to look for this hot phase is through X-ray emission. Unlike absorption or the SZ-effect, this scales with the density squared, and is therefore best suited for studying dense gas. However, if observed, it can give a more detailed image of a system than absorption along a single sightline. Emission around giant spirals, such as the very massive ($M_{\star} = 3 \times 10^{11} M_{\odot}$) isolated spiral galaxy NGC 1961, has been detected ([Anderson et al. 2016](#)). Around lower-mass spirals, such hot haloes have proven difficult to find: [Bogdán et al.](#)

(2015) stacked *Chandra* observations of eight $M_{\star} = 0.7\text{--}2 \times 10^{11} M_{\odot}$ spirals and found only upper limits on the X-ray surface brightness beyond the central galaxies. Anderson et al. (2013) stacked *ROSAT* images of a much larger set of galaxies (2165), and constrained the hot gas mass in the inner CGM.

In this work, we will focus on metal-line absorption. O VI absorption has been studied extensively using its FUV doublet at 1032, 1038Å at low redshift. It has been the focus of a number of observing programmes with the *Hubble Space Telescope's Cosmic Origins Spectrograph* (HST-COS) (e.g., Tumlinson et al. 2011; Johnson et al. 2015, 2017). A complication with O VI is that the implications of the observations depend on whether the gas is photo-ionised or collisionally ionised. This is often uncertain from observational data (e.g., Carswell et al. 2002; Tripp et al. 2008; Werk et al. 2014, 2016), and simulations find that both are present in the CGM (e.g., Tepper-García et al. 2011; Rahmati et al. 2016; Oppenheimer et al. 2016, 2018; Roca-Fàbrega et al. 2019). The uncertainty in the ionization mechanism leads to uncertainties in which gas phase is traced, and how much mass is in it.

The hot phase of the CGM, predicted by analytical arguments (the virial temperatures of haloes) and hydrodynamical simulations is difficult to probe in the FUV, since the hotter temperatures expected for $\sim L_{\star}$ galaxies' CGM imply higher-energy ions. One option, proposed by Tepper-García et al. (2013) and used by Burchett et al. (2019), is to use HST-COS to probe the CGM with Ne VIII (770, 780Å) at higher redshifts ($z > 0.5$). These lines in the extreme ultraviolet (EUV) cannot be observed at lower redshifts, so for nearby systems a different approach is needed.

Many of the lines that might probe the CGM hot phase have their strongest absorption lines in the X-ray regime (e.g., Perna & Loeb 1998; Hellsten et al. 1998; Chen et al. 2003; Cen & Fang 2006; Branchini et al. 2009). Some extragalactic O VII, O VIII and Ne IX X-ray line absorption has been found with current instruments, but with difficulty. Kovács et al. (2019) found O VII absorption by stacking X-ray observations centred on H I absorption systems near massive galaxies, though they targeted large-scale structure filaments rather than the CGM, while Ahoranta et al. (2020) found O VIII and Ne IX at the redshift of an O VI absorber. Bonamente et al. (2016) found likely O VIII absorption at the redshift of a broad Lyman α absorber. These tentative detections demonstrate that more certain, and possibly blind, extragalactic detections of these lines might be possible with more sensitive instruments.

The hot CGM of our own Milky Way galaxy can be observed more readily. Absorption from O VII has been found by e.g., Bregman & Lloyd-Davies (2007) and Gupta et al. (2012, also O VIII), and Hodges-Kluck et al. (2016) studied the velocities of O VII absorbers. Gatuza & Churazov (2018) studied Ne IX absorption alongside O VII and O VIII, focussing on the hot CGM and the ISM. The Milky Way CGM has also been probed with soft X-ray emission (e.g., Kuntz & Snowden 2000; Miller & Bregman 2015; Das et al. 2019), and studied using combinations of emission and absorption (e.g., Bregman & Lloyd-Davies 2007; Gupta et al. 2014; Miller & Bregman 2015; Gupta et al. 2017; Das et al. 2019).

Previous theoretical studies of CGM X-ray absorption include analytical modelling, which tends to focus on the Milky Way. For example, Voit (2019), used a precipitation-limited model to predict absorption by O VI– VIII, N V, and Ne VIII, and Stern et al. (2019) compared predictions of their cooling flow model to O VII and O VIII absorption around the Milky Way. Faerman et al. (2017) constructed a phenomenological CGM model, based on O VI– VIII absorption and O VII and O VIII emission in the Milky Way. Nelson et al. (2018) studied O VII and O VIII in IllustrisTNG, but focused on a wider range of halo masses: two orders of magnitude in halo mass around L_{\star} .

In [Wijers et al. \(2019\)](#) we used the EAGLE hydrodynamical simulation to predict the cosmic distribution of O VII and O VIII for blind observational surveys. We found that absorbers with column densities $N_{\text{O VII, O VIII}} \gtrsim 10^{16} \text{ cm}^{-2}$ typically have gas overdensities $\gtrsim 10^2$, and that absorbers with overdensities ~ 10 may be difficult to detect at all in planned surveys. Therefore, we expect that a large fraction of the X-ray absorbers detectable with the planned Athena X-IFU ([Barret et al. 2016](#)) survey, and proposed missions such as Arcus ([Brenneman et al. 2016](#); [Smith et al. 2016](#)), are associated with the CGM of galaxies. Until such missions are launched, progress can be made with deep follow-up of FUV absorption lines with current X-ray instruments. The simulations can also help interpret the small number of absorbers found with current instruments (e.g., [Nicastrò et al. 2018](#); [Kovács et al. 2019](#); [Ahoranta et al. 2020](#)); e.g. [Johnson et al. \(2019\)](#) used galaxy information to re-interpret the lines found by [Nicastrò et al. \(2018\)](#).

In this work, we will consider O VI (1032, 1038Å FUV doublet), Ne VIII (770, 780Å EUV doublet), O VII (He- α resonance line at 21.60Å), O VIII (18.9671, 18.9725Å doublet), Ne IX (13.45Å), and Fe XVII (15.01, 15.26Å). In collisional ionization equilibrium (CIE), the limiting ionisation case for high-density gas, these ions probe gas at temperatures $T \sim 10^{5.5} - 10^7 \text{ K}$, covering the virial temperatures of $\sim L_*$ haloes to smaller galaxy clusters (see Fig. [3.1](#) and Table [3.3](#)), as well as the ‘missing baryons’ temperature range in the warm-hot intergalactic medium (e.g., [Cen & Ostriker 1999](#)). We include O VI because this highly ionized UV ion has proved useful in *HST-COS* studies, and Ne VIII has been used to probe a hotter gas phase, albeit at higher redshifts. O VII, O VIII, and Ne IX are strong soft X-ray lines, probing our target gas temperature range, and have proven to be detectable in X-ray absorption. Fe XVII is expected to be a relatively strong line at higher energies ([Hellsten et al. 1998](#)), probing the hottest temperatures in the missing baryons range (close to 10^7 K), and is therefore also included.

We will predict UV and X-ray column densities in the CGM of EAGLE galaxies at $z = 0.1$, and explore the physical properties of the gas the various ions probe. We also investigate which haloes we are most likely to detect with the Athena X-IFU. In [§4.2](#) we discuss the EAGLE simulations and the methods we use for post-processing them. In [§4.4](#), we will discuss our results. We start with a general overview of the ions and their absorption in [§3.3.1](#), then discuss the baryon, metal, and ion contents of EAGLE haloes in [§3.3.2](#). Then, we discuss what fraction of absorption systems of different column densities are due to the CGM ([§3.3.3](#)) and how those column densities translate into equivalent widths, which are more directly observable. We then switch to a galaxy-centric perspective and show absorption profiles for galaxies of different masses ([§3.3.4](#)), and what the underlying spherical gas and ion distributions are ([§3.3.5](#)). In [§3.4](#) we use those absorption profiles and the relations we found between column density and equivalent width to predict what can be observed. In [§4.5](#) we discuss our results in the light of previous work, and we summarize our results in [§4.6](#).

Throughout this paper, we will use L_* for the characteristic luminosity in the present-day galaxy luminosity function ($\sim 10^{12} M_\odot$ haloes), and M_* for the stellar masses of galaxies. Except for centimetres, which are always a physical unit, we will prefix length units with ‘c’ if they are comoving and ‘p’ if they are proper/physical sizes.

3.2 Methods

In this section, we will discuss the cosmological simulations we use to make our predictions ([§4.2.1](#)), the galaxy and halo information we use ([§3.2.2](#)), and how we define the CGM ([§3.2.3](#)).

We explain how we predict column densities (§3.2.4 and §3.2.5), equivalent widths (§3.2.6), and absorption profiles (§3.2.7) from these simulations.

3.2.1 EAGLE

The EAGLE simulations (‘Evolution and Assembly of GaLaxies and their Environments’; Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016) are cosmological, hydrodynamical simulations. Gravitational forces are calculated with the GADGET-3 Tree-PM scheme (Springel 2005) and hydrodynamics is implemented using a smoothed particle hydrodynamics (SPH) method known as ANARCHY (Schaye et al. 2015, appendix A; Schaller et al. 2015). EAGLE uses a Λ CDM cosmogony with the Planck Collaboration et al. (2014) cosmological parameters: $(\Omega_m, \Omega_\Lambda, \Omega_b, h, \sigma_8, n_s, Y) = (0.307, 0.693, 0.04825, 0.6777, 0.8288, 0.9611, 0.248)$, which we also adopt in this work.

Here, we use the 100^3 cMpc^3 EAGLE simulation, though we made some comparisons to both smaller-volume and higher-resolution simulations to check convergence. It has a dark matter particle mass of $9.70 \times 10^6 M_\odot$, an initial gas particle mass of $1.81 \times 10^6 M_\odot$, and a Plummer-equivalent gravitational softening length of 0.70 pkpc at the low redshifts we study here.

The resolved effects of a number of unresolved processes (‘subgrid physics’) are modelled in order to study galaxy formation. This includes star formation, black hole growth, and the feedback those cause, as well as radiative cooling and heating of the gas, including metal line cooling (Wiersma et al. 2009a). To prevent artificial fragmentation of cool, dense gas, a pressure floor is implemented at ISM densities.

In EAGLE, stars form in dense gas, with a pressure-dependent star formation rate designed to reproduce the Kennicutt-Schmidt relation. They return metals to surrounding gas based on the yield tables of Wiersma et al. (2009b) and provide feedback from supernova explosions by stochastically heating gas to $10^{7.5}$ K, with a probability set by the expected energy produced by supernovae from those stars (Dalla Vecchia & Schaye 2012). Black holes are seeded in low-mass haloes and grow by accreting nearby gas (Rosas-Guevara et al. 2015). They provide feedback by stochastic heating as well (Booth & Schaye 2009), but to $10^{8.5}$ K. This stochastic heating is used to prevent overcooling due to the limited resolution: if the expected energy injection from single supernova explosions is injected into surrounding dense $\sim 10^6 M_\odot$ gas particles at each time step, the temperature change is small, cooling times remain short, and the energy is radiated away before it can do any work. This means self-regulation of star formation in galaxies fails, and galaxies become too massive. The star formation and stellar and black hole feedback are calibrated to reproduce the $z = 0.1$ galaxy luminosity function, the black-hole-mass-stellar-mass relation, and reasonable galaxy sizes (Crain et al. 2015).

3.2.2 Galaxies and haloes in the EAGLE simulation

We use galaxy and halo information from EAGLE in two ways. First, we look at the properties of gas around haloes. We obtain absorption profiles (column densities as a function of impact parameter), as well as spherically-averaged gas properties as a function of (3D) distance to the central galaxy. Second, we investigate what fraction of absorption in a random line of sight with a particular column density is, on average, due to haloes (of different masses), to help interpret what might be found in a blind survey for line absorption.

We use the EAGLE galaxy and halo catalogues, which were publicly released as documented by McAlpine et al. (2016). The haloes are identified using the Friends-of-Friends

Table 3.1: The halo sample size from EAGLE L0100N1504 at $z = 0.1$, with the total number of haloes (equal to the number used for the 2D radial profiles), the number outside R_{200c} of any 6.25 cMpc slice edge, and the number used for 3D radial profiles.

M_{200c} $\log_{10} M_{\odot}$	total	off edges	3D profiles
11.0–11.5	6295	6044	1000
11.5–12.0	2287	2159	1000
12.0–12.5	870	792	870
12.5–13.0	323	288	323
13.0–13.5	119	103	119
13.5–14.0	26	20	26
≥ 14.0	9	8	9

(FoF) method (Davis et al. 1985), which connects dark matter particles that are close together (within 0.2 times the mean inter-particle separation, in this case), forming haloes defined roughly by a constant outer density. Other simulation particles (gas, stars, and black holes) are linked to a FoF halo if their closest dark matter particle is. Within these haloes, galaxies are then identified as subhaloes recovered by SUBFIND (Springel et al. 2001; Dolag et al. 2009), which identifies self-bound overdense regions within the FoF haloes. The central galaxy is the subhalo containing the particle with the lowest gravitational potential.

Though SUBFIND and the FoF halo finder are used to identify structures, we do not characterise haloes using their masses directly. Instead, we use M_{200c} , for halo masses, which is calculated by growing a sphere around the FoF halo potential minimum (central galaxy) until the enclosed density is the target $200\rho_c$, where $\rho_c = 3H(z)^2 (8\pi G)^{-1}$ is the critical density, and $H(z)$ is the Hubble factor at redshift z . For stellar masses, we use the stellar mass enclosed in a sphere with a 30 pkpc radius around each galaxy’s lowest-gravitational-potential particle. We use centres of mass for the positions of galaxies, and the centre of mass of the central galaxy for the halo position.

Since the temperature of the gas is important in determining its ionization state, we also want an estimate of the temperature of gas in haloes of different masses. For this, we use the virial temperature

$$T_{200c} = \frac{\mu m_H}{3k} G M_{200c}^{2/3} (200\rho_c)^{1/3}, \quad (3.1)$$

where m_H is the hydrogen mass, G Newton’s constant, and k the Boltzmann constant. We use a mean molecular weight $\mu = 0.59$, which is appropriate for primordial gas, with both hydrogen and helium fully ionized.

We will look into the properties of haloes mostly as a function of M_{200c} . For this, we use halo mass bins 0.5 dex wide, starting at $10^{11} M_{\odot}$. Table 3.1 shows the sample size this yields for different halo masses. There is a halo with a mass $> 10^{14.5} M_{\odot}$, but we mostly choose not to include a separate bin for this single $10^{14.53} M_{\odot}$ halo, and group all haloes with $M_{200c} > 10^{14} M_{\odot}$ together instead. The second column shows the total number of haloes in the 100^3 cMpc³ volume we use, and the third column shows the number of haloes that are not ‘cut in pieces’ by the box slicing method we use to obtain column densities (§3.2.5). The sample size in the second column is used when calculating absorption as a function of impact parameter. However, to reduce calculation times, we use a subsample of 1000 randomly chosen haloes when we calculate total baryon and ion masses in the CGM, and gas properties as a function of (3D) radius. This is shown in the fourth column.

3.2.3 CGM definitions

Roughly speaking, the CGM is the gas surrounding a central galaxy, in a region similar to that of the dark-matter halo containing the galaxy. This definition is not very precise, because there is no clear physical boundary between the CGM and IGM or between the CGM and ISM. We will make use of a few different definitions. Here, we discuss how to identify individual SPH particles as part of the CGM. In §3.2.7 we discuss two methods for identifying (line-of-sight-integrated) absorption due to haloes. We mention the used definition in each figure caption, but summarize the definitions here.

The simplest approach we take is to ignore any explicit halo membership and just consider all gas as a function of distance to halo centres. We use this method for column densities and covering fractions as a function of impact parameter (though we do limit what is included along the line of sight; see §3.2.5), and for the temperature, density, and metallicity profiles we calculate. This is what we use in Fig. 3.5 the solid, black lines in Fig. 3.6 the solid lines in Fig. 3.8 Figs. 3.10, 3.11, 3.12, 3.13, 3.14 and 3.17 and the black lines in Fig. 3.18.

The first CGM definition we use is based on the friends-of-friends groups we discussed in §3.2.2. Here, we define the CGM as all gas in the FoF group defining a halo, as well as any other gas within the R_{200c} sphere of that halo. We use this definition when we want to identify all gas within a set of haloes (the haloes in different mass bins), because for each EAGLE gas particle, a halo identifier following this definition is stored (The EAGLE team 2017). We use this in Fig. 3.2 and in the halo-projection method discussed in §3.2.7 used in the brown and rainbow-coloured lines in Figs. 3.6 and 3.18 and the dashed lines in Fig. 3.8. This method is also one of the options explored in Fig. 3.16 (see also §3.2.7 and Appendix 3.B).

In §3.3.2, we also describe the composition of haloes using other CGM definitions. For Figs. 3.3 and 3.4, we define all gas within R_{200c} of the halo centre as part of the halo. When we split the gas mass into CGM and ISM in Fig. 3.3, we define the ISM to be all star-forming gas and the CGM to be all other gas inside the halo. In Fig. 3.4, we explore the ion content of the halo. Here, we roughly excise the central galaxy by excluding gas within $0.1 R_{200c}$ of the halo centre. However, we explore some variations of these definitions.

3.2.4 The ions considered in this work

We consider six different ions in this work: O VI, O VII, O VIII, Ne VIII, Ne IX, and Fe XVII. We list the atomic data we use for the absorption lines of these ions in Table 4.1. To calculate the fraction of each element in an ionisation state of interest, we use tables giving these fractions as a function of temperature, density, and redshift. These are the tables of Bertone et al. (2010a,b). The density- and redshift-dependence comes from the assumed uniform, but redshift-dependent Haardt & Madau (2001) UV/X-ray background. The tables were generated using using CLOUDY (Ferland et al. 1998), version c07.02.00. This is consistent with the radiative cooling and heating used in the EAGLE simulations (Wiersma et al. 2009a).

Unfortunately, this main set of tables we use does not include all the ionization states of oxygen, and we want to examine the overall partition of oxygen ions in haloes. Therefore, we also use a second set of tables, though only for the oxygen ions in Fig. 3.4. This second set of tables was made under the same assumptions as our main set: the uniform but time-dependent UV/X-ray background (Haardt & Madau 2001) used for the EAGLE cooling tables, assuming optically thin gas in ionization equilibrium. However, they were generated using a newer CLOUDY version: 13 (Ferland et al. 2013). We checked by comparing the tables and a smaller EAGLE simulation that the differences between these tables are small for O VI– VIII. In a part of a smaller EAGLE volume, and in the column density regimes of interest, the O VI

column densities differed by $\lesssim 0.1$ dex. The O VII and O VIII column densities differed even less. The tables differ most clearly in the photo-ionized regime, where the column densities are small.

3.2.5 Column densities from the simulated data

Using these ion fractions, we calculate column densities in the same way as in [Wijers et al. \(2019\)](#). In short, we use the ion fraction tables we described in [§3.2.4](#) which we linearly interpolate in redshift, log density, and log temperature to get each SPH particle's ion fraction. We multiply this by the tracked element abundance and mass of each SPH particle to calculate the number of ions in each particle.

We then make a 2-dimensional column density map from this ion distribution. Given an axis to project along and a region of the simulation volume to project, we calculate the number of ions in long, thin columns parallel to the projection axis. We then divide by the area of the columns perpendicular to the projection axis to get the column density in each pixel of a 2-dimensional map. In order to divide the ions in each SPH particle over the columns, we need to assume a spatial ion distribution for each particle. For this, we use the same C2-kernel used for the hydrodynamics in the EAGLE simulations ([Wendland 1995](#)), although we only input the 2-dimensional distance to each pixel centre.

A simple statistic that can be obtained from these maps is the column density distribution function (CDDF). This is a probability density function for absorption system column density, normalized to the comoving volume probed along a line of sight. The CDDF is defined by

$$f(N, z) = \frac{\partial^2 n}{\partial \log_{10} N \partial X}, \quad (3.2)$$

where N is the column density, n the number of absorbers, z the redshift, and X the absorption length given by

$$dX = (1 + z)^2 (H(0) / H(z)) dz, \quad (3.3)$$

where $H(z)$ is the Hubble parameter.

In practice, we make column density maps along the Z-axis of the simulation box, which is a random direction for haloes. We use 32000^2 pixels of size 3.125^2 cMpc² for the column density maps, and 16 slices along the line of sight, which means the slices are 6.25 cMpc thick.

[Wijers et al. \(2019\)](#) found that this produces converged results for O VII and O VIII CDDFs up to column densities $N \approx 10^{16.5} \text{ cm}^{-2}$. Here we mean converged with respect to pixel size, simulation size, and simulation resolution. By default, we set the temperature of star-forming gas to be 10^4 K, since the equation of state for this high-density gas does not reflect the temperatures we expect from the ISM. However, this has negligible impacts on the column densities of O VII and O VIII. Note that all our results do neglect a hot ISM phase, which is not modelled in EAGLE, but may affect column densities in observations for very small impact parameters.

[Rahmati et al. \(2016\)](#) used EAGLE to study UV ion CDDFs and tested convergences for O VI and Ne VIII. They used the same slice thickness at low redshift, but a lower map resolution: 10000^2 pixels. At that resolution, they find O VI CDDFs are converged to $N \approx 10^{15} \text{ cm}^{-2}$, and Ne VIII to $N \approx 10^{14.5} \text{ cm}^{-2}$. The volume and resolution of the simulation do affect CDDFs down to lower column densities. For O VI, resolution has effects down to $N \approx 10^{14} \text{ cm}^{-2}$.

We checked the convergence of Ne IX and Fe XVII CDDFs with slice thickness, pixel size, box size, and box resolution in the same way as [Wijers et al. \(2019\)](#). We found that Ne IX

Table 3.2: Atomic data for the absorption lines we study. For each ion, we record the wavelengths λ , oscillator strengths f_{osc} , and transition probabilities A we use to calculate the equivalent widths in Fig. 3.7. For resolved doublets, we only use the stronger line. The last column indicates the source of the line data: M03 for Morton (2003), V96 for Verner et al. (1996), and K18 for Kaastra (2018).

ion	λ Å	f_{osc}	A s^{-1}	source
O VI	1031.9261	0.1325	4.17×10^8	M03
Ne VIII	770.409	0.103	5.79×10^8	V96
O VII	21.6019	0.696	3.32×10^{12}	V96/K18
O VIII	18.9671	0.277	2.57×10^{12}	V96
	18.9725	0.139	2.58×10^{12}	V96
Ne IX	13.4471	0.724	8.90×10^{12}	V96/K18
Fe XVII	15.0140	2.72	2.70×10^{13}	K18

column densities are converged up to $N \approx 10^{16} \text{ cm}^{-2}$, with $\lesssim 20$ per cent changes in the CDDF at $N \gtrsim 10^{12} \text{ cm}^{-2}$ due to factor 2 changes in slice thickness. For Fe XVII, CDDFs are converged to $N \approx 10^{15.4} \text{ cm}^{-2}$, with mostly smaller dependences on slice thickness than the other X-ray ions. (We will later see that this ion tends to be more concentrated within haloes, so on smaller scales, than the others we investigate.) The trends of effect size with column density, and the relative effect sizes of changing pixel size, slice thickness, simulation volume, and simulation resolution on the CDDFs, are similar to those for O VII and O VIII. We note that the resolution test for Fe XVII may not be reliable, since at larger column densities, this ion is largely found in high-mass haloes which are very rare or entirely absent in the smaller volume (25^3 cMpc^3) used for this test.

3.2.6 Equivalent widths from the simulated data

In observations, column densities are not directly observable. Instead, they must be inferred from absorption spectra. The equivalent width (EW) can be calculated from the spectrum more directly, and for X-ray absorption, determines whether a line is observable. (Line widths can play a role, but for the Athena X-IFU, those will be below the spectral resolution of the instrument in all cases, as we will later show.)

We compute the EWs in mostly the same way as Wijers et al. (2019), using `specwiz` (e.g., Tepper-García et al. 2011, §3.1). Briefly, in Wijers et al. (2019) we extracted absorption spectra along 100 cMpc sightlines through the full EAGLE simulation box, then calculated the EW for the whole sightline, and compared that to the total column density calculated in the same code.

In `SPECWIZARD`, sightlines are divided into pixels (1-dimensional), and ion densities, ion-density-weighted peculiar velocities and ion-density-weighted temperatures are calculated in those pixels. The spectrum is then calculated by adding up the optical depth contributions from the position-space pixels in each spectral pixel. The optical depth profile used for each position-space pixel is Gaussian, with the centre determined by the pixel position and peculiar velocity, the width by the temperature (thermal line broadening only), and the normalization by the column density. Since, in reality, spectral lines are better described as Voigt profiles, a convolution of a Gaussian with a Cauchy-Lorentz profile, we convolve the (Gaussian-line) spectra from `SPECWIZARD` with the appropriate Cauchy-Lorentz profile for each spectral line, using the transition probabilities from Table 4.1.

Comparing EWs calculated over the full sightlines with and without the additional line broadening (eq. 3.5), we find that for O VI and Ne VIII, the differences are < 0.01 dex everywhere. For the X-ray ions, the vast majority of sightlines show differences < 0.1 dex, with larger differences occurring in $\lesssim 10$ sightlines at the highest column densities. The differences are largest for Fe XVII.

In this work, we do not measure column densities and EWs along full 100 cMpc sightlines. Instead, we use velocity windows around the line-of-sight velocity where the optical depth is largest. We calculate EWs in these velocity ranges by integrating the synthetic spectra over that velocity range. For the column densities in those windows, we use the fact that the total optical depth is proportional to the column density. Therefore, the fraction of the total column density in each velocity window is the same as the fraction of the total (integrated) optical depth contained within the window.

Note that we do not necessarily use all absorption systems in the sightline. This may bias our results, but so does using full sightline values. Identifying and fitting individual absorbers and absorption systems is beyond the scope of this paper. In Appendix 3.A we show that our results are insensitive to the precise choice of velocity window.

For the UV ions, we mimic velocity windows used to define absorption systems by observers: $\pm 300 \text{ km s}^{-1}$ (rest-frame). This matches how Burchett et al. (2019) defined absorption systems in their CASBaH study of Ne VIII. For O VI, Johnson et al. (2015) searched $\Delta v = \pm 300 \text{ km s}^{-1}$ regions around galaxy redshifts for the eCGM survey. Tumlinson et al. (2011) searched a larger region of $\Delta v = \pm 600 \text{ km s}^{-1}$ in the COS-Halos survey, but found that the absorbers were strongly clustered within $\Delta v = \pm 200 \text{ km s}^{-1}$.

For the X-ray lines, we want to use velocity windows resolvable by the Athena X-IFU: the FWHM resolution should be 2.5 eV (Barret et al. 2018). This corresponds to different velocity windows for the different lines (at different energies) we consider: $\approx 1200 \text{ km s}^{-1}$ for O VII, $\approx 1000 \text{ km s}^{-1}$ for O VIII, $\approx 800 \text{ km s}^{-1}$ for Fe XVII, and $\approx 700 \text{ km s}^{-1}$ for Ne IX at $z = 0.1$. Based on the dependence of the best-fitting b -parameters on the velocity ranges, we choose to use a half-width $\Delta v = \pm 800 \text{ km s}^{-1}$ for the X-ray ions. We discuss this choice in Appendix 3.A

We started with the sample of spectra for the sightlines used in Wijers et al. (2019) for $z = 0.1$. This sample was a combination of three subsamples, selected to have high column density in O VI, O VII, or O VIII. Subsamples were selected uniformly in log column density for $N \geq 10^{13} \text{ cm}^{-2}$ in each ion, iterating the selection until the desired total sample size of 16384 sightlines was reached. For this work, we added a sample of the same size, but with subsamples selected by Ne VIII, Ne IX, and Fe XVII column density. Some sightlines in the two samples overlapped, giving us a total sample of 31706 sightlines. For each ion, we only use the sightlines selected for that ion specifically. These subsamples contain ≈ 5600 sightlines each.

Table 4.1 lists the wavelengths, oscillator strengths, and transition probabilities we use for the ions. If an ion absorption line is actually a close doublet (expected to be unresolved), we calculate the equivalent widths from the total spectrum of the doublet lines. This is only the case for O VIII (e.g. fig. 4 of Wijers et al. 2019). For Fe XVII, the 15.26, 15.02Å doublet has a rest-frame velocity difference of $4.75 \times 10^3 \text{ km s}^{-1}$. This is well above the line widths we find, so the lines will not generally be intrinsically blended, and should be resolvable by the Chandra LETG¹ and the XMM-Newton RGS (den Herder et al. 2001, fig. 11). The Athena X-IFU will have a higher resolution (Barret et al. 2018). We only use the stronger component for the O VI 1031.9, 1037.6Å and Ne VIII 770.4, 780.3Å doublets, which are easily resolved

¹http://cxc.harvard.edu/cdo/about_cchandra/

with current FUV spectrographs.

We note that for Fe xvii, the atomic data for the line are under debate, with theoretical calculations and experiments finding different values (e.g., Gu et al. 2007; de Plaa et al. 2012; Bernitt et al. 2012; Wu & Gao 2019; Gu et al. 2019). Indeed the Kaastra (2018) wavelength and oscillator strength that we use for this ion do not agree with the Verner et al. (1996) values. The wavelengths only differ by 0.001Å (a relative difference of 0.007 per cent), but the oscillator strengths and transition probabilities differ by 8 per cent.

We will use these spectra to infer the relation between the more directly observable EWs, and the more physically interesting column densities we use throughout the paper. We parametrise this relation using the relation between column density and EW for a single absorber (so-called ‘curves of growth’), using line widths b . These relations are for a single Voigt profile (or doublet of Voigt profiles). They consist of a Gaussian absorption line convolved with a Cauchy-Lorentz profile. The line is described by a continuum-normalized spectrum $\exp(-\tau(\Delta v))$, where Δv is the velocity offset from the line centre and τ is the optical depth. The Gaussian part of the optical depth profiles is described by

$$\tau(\Delta v) \propto N b^{-1} \exp(-(\Delta v b^{-1})^2), \quad (3.4)$$

where N is the column density of the ion. The constant of proportionality is governed by the atomic physics of the transition in question. For such a line, FWHM = 1.67 b . However, the line is additionally broadened by the Cauchy-Lorentz component

$$f(v) = \frac{1}{4\pi^2} \frac{A}{(\Delta v)^2 + (A/4\pi)^2}, \quad (3.5)$$

where Δv is the frequency offset and A is the transition probability. When we fit b parameters, we model the Voigt profile of the lines (convolution of eqs. 3.4 and 3.5), and b refers to the width of the Gaussian component (eq. 3.4) alone.

We will fit these b parameters to the column densities and EWs measured along the different sightlines for the different ions, by minimizing

$$\sum_i (\log_{10} \text{EW}_i - \log_{10} \text{EW}(N_i, b))^2, \quad (3.6)$$

where the sum is over the sightlines, N is the column density, and $\text{EW}(N, b)$ is obtained by integrating the spectrum produced by the Voigt profile in eqs. 3.4 and 3.5. Fitting the EWs themselves instead of the log EWs makes little difference: only a few km s^{-1} . Using the velocity windows instead of the full sightlines only makes a substantial difference for O viii. We discuss the dependence of the best-fitting b values on the velocity range used in Appendix 3.A

Note that the indicative b -parameters we find here from the curve of growth should not be directly compared with observed values: in UV observations, line widths are often measured by fitting Voigt profiles to individual absorption components, instead of inferred from theoretically known column densities and EWs of whole absorption systems as we do here.

3.2.7 Absorption profiles

We extract absorption profiles around galaxies from the 2-dimensional maps described in §3.2.5. We extract profiles from both full maps and from maps created using only gas in

haloes in particular mass ranges (i.e., gas in the FoF groups or R_{200c} regions of these haloes; see §3.2.3). Given the positions of the galaxies, we obtain radial profiles by extracting column densities and distances from pixel centres to galaxy centres, then binning column densities by distance.

We use only 2-dimensional distances (impact parameters) here, but only use the column density map for the Z-coordinate range that includes the galaxy centre. We compared this method to two variations for obtaining radial profiles (not shown): adding up column densities from the two slices closest to the halo centre, and using only galaxies at least R_{200c} away from slice edges for radial profiles. We found that this made little difference for the median column densities: profiles excluding haloes close to slice edges were indistinguishable from those using all haloes, in part because the excluded haloes were only a small part of the sample (Table 3.1). The exceptions were the most massive haloes ($M_{200c} > 10^{13.5} M_{\odot}$), where larger haloes and small sample sizes mean the effect on the sample is larger. Even there, differences remained $\lesssim 0.2$ dex. Using two slices instead of one made a significant difference only where both predicted median column densities were well below observable limits we consider, and well below the highest halo column densities we find.

To obtain the contributions of different haloes to the CDDF, we use two approaches. In the first, which we call the *halo-projection method*, we make CDDFs by counting ions in long, thin, columns as for the total CDDFs, but we only use particles that are part of a halo's FoF group, or inside its R_{200c} sphere. Alternatively, we make maps describing which pixels in the full column density maps belong to which haloes, if any, by checking if a pixel is within R_{200c} of a halo (in projected distance r_{\perp}): the *pixel-attribution method*. To do this, we make 2d maps of the same regions, and at the same resolution, as the column density maps. These are simple True/False maps, and we make them for every set of haloes we consider. However, the map does not include any pixel that is closer, in units of R_{200c} , to a halo from a different mass-defined set. We compare these methods for splitting up the CDDFs in Appendix 3.B. Typically, the results are similar for larger column densities, but the halo-projection CDDFs contain more small column density values, coming largely from sightlines probing only short paths through the edges of the haloes.

The advantage of using the pixel-attribution method is that it is more comparable to observations, where large-scale structure around haloes will also be present. (Note, however, that we neglect peculiar velocities.) For the CDDFs, it also allows us to attribute specific pixels in the maps to a halo or the IGM, meaning we can truly split up the CDDF into different contributions. A downside is that some haloes will be close to an edge of the projected slice, meaning that absorption due to a halo in one slice will be missed, while that of another is underestimated. However, the fraction of such haloes is small (Table 3.1). On the other hand, absorption may also be attributed to haloes that just happen to be close (in projection) to the absorber. This is mainly an issue for lower-mass haloes. We also explore this effect Appendix 3.B.

3.3 Results

First, we investigate some of the simplest data on our ions: what temperatures and densities they exist at (§3.3.1). We then discuss the contents of haloes (§3.3.2). Then we discuss the contributions of different haloes to the ion CDDFs, and the relation between column densities and EWs (§3.3.3), absorption around haloes as a function of impact parameter (§3.3.4), and the 3D ion distribution around galaxies (§3.3.5). For predictions that should be comparable to observations, we refer the reader to §3.4. These results are for $z = 0.1$. In appendix 3.C.

Table 3.3: Data for the ions we study. E_{ion} is the energy needed to remove the least bound electron from each ion, and T_{CIE} is the preferred CIE temperature of the ions. The CIE ranges are the upper and lower temperatures at which the ion fraction is 10 per cent of the CIE maximum. Ionization energies are from [Lide \(2003\)](#).

ion	E_{ion} eV	T_{CIE} \log_{10} K
O VI	138.12	5.3–5.8
Ne VIII	239.10	5.6–6.1
O VII	739.29	5.4–6.5
O VIII	871.41	6.1–6.8
Ne IX	1195.83	5.7–6.8
Fe XVII	1266	6.3–7.0

we compare some results to those for $z = 0.5$.

3.3.1 Ion properties

First, we will examine at which densities and temperatures the ions we investigate exist in meaningful quantities, which can be used to make a simple estimate of which ions are most prominent in which haloes. [Table 3.3](#) and [Fig. 3.1](#) show the energies and temperatures associated with each ion. [Fig. 3.1](#) visualizes the [Bertone et al. \(2010a,b\)](#) ionization tables we use throughout the paper.

The shaded regions for each ion in [Fig. 3.1](#) show the temperatures and densities where the ion fraction is at least 0.1 times the maximum fraction in CIE. The temperature range this corresponds to in CIE is given in [Table 3.3](#).

[Fig. 3.1](#) shows two regimes for each ion. The first is the high-density regime where ionization by the UV/X-ray background is negligible compared to ionization by electron-ion collisions. Since recombinations and ionizations both increase as n_{H}^2 in this regime, ion fractions are only dependent on temperature here. Since we assume ionization equilibrium, this is the collisional ionization equilibrium (CIE) regime. The second is the low-density regime where ionization by the UV/X-ray background dominates, and the density of the gas becomes important. This is the photo-ionization equilibrium (PIE) regime. The transition between these regimes occurs at $n_{\text{H}} \sim 10^{-5} \text{ cm}^{-3}$.

The long, coloured tick marks on the right axis indicate the temperature where each ion’s fraction is largest in CIE, and the right axis shows the halo mass with T_{200c} (eq. [4.1](#)) corresponding to the temperature on the left axis. Since the densities in the CGM are typically $n_{\text{H}} \gtrsim 10^{-5} \text{ cm}^{-3}$ (see [§3.3.5](#)), comparing the halo masses on the right axis to the temperatures where the ion fractions are high in CIE gives a reasonable estimate of which haloes contain the highest masses of the different ions, and have the highest column densities of those ions (as shown later in [Figs. 3.2](#) and [3.8](#)).

3.3.2 The baryonic content of haloes

Next, we look into how the ions relate to haloes in EAGLE. [Fig. 3.2](#) shows the contributions of haloes of different masses to the total mass and ion budget in the simulated 100^3 cMpc^3 . An SPH particle is considered part of a halo if it is within the halo’s FoF group or R_{200c} region. We include the 14.5–15 bin for consistent spacing, but this bin contains only a single halo

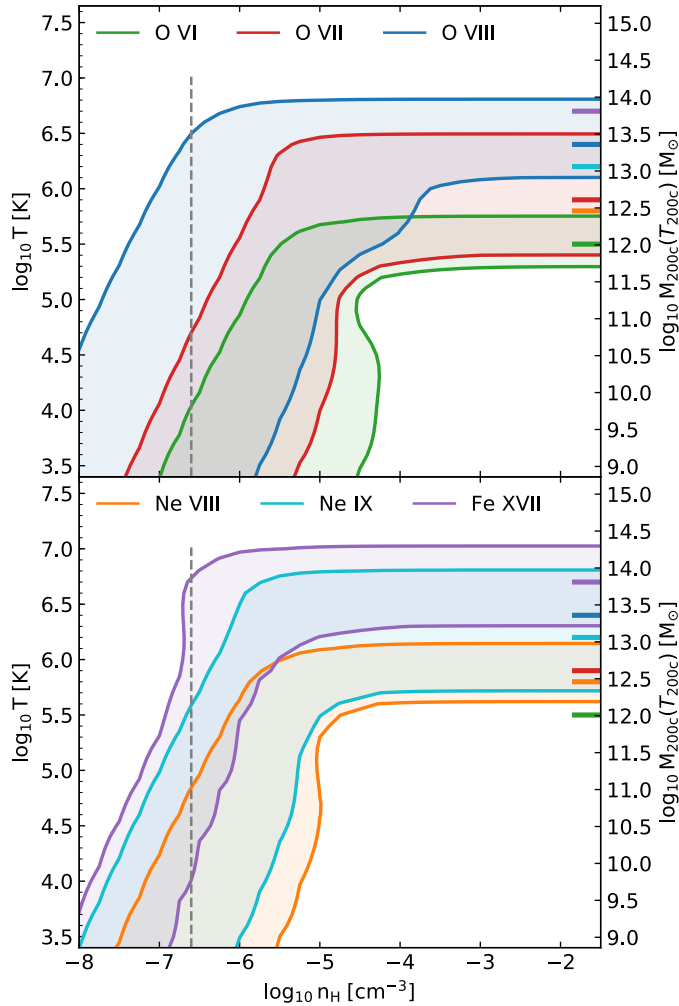


Figure 3.1: The temperatures and densities where different ions occur at $z = 0.1$, assuming a [Haardt & Madau \(2001\)](#) UV/X-ray background as the only photo-ionizing source. The contours for each of the indicated ions are at 10 per cent of the maximum ion fraction in CIE. The vertical, dashed line indicates the cosmic average baryon density. The right axis indicates the halo masses with virial temperatures (eq. 4.1) matching the temperatures on the y-axis, and the coloured ticks indicate where each ion's fraction peaks in CIE.




Figure 3.2: The fraction of total gas mass and the (gas-phase) elements and ions we investigate contributed by haloes of different masses at $z = 0.1$ in the EAGLE simulation. Colours indicate halo masses according to the colour bar, with grey indicating gas that does not belong to any halo. Gas is considered part of a halo if it is part of its FoF group or R_{200c} sphere. Neon and iron (not shown) are distributed similarly to oxygen.

with $M_{200c} = 10^{14.53} M_{\odot}$, so in rest of the paper, we will group all nine haloes with masses $M_{200c} \geq 10^{14} M_{\odot}$ into one halo mass bin.

Fig. 3.2 shows the ions inside haloes are mostly found at halo masses where $T_{200c} \sim T_{\text{CIE}}$. The differences between ions, and between the ion, metal, and mass distributions show that these trends are not simply a result of the ions tracing mass or metals. The importance of haloes with $T_{200c} \sim T_{\text{CIE}}$ can be explained by a few factors. First, the temperature of the warm/hot gas in haloes is roughly T_{200c} . Second, in haloes, the ions are mostly found in whatever gas there is at $\sim T_{\text{CIE}}$. This is because, third, the density of the warm/hot phase is mostly high enough that the gas is collisionally ionized. (In lower-mass haloes, with $M_{200c} \lesssim 10^{12} M_{\odot}$, and/or gas at $\sim R_{200c}$, photo-ionization does become relevant.) This means that haloes with $T_{200c} \sim T_{\text{CIE}}$ contain larger amounts of ion-bearing gas than haloes at higher or lower temperatures (masses). We will demonstrate these properties of the halo gas in Fig. 3.13.

Besides all gas, we also want to investigate the gas in the CGM specifically. We show the mass fraction in different baryonic components as a function of halo mass in the left panel of Fig. 3.3. Here, we consider everything within R_{200c} of the central galaxy to be part of the halo. The black hole contribution is too small to appear on the plot. The total baryon fraction increases with halo mass, and is substantially smaller than the cosmic fraction for $M_{200c} < 10^{13} M_{\odot}$. The trend at lower halo masses ($M_{500c} < 10^{13} M_{\odot}$) is not currently constrained by observations. The EAGLE baryon fractions are somewhat too high for $M_{200c} > 10^{14} M_{\odot}$ (Barnes et al. 2017). The observations do support the trend of rising baryon fractions with halo mass at high masses.

The CGM mass fraction increases with halo mass, while the stellar and ISM fractions peak at $M_{200c} \sim 10^{12} M_{\odot}$, with the ISM fraction declining particularly steeply towards higher masses. This is likely a result of star formation quenching starting in $\sim L_{*}$ galaxies. The ‘missing baryons’ CGM at $10^{5.5} - 10^7$ K dominates for halo masses $M_{200c} \sim 10^{12} - 10^{13.5} M_{\odot}$,




Figure 3.3: *Left panel:* The fraction of halo mass (i.e., $< R_{200c}$) in stars, ISM, and CGM as a function of halo mass in the EAGLE simulation at $z = 0.1$. The grey line shows the total baryon mass fraction, and the purple line shows CGM gas with temperatures in the $10^{5.5}-10^7$ K range. The dashed, horizontal line indicates the cosmic baryon fraction. *Right panel* (note the different y-axis range): The fraction of halo oxygen mass (oxygen ejected by stars, currently within R_{200c}) in stars, ISM, and CGM in haloes of different masses. The halo oxygen budget (total and in stars) does not include metals produced in stars that have never been ejected, or any oxygen captured by black holes. The solid lines show medians and shaded regions show the 80 per cent halo-to-halo scatter in each halo mass bin; the shading is omitted for legibility for the total baryons and $10^{5.5}-10^7$ K CGM. Here, ISM is defined as all star-forming gas and CGM as the other gas. We use 0.1 dex halo mass bins for $M_{200c} < 10^{13} M_{\odot}$ haloes, then 0.25 dex bins, and one bin for the haloes above $10^{14} M_{\odot}$. The CGM is typically the largest baryon mass component in haloes, and typically contains more metals than the ISM at all halo masses we study.

which is what we would expect according to T_{200c} . The $M_{200c} \sim 10^{12}-10^{13.5} M_{\odot}$ haloes where this gas dominates are indeed the ones that dominate the ion budgets in Fig. 3.2, except for O VI, which probes cooler gas, and Fe XVII, which probes gas in this temperature range, but where the dominant haloes include some higher-mass ones, in agreement with T_{200c} (Fig. 3.1).

The right panel of Fig. 3.3 similarly shows the fraction of oxygen in different baryon components for haloes of different masses. Oxygen produced in stars, but never ejected is not counted. A smaller fraction of the oxygen that was swallowed by black holes is not tracked in EAGLE. The fraction in stars therefore reflects the metallicity of the gas the stars were born with. The fractions for neon are nearly identical to those for oxygen, while the curves for iron have the same shape, but with a somewhat smaller mass fraction in stars and more in CGM and ISM.

We see that at lower halo masses, most of the metals in haloes reside in stars, while for $M_{200c} \gtrsim 10^{13} M_{\odot}$, more metals are found in the CGM. The changes with halo mass seem to be in line with the overall mass changes in ISM and CGM as halo mass increases (Fig. 3.3), though the stars and ISM contain higher metal fractions than mass fractions, reflecting their higher metallicities. Interestingly, there are more metals in the CGM than in the ISM for all halo masses, though the difference is small for $M_{200c} < 10^{12} M_{\odot}$. This is similar to what Oppenheimer et al. (2016) found for a smaller set of haloes with EAGLE-based halo zoom simulations. They considered all the oxygen produced in galaxies within R_{200c} , in 20 zoom

simulations of $M_{200c} = 10^{11} - 10^{13} M_{\odot}$ haloes, and found that a substantial fraction of that oxygen ($\sim 30 - 70$ per cent) is outside R_{200c} at $z = 0.2$. That oxygen is not included in the census in Fig. 3.3.

The mass and oxygen fractions in the CGM and ISM do depend somewhat on the definition of the ISM. (The CGM is all gas within R_{200c} that is not ISM in all our definitions.) In Fig. 3.3 we define the ISM as all gas with a non-zero star formation rate. Since the minimum density for star formation in EAGLE is lower for higher metallicity, higher-metallicity gas is more likely to be counted as part of the ISM. If we define the ISM as gas with $n_{\text{H}} > 10^{-1} \text{ cm}^{-3}$ instead, the mass fractions change. Per halo, the ISM mass changes by a median of $\approx -30 - 50$ per cent for $M_{200c} \lesssim 10^{12} M_{\odot}$, ≈ 0 per cent at $\sim 10^{13} M_{\odot}$, and up to $+30$ per cent at higher masses. The central 80 per cent range is large, including differences comparable to the total ISM mass using the star-formation definition in both directions. The scatter in differences is largest at low masses. The median trend with halo mass makes sense given the higher central metallicities (meaning lower minimum n_{H} for star formation) we find in lower-mass haloes (Fig. 3.13). If we count gas that is star-forming or meets the n_{H} threshold as ISM, the ISM mass can only increase relative to the star-forming definition. Median differences are $\lesssim 3$ per cent at $M_{200c} \lesssim 10^{12} M_{\odot}$, but increase to $\approx 30 - 60$ per cent at $M_{200c} \gtrsim 10^{13} M_{\odot}$. Since the CGM contains more mass overall, differences in the CGM mass using the two alternative ISM definitions are typically $\lesssim 11$ per cent (central 80 per cent of differences).

The ISM definitions also affect how oxygen is split between the ISM and CGM. Using the $n_{\text{H}} > 10^{-1} \text{ cm}^{-3}$ definition results in lower ISM oxygen fractions, with median per-halo differences $\approx -20 - 55$ per cent, and a central 80 per cent range of differences mostly between ≈ -10 and -90 per cent. CGM fractions are consistently higher, with median per-halo differences of up to ≈ 40 per cent at $M_{200c} < 10^{12.5} M_{\odot}$, decreasing to close to zero between $M_{200c} = 10^{12}$ and $10^{13} M_{\odot}$. Using the combination ISM definition ($n_{\text{H}} > 10^{-1} \text{ cm}^{-3}$ or star-forming) does not change the oxygen masses by much, since dense, but non-star-forming gas has a low metallicity. The central 80 per cent of per halo differences is < 1 per cent at all M_{200c} .

Finally, since we are primarily interested in ions in this work, we look into the ionization states of the metals in haloes of difference mass.² The bottom panels of Fig. 3.4 show the fraction of ions in the CGM (all gas at $0.1 - 1 R_{200c}$) as a function of halo mass, compared to the CIE ion fractions at the halo virial temperatures in the top panels. In the left panels, we show median ion fractions with the 10th–90th percentile range, for the ions we focus on in this work. In the right panels, we show the average fractions of all the ionization states of oxygen. Note that the ionization table we use does not include the effects of self-shielding (or local radiation sources), so the lowest ionization state, O I, could be underestimated.

For the ions we focus on in this work, including the gas within $0.1 R_{200c}$ has a negligible effect, since there is very little highly ionized gas there (Fig. 3.12). Including gas out to $2 R_{200c}$ does make a difference. If that gas is included, these ion fractions rise, especially at the low- and high-mass ends, and the peaks of the ionization curves shift to slightly higher masses. The larger overall ion fractions are likely due to the increased amount of gas photo-ionized to the higher states we examine here at larger distances. The slight shifts are likely due to the lower gas temperatures in the same haloes at larger distances (Fig. 3.13).

For the lower ionization states in the bottom right panel, whether or not we include gas at radii $< 0.1 R_{200c}$ has more of an effect: including this gas increases the O I and O II content by large amounts; the fraction of the total increases by $\approx 0.2 - 0.4$ for $M_{200c} \sim 10^{11} -$

²Here, we use CLOUDY version 13 ionization tables for the oxygen ions, instead of the version 7.02 tables used in the rest of the paper.




Figure 3.4: Total fraction of each element in the ionization states indicated in the legends in the $z = 0.1$ EAGLE CGM, as a function of halo mass ($0.1-1 R_{200c}$; lower panels). The top panels show the CIE ion fractions as a function of temperature, assuming the temperatures match T_{200c} for the different halo masses. The left panels show these fractions for the six ions we focus on in this work. The solid lines show median ion fractions in different mass bins, while the shaded regions in the same colour show the scatter (percentiles 10–90). The virial temperature and CIE ionization curves predict the qualitative trends of ionization fraction as a function of halo mass, but strongly underestimate the ion fractions in low-mass haloes. The right panels show the average fraction of oxygen in the indicated ionization states in the CGM. The ions O I– VIII are shown from bottom to top in order.

$10^{12} M_{\odot}$, with the effect decreasing toward higher halo masses. The difference will be due to the fact that the central galaxy contains plenty of cold gas, but very little of the more highly ionized species. (Wijers et al. (2019) verified that the O VII and O VIII CDFs are negligibly impacted by whether or not star-forming gas is accounted for.) Including gas at larger radii (out to $2 R_{200c}$) increases the fraction of oxygen in the O VI– VIII states, at the cost of gas in lower states, but also at the cost of O IX at $M_{200c} \gtrsim 10^{12} M_{\odot}$.

For the high ions in the left panels, we confirm by comparing the top and bottom panels that the CIE ionization peak and halo virial temperatures are good predictors of the qualitative trends of halo ion content as a function of halo mass, but the CIE($T = T_{200c}$) curves strongly underestimate the ion fractions at low mass, where photo-ionization dominates.

The CIE curves peak at slightly larger halo masses than EAGLE haloes show. This might be because the temperature inside R_{200c} is typically higher than T_{200c} . We will show this using the mass- and volume-weighted temperature profiles in Fig. 3.13. Alternatively, or additionally, photo-ionization may be responsible, by lowering the typical temperature at which the ions are preferentially found. Fig. 3.13 shows this would mostly be important at lower halo masses ($M_{200c} \lesssim 10^{12} M_{\odot}$) or at radii approaching R_{200c} .

For O VI, we do not find a peak at all in the halo mass range we examine. This is due to photo-ionization becoming important at and below the halo masses where CIE would produce an O VI peak, in the same regime where other halo ion fractions flatten out.

As in the left panels, the CIE curve for a single temperature predicts much more extreme ion fractions than we see in the EAGLE haloes. In particular, Fig. 3.4 shows that lower-mass haloes contain many high ions, and that the lowest ionization states peak at much higher masses than CIE($T = T_{200c}$) predicts, suggesting the presence of significant amounts of gas with $T \ll T_{200c}$.




Figure 3.5: The CDDFs for the ions we consider in this paper, for the EAGLE simulation at $z = 0.1$. Coloured ticks on the x-axis roughly indicate the positions of breaks in the CDDFs (determined visually), which serve as reference points in further figures. These are at the same position for O VII and O VIII, but the ticks are slightly offset for legibility.

On the other hand, the higher high-ion fractions than suggested by the CIE curves indicate the presence of $T \gg T_{200c}$ gas in sub- L_* haloes. This is likely a result of gas heating by stellar (and at higher masses, AGN) feedback. Temperature distributions indicate this is not only a result of the direct heating of particles due to feedback in EAGLE, but that sub- L_* haloes have smooth mass- and volume-weighted temperature distributions that can extend to $\sim 10^6$ K or somewhat higher at $\sim R_{200c}$. Besides this hotter gas, photo-ionized gas close to R_{200c} also plays a part: at these radii in $M_{200c} \lesssim 10^{12} M_\odot$ haloes, gas densities can reach $n_H \sim 10^{-5} \text{ cm}^{-3}$ (Fig. 3.13), where photo-ionization becomes important. The importance of photo-ionization for the CGM ion content was previously pointed out by Faerman et al. (2020) in their isentropic model of the CGM of an L_* galaxy.

3.3.3 Column density distributions and equivalent widths

Before we look into metal-line absorption around haloes, we consider metal-line absorption at random locations. We consider how their column densities relate to the more directly observable equivalent widths of absorption systems, and how haloes contribute to the absorbers we expect to find in a blind survey.

In Fig. 3.5 we show the column density distributions for the six ions we focus on. The ions all show distributions with roughly two regimes, with a shallow and steep slope at low and high column densities, respectively. The coloured ticks on the x-axis indicate the ‘knees’ which mark the transition between these regimes, determined visually. The ticks are for reference in other figures.

In Fig. 3.6 we explore how haloes contribute to this absorption along randomly chosen sightlines. It shows the contributions of different halo masses to the CDDFs of our six ions.




Figure 3.6: The contribution of absorption by haloes of different masses to the column density distributions of the ions indicated in the panels at $z = 0.1$ in the EAGLE simulations. The black line indicates the distribution of all absorption systems, while the brown, dashed line indicates the contribution of all haloes (including those with $M_{200c} < 10^{11} M_{\odot}$). The colour bar indicates the mass range for which each solid, coloured line represents the contribution to the CDDF. Contributions are determined by computing CDDFs from column density maps made with only gas in each halo mass range (in a FoF group or R_{200c} sphere): the halo-projection method in §3.2.7

The CDDFs for each halo mass bin are generated from the simulations in the same way as the total CDDFs, but using only SPH particles belonging to a halo of that mass (the halo-projection method from §3.2.7). An SPH particle belongs to a halo if it is in the halo’s FoF group, or within R_{200c} of the halo centre.

From Fig. 3.6 we see that for the X-ray ions, most absorption at column densities higher than the knee of the CDDF is due to haloes. This confirms the suspicion of Wijers et al. (2019) that this was the case for O VII and O VIII, based on the typical gas overdensity of absorption systems at these column densities. However, for the FUV/EUV ions O VI and Ne VIII, there is a substantial contribution from gas outside haloes at these relatively high column densities.

For all these ions, we also note the following trend. The absorption at higher column densities tends to be dominated by more massive haloes until a turn-around is reached. These turn-around masses are consistent with the temperatures preferred by the ions, suggesting they are being driven by the increase in virial temperature with halo mass (compare to Fig. 3.1). We have verified that trends with halo mass are not driven simply by the covering fraction of haloes of different masses.

To get a sense of what column densities might be detectable with different instruments (§3.4) we look into what rest-frame EWs these column densities typically correspond to.




Figure 3.7: Rest-frame EWs for the ions we investigate as a function of ion column density at $z = 0.1$ in the EAGLE simulation. The left axes show EWs in \log_{10} mÅ, the right axes show \log_{10} eV. The solid, grey line shows the median EW in bins of 0.1 dex in column density, while the shading shows the central 80 per cent (darker grey) and central 96 per cent (lighter grey) of the EWs in the bins. For EWs outside these ranges, and column density bins with fewer than 50 sightlines, we show each sightline as a single grey point. We also show best-fitting values (using eq. 3.6) for the Gaussian line broadening b (eq. 3.4) in blue dot-dashed lines. The best-fitting values are indicated in the bottom right of each panel. The relation for unsaturated absorption is shown with a dotted green line. The orange, dashed line shows the thermal broadening for ions at the temperature where their ion fraction is at a maximum in CIE (equation 3.7, Fig. 3.1). The various dotted brown lines show the column-density-EW relation for Voigt profiles with different Gaussian line broadening values (i.e., b parameters): 10, 20, 50, 100, and 200 km s^{-1} , from bottom to top in the panels. The spectra and column-density-EW relations are for absorption lines at a single rest-frame wavelength, except for O VIII, where we model doublet absorption.

Though we will work with column densities in the rest of this paper, the fits we find can be used to (roughly) convert between the two. Fig. 3.7 shows typical EW as a function of column density.

We parametrize the column-density-EW relation using the width of the Gaussian part of the Voigt profile b , as described in §3.2.6. We list the best-fitting parameters in Table 3.4 and show the relation for these parameters in Fig. 3.7. The shadings in Fig. 3.7 give an indication of how broad the b -parameter distribution is (10th and 90th, 2nd and 98th percentiles). We will use these best-fitting b -parameters in §3.4 to estimate the minimum column densities observable with the Athena X-IFU, Arcus, and the Lynx XGS. We explore the dependence of the best-fitting values on the velocity windows in which we measure column densities and EWs in Appendix 3.A.

Generally, the thermal line broadening expected at the temperature where the ion frac-

Table 3.4: Best-fitting b parameters to the column-density-EW relation for the different ions, derived from EAGLE mock spectra at $z = 0.1$. Rest-frame EWs are calculated for each ion (first column) using the absorption lines in Table 4.1. The second column indicates the half width of the velocity windows used to calculate the EWs. Column 3 shows the best-fitting b parameters (eq. 3.6), using the velocity windows in column 2.

ion	Δv km s ⁻¹	$b(\Delta v)$ km s ⁻¹
O VI	300	28
Ne VIII	300	37
O VII	800	83
Ne IX	800	82
O VIII	800	112
Fe XVII	800	92

tion peaks in CIE,

$$b(T_{\max, \text{CIE}}) = \sqrt{2kT_{\max, \text{CIE}} m_{\text{ion}}^{-1}}, \quad (3.7)$$

gives a good lower limit³ to the EWs (dashed orange lines). Here, m_{ion} is the ion mass. For O VII, Ne IX, Fe XVII, and particularly O VI, lower values do occur. For O VII, Ne IX, and Fe XVII, this is still consistent with the lower end of the CIE temperature range in Table 3.3: $b = 16, 20,$ and 24 km s^{-1} , respectively. For O VII, this was previously described by Wijers et al. (2019). For O VI, the lower CIE end gives $b = 14 \text{ km s}^{-1}$, which does not cover this range. Such low b values are rare for this ion, but their occurrence suggests at least some high-column-density O VI is photo-ionised.

In Fig. 3.7 we can also see the importance of Lorentz broadening for the EWs of the different absorption lines. The single-component absorber curves (all lines except the grey ones) show an upturn where the ‘wings’ of the Voigt profile become important. This becomes relevant for narrow, high-column-density absorbers for the X-ray lines, especially O VIII and Fe XVII. For the UV lines, the effect of the Lorentz broadening is negligible, since the extra broadening is smaller relative to their wavelengths compared to the X-ray lines (Table 4.1).

3.3.4 Column density profiles

The CDDFs we examined are predictions for finding absorption lines at random in the spectra of background sources. However, it is also common to look for absorption close to galaxies specifically, especially in stacking studies. Therefore, we consider what we might find if we looked for absorption around haloes of different masses. For this, we use the radial profiles computed as described in §3.2.7. The column density radial profiles are shown in Fig. 3.8. The solid lines show absorption by all gas in the same 6.25 cMpc slice as the halo centres, while the dashed lines show absorption only by gas in a halo (FoF group or otherwise inside R_{200c}) with M_{200c} in the matched halo mass range. In principle, this means that single-halo profiles might include absorption by gas in different haloes of similar mass, but the fact that the dashed lines for all ions drop off sharply at the same $r_{\perp} \approx 1.5 R_{200c}$ indicates that this effect is negligible, at least for the median profiles.

³For a given column density, non-thermal broadening or multiple absorption components spread out the ions in velocity space, meaning the absorption is less saturated. Therefore, a single line or doublet with only thermal broadening should give a lower limit to the EW of an absorption system at fixed column density.




Figure 3.8: Radial column density profiles for the different ions around central galaxies in haloes of different masses at $z = 0.1$ in the EAGLE simulation. The median column density at different impact parameters is shown, with shaded regions showing the 10th–90th percentile range for a halo mass bin with a high central peak column density. Column densities are measured in the single 6.25 cMpc slice of the box that contains each central galaxy’s centre of mass. The horizontal, grey ticks on the y-axis indicate roughly where breaks in the CDDFs for the different ions occur (see Fig. 3.5). Solid lines include all gas from a given slice, while dashed lines show absorption around haloes in each slice coming only from gas within haloes (FoF group or within R_{200c}) in the same M_{200c} range. For legibility, we only show the halo contributions for three halo mass bins: the mass yielding the maximum median column density for each ion, and the highest and lowest masses.

We see a clear pattern: the median column density increases with halo mass until it reaches a peak, which corresponds to the halo mass where the relative contribution to the CDDF (at higher column densities) peaks in Fig. 3.6. This again supports the idea that the column densities of these haloes are largely driven by the halo virial temperature.

We also note more qualitative trends. Column densities at large distances ($\gtrsim 2 R_{200c}$) increase considerably less with halo mass than central column densities do. At halo masses beyond the peak, the median column density declines and the profile flattens within R_{200c} , even having a deficit of absorption somewhere in the range ~ 0.1 – $1 R_{200c}$ compared to $\sim R_{200c}$ for the lower-energy ions in the largest halo mass bins. We will examine the causes of these trends in §3.3.5 using (3D) radial profiles of the halo gas properties.

The fraction of absorption caused by gas in the haloes (dashed curves) also shows a clear trend: the halo contributions are largest in halo centres, and for haloes at the mass where the median column density peaks. Halo contributions drop as typical column densities decrease, towards both higher and lower halo masses.

Comparing to the column densities where breaks in the CDDFs occur (long horizontal

grey ticks on the left), we see the absorption in the high column density tails of the overall distribution, at column densities above those indicated, comes from absorbers that are stronger than typical for haloes of any mass. Therefore, the low occurrence of stronger absorbers does not simply reflect the low volume density of haloes in the ion's preferred mass range, it is also due to the fact that they are relatively high column density absorbers for such haloes. Note that the scatter here includes both inter-halo and intra-halo scatter, so it is possible that such absorbers are more common in a subset of haloes at some halo mass.

In observations, halo masses can be uncertain, especially around low-mass galaxies. Therefore, we also show radial profiles in bins of central galaxy stellar mass, as a function of projected distance to the galaxy centre of mass. Unlike before, we obtain the median and scatter in column density in bins of physical impact parameter. We only consider central galaxies here. The profiles are shown in Fig. 3.10.

We use bins spaced by 0.5 dex in stellar mass. However, we do not use a separate bin for $M_\star = 10^{11.5} - 10^{11.7} M_\odot$, since this bin would only contain six galaxies. Instead, we group all $M_\star > 10^{11} M_\odot$ galaxies into one bin. We use at most 1000 galaxies (randomly selected) for the profiles for each M_\star bin, which is relevant for galaxies with $M_\star < 10^{10.5} M_\odot$.

Fig. 3.9 shows the stellar-mass-halo-mass relation for EAGLE central galaxies as a 'confusion matrix'. It shows how the stellar mass bins we use in this section map onto the halo mass bins used in the rest of the paper. According to Schaye et al. (2015), the galaxy stellar mass function is converged with resolution down to stellar masses $M_\star \approx 2 \times 10^8 M_\odot$, though for other properties, such as star formation rates, the lower limit is $\sim 10^9 M_\odot$ or somewhat more massive. In the lowest halo mass bin we considered ($M_{200c} = 10^{11} - 10^{11.5} M_\odot$), we do find a substantial contribution from $M_\star < 10^9 M_\odot$ galaxies, but most central galaxies in this halo mass bin have $M_\star > 2 \times 10^8 M_\odot$. Fig. 3.9 also shows that the highest three halo mass bins will have little impact outside the largest stellar mass bin, and the very largest halo mass bin contains too few galaxies to contribute significantly for any stellar mass.

Fig. 3.10 shows the same main trends of column density with M_\star in physical distance units as Fig. 3.8 showed for normalized distance and halo mass. However, for O VIII, Ne IX, and Fe XVIII, the fact that the highest-stellar-mass bin contains mostly $M_{200c} < 10^{13.5} M_\odot$ haloes means we do not see a decrease in column density towards the highest stellar masses. The overall correspondence implies that, with sufficiently sensitive instruments and large enough sample sizes, the column density trends with halo mass should be observable.

Note that the innermost parts of these profiles ($r_\perp \ll 10$ pkpc) might be less reliable, where they probe the central galaxy or gas close to it. Wijers et al. (2019) found that including or excluding star-forming gas altogether has very little effect on the CDDFs of O VII and O VIII. Indeed, in making the column density maps, we assumed all this gas had a temperature of 10^4 K, too cool for these ions at these high densities (Fig. 3.1). However, in reality, a hot phase in the ISM may contain such ions. On the other hand, in EAGLE, there is hot, low-density gas in halo centres in (§3.3.5), which may have been directly heated by star formation or AGN feedback, and might cause absorption that is sensitive to the adopted subgrid heating temperatures associated with these processes.

Now that we have examined median column densities, we consider the extreme end of the distribution: how much absorption we find at very high column densities as a function of impact parameter. By very high column densities, we mean those above the CDDF breaks in Fig. 3.5. The values are listed in Table 3.5 and shown in Fig. 3.11. For a number of ions, the absorption lines we analyse here (Table 4.1) (start to) become saturated at these column densities (Fig. 3.7). For unresolved X-ray lines, these covering fractions might therefore be difficult to measure observationally as long as the widths of the absorption components re-




Figure 3.9: The EAGLE M_* - M_{200c} relation for central galaxies at $z = 0.1$ shown as a confusion matrix, demonstrating how our M_* and M_{200c} bins compare. The number of central galaxies in each M_* , M_{200c} bin is shown. The lowest-mass bins have no lower limit, and include galaxies and haloes that are unresolved in the simulation. The colours show what fraction of galaxies in each M_* bin are in haloes in different M_{200c} bins.

Table 3.5: Threshold column densities ($\log_{10} N$ [cm^{-2}]) used for covering fractions for the different ions we show in Figs. 3.11 and 3.14. The EW cited in the top line of the table is an observer-frame (redshifted) value. It was converted into column densities using the lines from Table 4.1 at $z = 0.1$, and the best-fitting b parameters from Table 3.4 (blue, dot-dashed lines in Fig. 3.7). The sources for the data are described in §3.4.

	O VI	Ne VIII	O VII	Ne IX	O VIII	Fe XVII
EW = 0.18 eV			15.4	15.4	15.6	14.8
HST-COS	13.5	13.5				
CDDF break	14.3	13.7	16.0	15.3	16.0	15.0

main unresolved, which is expected even for the Athena X-IFU (Wijers et al. 2019, fig. 4).

In Fig. 3.11 we see that the covering fractions above the CDDF break typically peak close to galaxies. However, the relatively small cross-section of these central regions means that absorption above the break in the CDDF for blind surveys is dominated by regions outside the inner 30 pkpc around galaxies. We determined this from the covering fraction profiles at different M_* , and the total CDDFs for the ions. We compared different sets of absorbers. The first set are the absorbers in the central regions. These are absorbers in the same 6.25 cMpc slice of the simulation, with impact parameters $r_\perp < 30$ pkpc (< 30 pkpc absorbers). The second set is similar, but contains absorbers with $r_\perp \lesssim R_{200c}$. In each stellar mass bin, we use the median R_{200c} of the parent haloes to define this edge. We estimate the number of absorbers above the column density breaks in the two r_\perp ranges from the covering fraction profiles.




Figure 3.10: Median column density as a function of impact parameter (physical distance) and stellar mass of the central galaxy at $z = 0.1$ in the EAGLE simulation. Different panels are for different ions. Long ticks on the y-axes (left) indicate approximately where the breaks in the CDDFs occur (Fig. 3.5). The profiles extend out to $3R_{200c}$ of the 99th percentile of the M_{200c} distribution of each M_{\star} bin (see Fig. 3.9). Points on each curve mark the median virial radius in each M_{\star} bin, and horizontal lines show the central 80 per cent of virial radii in those bins. Shaded regions show the 10th–90th percentile range for a stellar mass bin with a high central column density.

The < 30 pkpc absorbers contain $\lesssim 10$ per cent of the absorption above the CDDF break in the $\lesssim R_{200c}$ sample, at least in M_{\star} bins responsible for > 10 per cent of the total absorption above the CDDF breaks. For M_{\star} bins responsible for less of the total absorption, the < 30 pkpc absorbers make up $\lesssim 33$ per cent of the $\lesssim R_{200c}$ absorbers (with one exception of 42 per cent: Fe xvii around $M_{\star} = 10^{10} - 10^{10.5} M_{\odot}$ galaxies). Looking back to Fig. 3.10, this also means that absorption above the CDDF break is indeed dominated by scatter in column densities around galaxies at larger radii, rather than typical absorption where column densities are highest.

3.3.5 Halo gas as a function of radius

In order to better understand the overall contents of haloes, as well as their absorption profiles, we examine the gas and ions in haloes as a function of (3-D) radius. In Fig. 3.12 we show various cumulative 3-D profiles for each halo. These profiles come from averaging individual haloes’ radial mass distributions, after normalizing those distributions to the amount within R_{200c} . This means that the combined profiles reflect typical (ion) mass distributions, without weighting by halo mass, baryon fraction, or halo ionization state.

Most of the ions in these haloes lie in the outer CGM ($r \gtrsim 0.3 R_{200c}$). This explains the relatively flat absorption profiles out to $\sim R_{200c}$ in Fig. 3.8. The S-shaped cumulative profiles




Figure 3.11: Covering fractions around central galaxies of different stellar masses in EAGLE at $z = 0.1$ as a function of physical impact parameter. The covering fraction is the fraction of sightlines with column densities larger than a threshold value at each impact parameter. The threshold column densities, in $\log_{10} \text{cm}^{-2}$ units, are shown in the panels. These covering fractions are for column densities equal to the respective CDDF breaks. Points on each curve mark the median virial radius in each M_{\star} bin, and horizontal lines show the central 80 per cent of virial radii in those bins. (Some are outside the range of the plot.)




Figure 3.12: Average cumulative volume, gas mass, and ion mass 3-D profiles for the different ions in the EAGLE simulation at $z = 0.1$. We show the average enclosed fraction of each quantity, normalized for each halo by the amount enclosed within R_{200c} . This shows that a large fraction of the ions is near R_{200c} . The different panels are for different M_{200c} ranges, shown in the top left corners in units $\log_{10} M_{200c} / M_{\odot}$. Since the trend with halo mass is weak, we only show three halo mass bins.

at large halo masses explain the second peaks around R_{200c} in the radial profiles of some of the high-mass haloes: most of the lower-energy ions, like O VI, in these haloes lie in a shell at large radii, which leads to a peak in the 2D-projected column densities. The enclosed ion fractions generally fall between the enclosed mass and volume fractions. Exceptions are lower ions in the inner CGM of high-mass haloes. Also, Fe XVII is more centrally concentrated than the other ions and gas overall, as Fig. 3.8 also showed. We will discuss this in more detail later. The high spike in O VI mass at large radii in low-mass haloes is not present in a small, random sample of individual halo O VI profiles, and is therefore not a typical feature for this halo mass.

In Fig. 3.13, we show mass- and volume-weighted median temperature, density, and metallicity profiles (left column). For the temperature profiles, the dotted lines show the simple prediction: T_{200c} , as calculated from eq. 4.1. The colours match the median halo mass in each bin. The profiles show a general rising trend with halo mass, with temperatures at R_{200c} matching T_{200c} reasonably, and following the T_{200c} trend. However, the temperature clearly decreases with radius in most cases. The exceptions are the inner parts of the profiles for low-mass haloes, for which the volume-weighted median temperatures can be much higher than the mass-weighted ones. The haloes are in fact multiphase, with cool gas at $\sim 10^4$ K, some gas at $\sim 10^5$ K, and the hotter volume-filling phase. Sharp transitions in the median profiles occur when the median switches from one phase to another. The multiphase nature is particularly prominent at low mass ($M_{200c} \ll 10^{13} M_{\odot}$).

The gas density also decreases with radius. It is generally higher in higher-mass haloes around R_{200c} , but at larger radii it also drops much faster than in lower-mass systems. Volume-weighted densities can be considerably lower than mass-weighted densities, reflecting the multi-phase nature of the gas. In the centres of low-mass haloes (especially at smaller radii than shown), median temperatures tend to increase as densities drop. This is likely the result of stellar and/or AGN feedback heating some gas in the halo centres, increasing its temperature and volume. These large volumes for particles centred close to the halo centres may dominate the volume-weighted stacks. Indeed, in this regime, it seems gas at a few discrete temperatures sets these trends (including at $10^{7.5}$ K, the heating temperature for stellar feedback), and medians from a few randomly chosen individual galaxies do not show this trend.

The metallicities also tend to decline with radius, with larger differences in lower-mass haloes. Evidently, the metals are better mixed in high-mass haloes, likely because star formation and the accompanying metal enrichment tend to be quenched in these systems, while processes such as mergers and AGN feedback continue to mix the gas. The mass-weighted metallicities are higher than the volume-weighted ones in the inner parts of the halo, while in the outer halo and beyond R_{200c} , the differences depend on halo mass. For high-mass haloes, the median metallicity in the volume-filling phase drops sharply somewhat outside R_{200c} . However, the scatter in metallicity at large radii is very large, particularly at high masses.

Fig. 3.13 also shows the corresponding ion-mass-weighted temperature, density and metallicity profiles as a function of halo mass. The coloured, dotted lines in the temperature profiles show T_{200c} at the median mass in each bin. The black, dotted lines indicate the CIE temperature range for each ion. Abrupt temperature changes are again a result of the median switching between different peaks in the temperature distribution of multiphase gas.

The ion-weighted temperature mostly follows the CIE temperature range (black dotted lines), rather than the T_{200c} range for that set of haloes (coloured dotted lines) within R_{200c} . Higher ion-weighted temperatures do occur, but in radial regions that contain relatively few ions. Ion-weighted temperatures below the CIE range mainly occur at radii




Figure 3.13: Median temperatures, densities, and metallicities of the halo gas in the EAGLE simulation at $z = 0.1$. The left panels show mass-weighted (dashed) and volume-weighted (solid) medians, while the other columns show medians weighted by the different ions. To obtain the medians, we first computed temperature, density, and metallicity histograms of all gas in each radial bin, weighted by mass, volume, or ion mass, for each halo. We then normalized these histograms by the total weighting quantity within R_{200c} , and averaged the normalized histograms of the haloes in each mass bin. We computed the medians from these stacked histograms. We used the same radial bins in R_{200c} units for all haloes. The dotted lines show T_{200c} (eq. 4.1) for the median halo mass in each bin, in colours matching the full range. The densities were calculated as mass densities and converted to hydrogen number densities assuming a primordial hydrogen mass fraction of 0.752. The metallicities are oxygen mass fractions (`SMOOTHEDELEMENTABUNDANCE/OXYGEN`), normalized to a 0.00549 solar mass fraction (Allende Prieto et al. 2001). The neon and iron mass fractions (not shown) follow similar radial and halo mass trends, though the values differ somewhat. We show a subset of halo masses for legibility. The lowest-radius bin contains all mass/volume/ions within $0.1 R_{200c}$. The ion temperatures are mostly set by CIE, while their densities roughly follow the volume-weighted density profile (i.e., the hot gas profile). The ion-weighted metallicities are biased high compared to the mass- and volume-weighted metallicities.

$\gtrsim R_{200c}$, where ion-weighted hydrogen number densities reach the $\sim 10^{-5} \text{ cm}^{-3}$ regime where photo-ionization becomes important and lower-temperature gas can become highly ionized (Fig. 3.1). For O VI, ions at lower temperatures do persist at smaller radii, within the 90 per cent scatter of the ion-weighted temperature, and especially in lower-mass haloes.

The ion-weighted densities in the CGM reflect the halo’s physical properties: they follow the halo gas density distribution, and in particular, the volume-filling hot phase in cases where the mass- and volume-weighted gas distributions differ. They are however biased to the temperature ranges favoured by CIE and the metallicities are biased high compared to the mass- and volume-weighted values shown in Fig. 3.13.

These temperature and density effects may explain the ‘shoulders’ around R_{200c} in the absorption profiles at some halo masses seen in Fig. 3.8. This phenomenon occurs at halo masses around or above those for which T_{200c} matches the CIE peak for each ion. Since the temperature of the CGM decreases with radius, the ions will preferentially be present at larger radii in higher-mass haloes, where they form a ‘shell’, which produces large column densities at projected radii close to the shell radius. This is visible in Fig. 3.12, where lower ions in higher-mass haloes have S-shaped cumulative ion mass distributions, with relatively little ion mass in the too-hot inner CGM.

However, around R_{200c} , as the halo-centric radius increases, the effect of the declining gas temperature is countered by photo-ionization of the cold phase, which also starts to become important around R_{200c} . This drives the preferred temperatures of the ions down with radius, along with the gas temperature. The ‘shoulders’ are strong in Fe XVII and Ne IX profiles; these ions have the highest ionization energies (Table 3.3) and are photo-ionized at lower densities (Fig. 3.1) than the others.

The sharp drops in the absorption profiles at large radii in Fig. 3.8 may also be explained by these halo properties: the gas density drops outside R_{200c} , and more sharply for higher halo masses. Similarly, the gas metallicity drops rapidly around these radii in the high-mass haloes. The differences between the ions seem to be consistent with the more easily photo-ionized ones producing more absorption in the cooler, lower-density gas around R_{200c} . However, the way the CIE temperature range lines up with gas temperatures depends on both the ion and the halo mass, so ion and mass trends are difficult to disentangle.

Though Fe XVII seems like an outlier in Fig. 3.12, in that it is more concentrated in halo centres than the (total) gas mass, this does fit into these trends: the outskirts of most haloes at the masses we consider are simply too cool for this ion. However, in $M_{200c} > 10^{14} M_{\odot}$ haloes, which have T_{200c} above the preferred range of Fe XVII, the absorption does extend out to R_{200c} , albeit at lower column densities.

We note that the sharp drops in mass- and volume-weighted median metallicity are not in contradiction with the flat ion-mass-weighted metallicities outside R_{200c} : there is very large scatter in the metallicity at large radii, and metal ions will preferentially exist in whatever metal-enriched gas is present.

3.4 Detection prospects

To predict what might be observable with different instruments, we first estimate the minimum observable column densities for the different ions. We use column density thresholds which correspond roughly to the detection limits of current, blind (UV) and upcoming (X-ray) surveys. We then use these limits to predict how many absorbers and haloes we should be able to detect per unit redshift, and out to what impact parameters we can expect to find measurable absorption.

3.4.1 Detection limits for different instruments

For the X-ray lines, we estimate the minimum detectable column density from the minimum detectable EW and the b -parameters from Table 3.4 assuming a single Voigt profile (or a doublet, for O VIII). Since these minima depend not just on the instrument, but on the observations (e.g., exposure time, background source flux and spectrum), we take the minimum EWs from the instrument science requirements, which assume a planned observing campaign as well as instrument properties. These are observer-frame minima, which we convert to rest-frame minimum EWs assuming $z = 0.1$, the redshift we assume throughout this work.

We focus on what should be detectable with the X-IFU on the planned Athena mission. Here, weak lines around 1 keV should be detectable at 5σ significance at observer-frame EWs of 0.18 eV. This is for 50 ks exposure times and a quasar background source with a 2–10 keV flux of 10^{-11} erg cm $^{-2}$ s $^{-1}$ and a photon spectral index $\Gamma = 1.8$ (Lumb et al. 2017). Blind detections of pairs of O VII and O VIII absorption lines should be possible at lower EWs than this, at least against bright Gamma-ray burst background sources (Walsh et al. 2020). We convert these minimum EWs to minimum column densities using the best-fitting relations shown in Fig. 3.7 (blue, dot-dashed lines) and Table 3.4. The minimum column densities are shown in Table 3.5.

A minimum EW estimate of 0.18 eV is on the rough side for Fe XVII and Ne IX, since the oxygen lines have been the main focus of WHIM and hot CGM detection plans. These lines are at different energies, so the energy-dependence of the sensitivity of the instrument and the spectrum of the background source and Galactic absorption, mean that 0.18 eV might not be a fully appropriate minimum EW for Ne IX and Fe XVII. Besides that, the relation between column density and EW has enough scatter above the minimum observable EW that it does not quite translate into a unique minimum column density, but it is an acceptable approximation in this regime (Wijers et al. 2019 appendix B).

We also make predictions for the proposed Arcus (Smith et al. 2016; Brenneman et al. 2016) mission, and the X-ray Grating Spectrometer (XGS) on the proposed Lynx mission (The Lynx Team 2018). For Arcus, we assume a minimum detectable EW of 4 mÅ (for 5σ detections). This is based on bright AGN background sources, which were selected to have a high flux between 0.5 and 2 keV (Brenneman et al. 2016), and exposure times < 500 ks (Smith et al. 2016). At least 40 blazars matching the brightness requirements are known (Smith et al. 2016). These estimates are based only on O VII and O VIII (and C VI), so this minimum EW may not apply to the Ne IX and Fe XVII lines at smaller wavelengths.

Note that Arcus not only aims to find weaker absorption lines than the Athena X-IFU, it is also meant to characterise them in more detail using its higher spectral resolution. Arcus has a $\approx 6\text{--}8\times$ higher spectral resolution than Athena at the wavelengths of O VII and O VIII at $z = 0.1$, which is sufficient ($\approx 120\text{--}150$ km s $^{-1}$, Smith et al. 2016) to determine if absorbers are associated with L_* galaxy haloes that have typical virial velocities of $150\text{--}300$ km s $^{-1}$, while the Athena X-IFU's resolution ($\approx 900\text{--}1000$ km s $^{-1}$, Barret et al. 2018) would be insufficient to determine if absorbers belong to individual galactic haloes.

For the Lynx XGS, the requirement is a detectable EW of 1 mÅ for O VII and O VIII (The Lynx Team 2018). This applies to 80 bright AGN background sources in a 5 Ms survey, focussed on detecting the CGM of $\sim L_*$ galaxies in absorption.

For Arcus and Lynx, we therefore limit our predictions to O VII and O VIII. The minimum EWs for Arcus translate to column densities of $10^{15.3}$ and $10^{15.6}$ cm $^{-2}$ for O VII and O VIII, respectively. For Lynx, the values are, respectively, $10^{14.6}$ and $10^{14.9}$ cm $^{-2}$.

For the FUV ions, we choose column densities based on what is currently observed with the Hubble Space Telescope's Cosmic Origins Spectrograph (HST-COS). We base estimates

on observed column densities and upper limits, and column densities used for covering fractions by observers. We use the data of Tumlinson et al. (2011) and Prochaska et al. (2011) for O VI, and of Burchett et al. (2019) and Meiring et al. (2013) for Ne VIII. Note that our limits are for $z = 0.1$ for consistency, but the EUV line we discuss for Ne VIII is only observable at higher redshifts. We explore the redshift evolution of the absorption in Appendix 3.C.

3.4.2 Halo detection rates

Based on the CDDFs for gas coming from haloes of different masses, we can estimate how many haloes of different masses should be detectable with the Athena X-IFU over a given total redshift path dz . Here, it does matter if we use the CDDFs based on the halo-projection method (such as in Fig. 3.6), or the pixel-attribution method, where we base the CDDF on column density maps including all gas, but only counting pixels with impact parameter $r_{\perp} \leq R_{200c}$, and that are not closer to another halo in r_{\perp} / R_{200c} units (see §3.2.7).

Using estimated minimum column densities of $10^{15.4} \text{ cm}^{-2}$ for O VII and $10^{15.6} \text{ cm}^{-2}$ for O VIII for the Athena X-IFU, we expect to find, in total, 2.3 O VII absorbers and 1.0 O VIII absorbers per unit redshift. Of those, 46 and 63 per cent are within R_{200c} of a central galaxy with $M_{200c} > 10^{11} M_{\odot}$, respectively. For O VII, 41 per cent of all the absorbers is attributed to haloes with $M_{200c} = 10^{12} - 10^{13.5} M_{\odot}$, and for O VIII, 53 per cent comes from $M_{200c} = 10^{12.5} - 10^{13.5} M_{\odot}$ haloes. (Since the halo-projection CDDFs do not add up to the all-gas CDDF, we do not attempt to derive such fractions from the halo-projection CDDFs.)

We also estimate the total density along lines of sight of observable absorbers coming from haloes of different masses. Here, both the halo-projection CDDFs and the pixel-attribution CDDFs are reasonable starting points. We expect to find 0.30 (0.34), 0.39 (0.61), and 0.26 (0.42) O VII absorbers per unit redshift with the Athena X-IFU blind survey based on the pixel-attribution CDDFs (based on the halo-projection CDDFs), in haloes of $M_{200c} = 10^{12} - 10^{12.5}$, $10^{12.5} - 10^{13}$, and $10^{13} - 10^{13.5} M_{\odot}$, respectively. For O VIII, we similarly expect 0.21 (0.20) and 0.31 (0.37) absorbers per unit redshift in M_{200c} bins of $10^{12.5} - 10^{13}$ and $10^{13} - 10^{13.5} M_{\odot}$, respectively. This is assuming most of the redshift path searched is close to $z = 0.1$, which is the redshift of the EAGLE snapshot we extracted the CDDFs from. We compare the absorption to $z = 0.5$ in Appendix 3.C.

Therefore, halo absorbers seem to be somewhat rare, but in a search for O VII and O VIII absorbers against 100 BLLacs and 100 gamma ray bursts (Lumb et al. 2017), we can reasonably expect to find quite a few of these absorbers. There are differences between the different determinations of halo absorption, but they are mostly not too severe. We have also assumed that a single limiting column density is a good estimate for detectability of these absorbers. We can see from Fig. 3.7 and fig. B1 of Wijers et al. (2019), that this should be a reasonable approximation for the limiting-case absorption lines.

3.4.3 Extent of detectable absorption

In Fig. 3.14, we investigate how far from halo centres this absorption occurs: we show the fraction of absorbers at different radii above the estimated minimum observable column density for the Athena X-IFU. Those minima are indicated in the panels. For O VI and Ne VIII, these are more optimistic (but achievable) minimum column densities for HST-COS, reflecting currently possible observations. Many of these have also been done. We discuss some EAGLE data comparisons in §4.5. For $\sim L_*$ galaxies, O VI should be widely observable according to EAGLE, and Ne VIII should be widely observed in the L_* -group mass range. The Burchett et al. (2019) observations (at $z \approx 0.5 - 1.1$) report Ne VIII absorbers and upper limits




Figure 3.14: Covering fractions around central galaxies of different stellar masses in EAGLE as a function of physical impact parameter at $z = 0.1$. The covering fraction is the fraction of sightlines with column densities larger than a threshold value at each impact parameter. The threshold column densities, in $\log_{10} \text{ cm}^{-2}$ units, are shown in the panels. Points on each curve mark the median virial radius in each M_{\star} bin, and horizontal lines show the central 80 per cent of virial radii in those bins. (Some are outside the range of the plot.) These covering fractions are for optimistic HST-COS detection limits (O VI, Ne VIII) and estimated Athena X-IFU limits (X-ray lines). The Athena limits are for observer-frame EWs of 0.18 eV (0.16 eV rest-frame). That is the expected minimum for 5σ detections, assuming 50 ks exposure times and a quasar background source with a 2–10 keV flux of $10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ and a photon spectral index $\Gamma = 1.8$ (Lumb et al. 2017).

around galaxies in a wide mass range ($M_{\star} \sim 10^9\text{--}10^{11} M_{\odot}$), which covers a large range in covering fractions in our predictions and makes a direct comparison to our figures difficult.

For the X-ray lines, Fig. 3.14 shows estimated detection limits for the Athena X-IFU. Detection prospects for CGM out to large distances (close to R_{200c}) look good for O VII and O VIII, and in a narrower mass range, for Fe XVII. Ne IX might, however, prove more difficult to detect at larger impact parameters. However, these limits are for blind detections, and hence conservative for targeted observations. It might thus be possible to find Ne IX counterparts to absorbers in e.g., O VIII, or to detect weaker lines by searching at known galaxy redshifts. For most of the X-ray ions, the detection thresholds are not very far from the CDDF breaks, so small changes in sensitivity could make large differences for detection prospects, as the difference between Figs. 3.11 and 3.14 shows.

Using the median column density profiles from Fig. 3.10 we find the radii where the covering fractions at the detection limits for the different instruments reach 0.5. Note that this is the covering fraction in annuli at, not circles of, these radii. These are shown in Table 3.6.

This confirms that with the Athena X-IFU, we should be able to find O VII and O VIII ab-

Table 3.6: Impact parameters r_{\perp} / pkpc, rounded to 10 pkpc, where the covering fraction $f(> N) = 0.5$ for different ions, galaxy masses, and column density limits in EAGLE at $z = 0.1$. The column density limits (indicated in the third row in $\log_{10} \text{ cm}^{-2}$) are calculated for the different ions and instruments as explained in the text. We estimate the limits for the Athena X-IFU (Athena), Arcus, and the Lynx XGS (Lynx). A dash ('-') means that the covering fraction is below 0.5 at all radii.

M_{\star} $\log_{10} M_{\odot}$	O VII			O VIII		
	Athena	Arcus	Lynx	Athena	Arcus	Lynx
	15.4	15.3	14.6	15.6	15.6	14.9
10.0–10.5	30	50	340	-	-	140
10.5–11.0	260	310	680	50	50	560
11.0–11.7	460	540	930	410	410	910

sorption out to close to R_{200c} for $M_{\star} > 10^{10.5} M_{\odot}$ galaxies. Arcus performs similarly; we note that compared to the Athena X-IFU, it does have a much higher spectral resolution for soft X-rays. Lynx should be sensitive to much weaker absorption lines, and should therefore be able to find absorption systems beyond the virial radii of $M_{\star} > 10^{10.5} M_{\odot}$ galaxies, and roughly up to the virial radii of $M_{\star} = 10^{10} - 10^{10.5} M_{\odot}$ galaxies. Indeed, one of the mission goals is to characterise the CGM of $\sim L_{\star}$ galaxies using these absorption lines (The Lynx Team 2018).

3.5 Discussion and comparison to previous works

All of the predictions reported in the previous sections are based on the EAGLE simulation. EAGLE reproduces a number of galaxy and diffuse gas properties, to differing degrees. At the highest halo masses we study ($M_{500c} \gtrsim 10^{13.5} M_{\odot}$), there are observations available to which to compare hot gas properties of the ICM. At lower masses, the gas properties are less well-constrained. Schaye et al. (2015) showed that EAGLE matches the relation between M_{500c} (X-ray) and I-band luminosity ($\sim M_{\star}$) well, but overestimates ICM gas fractions at fixed halo mass (from X-ray measurements) and soft X-ray luminosities at fixed X-ray spectroscopic temperature above 1 keV. Overall, the higher gas fractions mean we might overestimate column densities from these simulations.

Other authors have studied UV absorption in simulations, particularly CDDFs and absorption around L_{\star} to group-mass haloes. Rahmati et al. (2016, EAGLE) found that the O VI CDDF is somewhat underpredicted at the high column density end, though uncertainties in the observations and oxygen yields may mean the difference is not severe. They found the cosmic Ne VIII density agreed with that measured by Meiring et al. (2013), though the measurements and comparisons are somewhat uncertain.

Oppenheimer et al. (2016) made a comparison of the absorption in $\sim L_{\star}$ to group-mass halo zoom-in simulations using the EAGLE simulation code to COS-Halos (Tumlinson et al. 2011) and found similar trends with star formation rates, but overall somewhat too low O VI column densities in $\sim L_{\star}$ haloes. Oppenheimer et al. (2018), using similar EAGLE-based zooms, found that those higher O VI columns could be achieved though ionization by radiation from flickering AGN, in combination with non-equilibrium ionisation, but that the CDDF discrepancy is likely not resolvable this way. This radiation and non-equilibrium ionisation also affects other ions. The effect on O VII column densities is, however, small: they decrease by $\lesssim 0.1$ dex.

Nelson et al. (2018) compared their IllustrisTNG data to COS-Halos survey results for

O VI (Tumlinson et al. 2011), complemented by galaxy data, and to the eCGM survey data (Johnson et al. 2015), and found generally good agreement. For their calibrated TNG100-1 volume, the O VI CDDF might be too large at high column densities.

For O VI, there are long-standing difficulties in modelling observed absorbers (in combination with lower ions) due to the uncertain ionization mechanism (e.g., Werk et al. 2016). We find that, in EAGLE, the O VI is mostly collisionally ionized in the inner regions, but photo-ionized at $\gtrsim R_{200c}$, and that photo-ionized O VI is present in the inner regions of lower-mass haloes. In IllustrisTNG CIE is most important mechanism in the CGM (Nelson et al. 2018), as in EAGLE, assuming ionization equilibrium. However, the results of Oppenheimer et al. (2018) demonstrate that, at least for the CGM of $\sim L_*$ galaxies, such equilibrium assumptions may underestimate the effect of photo-ionization. They used non-equilibrium ionization, and a flickering AGN as an additional ionization source, in zoom simulations otherwise using the EAGLE code and physics. In NIHAO and VELA zoom simulations, Roca-Fàbrega et al. (2019) found a roughly equal mix of collisionally and photo-ionized O VI in $z = 0$, $M_{200c} = 10^{11} - 10^{12.6} M_{\odot}$ haloes. Overall, this supports the picture that observed O VI does not have one single origin.

Nelson et al. (2018 IllustrisTNG) also looked into O VII and O VIII absorption in the CGM. They found column densities that peak at similar masses as we find, and ion fraction trends with halo mass similar to ours except at the lowest masses we examine, though they measure the fractions in a somewhat different gas selection. Martizzi et al. (2019), also studying the TNG100-1 volume, look into the contributions of large-scale structures to the O VII and Ne IX CDDFs. Rather than halo contributions, they split the CDDFs into contributions from larger-scale cosmic web structures (knots, filaments, sheets, and voids), and find that absorption at high column densities mainly comes from knots and filaments, and that more collapsed structures contribute more as the column density of absorbers increases.

3.6 Conclusions

Using the EAGLE simulation, we investigate the contents and properties of the CGM of $M_{200c} > 10^{11} M_{\odot}$ haloes, and how they are probed by O VI (FUV), Ne VIII (EUV) and O VII, O VIII, Ne IX, and Fe XVII (X-ray) line absorption at $z = 0.1$. With future X-ray instruments like the Athena X-IFU, Arcus, and the Lynx XGS, we expect that some of these absorption lines can be used to study the hot CGM. The mass of this CGM phase in $\sim L_*$ and group-mass haloes is largely unconstrained by current observations, and differs in different cosmological simulations. Determining the mass and metal content of the hot CGM will therefore provide important constraints for our understanding of structure and galaxy formation.

For the baryons, gas, and metals in haloes, we find that:

- The CGM (non-star-forming gas) is the largest baryonic mass component within R_{200c} in EAGLE haloes at all masses we investigate ($M_{200c} > 10^{11} M_{\odot}$), and is particularly dominant at $M_{200c} \gtrsim 10^{13} M_{\odot}$ (Fig. 3.3).
- Within R_{200c} , the CGM (non-star-forming gas) also contains more oxygen than the ISM (star-forming gas) for all halo masses, though differences are small at $M_{200c} \lesssim 10^{12} M_{\odot}$. However, up to $M_{200c} \approx 10^{13} M_{\odot}$, stars contain most of the oxygen ejected by earlier stellar generations that remains within R_{200c} (Fig. 3.3).
- The ions we study mainly trace CGM gas at $10^{5.5} - 10^7$ K (Figs. 3.1 and 3.13), which

constitutes a large fraction of the non-star-forming gas within R_{200c} in $M_{200c} \approx 10^{12} - 10^{13.5} M_{\odot}$ haloes (Fig. 3.3).

- The mass ranges for which median column densities are highest (Fig. 3.8) are in line with simple predictions comparing the virial temperature with the temperature where the ion fraction peaks in CIE (Fig. 3.1). This is because these ions mainly trace gas at temperatures around the CIE peak in the volume-filling phase of the CGM (Fig. 3.13).
- In the inner CGM, these ions are all mainly collisionally ionized (although some O VI is photo-ionized), but close to R_{200c} , photo-ionization becomes relevant (Figs. 3.1 and 3.13). The combination of multi-phase gas, a temperature gradient in the volume-filling phase, and different ionization mechanisms means that the haloes (gas at $0.1 - 1 R_{200c}$) exhibit a larger diversity of ions than the single-temperature CIE model alone would predict (Fig. 3.4).

We note an interesting feature in the median column density profiles: the column densities do not always decrease as the halo-centric impact parameter increases. This occurs in haloes with T_{200c} above the CIE peak temperature of the absorbing ion, where we find peaks or ‘shoulders’ in the column density profiles (Fig. 3.8). This occurs because the temperature of the volume-filling phase declines towards the outskirts of the halo (Fig. 3.13), causing the ion fraction to be larger there. Despite the decline of gas density with radius, this leads to a ‘shell’ around the galaxy where most of the metals in a particular ionization state are found (Fig. 3.12). In projection, depending on the strength of the shell feature, this leads to a peak or flattening of the column density as a function of halo-centric radius, typically around R_{200c} .

When we examine absorption as in blind surveys, we find that:

- The CDDFs have shallow slopes at lower column densities and a ‘tail’ with a steep slope at high column densities (Fig. 3.5).
- For the X-ray ions, the high-column-density tail of the CDDF is produced mostly by CGM gas (Fig. 3.6): 70–80 per cent has an impact parameter $r_{\perp} < R_{200c}$ for a halo with $M_{200c} > 10^{11} M_{\odot}$.

Finally, we make predictions for observational detections:

- For most of these ions, column densities remain large out to $\sim R_{200c}$ in haloes where T_{200c} is around the CIE peak temperature for that ion (Figs. 3.8 and 3.1), and O VII and O VIII should be detectable with Athena that far out around $M_{\star} > 10^{10.5} M_{\odot}$ galaxies (Fig. 3.14). However, Fe XVII absorption is more centrally concentrated, and more confined to haloes than the other ions in general (Fig. 3.2).
- We expect that the Athena X-IFU can detect O VII absorption in 77 (59) per cent of sightlines passing central galaxies with stellar masses $M_{\star} = 10^{10.5-11.0} M_{\odot}$ ($M_{\star} > 10^{11.0} M_{\odot}$) within 100 pkpc. For O VIII, this is 46 (82) per cent. Hence, the X-IFU will probe covering fractions comparable to those detected with the Cosmic Origins Spectrograph for O VI.
- Ne IX and Fe XVII might prove more difficult to find in the CGM with Athena, because the (roughly estimated) minimum observable column densities for these ions are close to, or in, the high-column-density tail of the CDDF.

- At column densities expected to be detectable with the Athena X-IFU, some of the absorption lines will be saturated. However, the saturation is less severe than thermal line broadening would predict for the temperatures where the CIE ionization fractions peak (Fig. 3.7).
- Our set of ions is mostly suited to probe haloes of $M_{200c} \sim 10^{12} - 10^{13.5} M_{\odot}$, though O VI, and Ne VIII at $z = 0.5$, also probe lower halo masses, and Fe XVII also probes somewhat hotter haloes.

Acknowledgements

We would like to thank Rob Crain and Fabrizio Nicastro for their involvement in the beginning of the project. NAW and BDO would like to thank Akos Bogdan, Ralph Kraft, Randall Smith, and Alexey Vikhlinin for useful discussions. We thank Yakov Faerman and the anonymous referee for useful comments. We would additionally like to thank John Helly for programs we used to access EAGLE data (READ_EAGLE), and Volker Springel for the original version of the code we use to project particles onto a grid (HSM LANDPROJECT). Ali Rahmati helped NAW test the code we use to make projections and helped set up SPECWIZARD. We used the NUMPY (Oliphant 2006), SCIPY (Jones et al. 2001), H5PY (Collette 2013), and MATPLOTLIB (Hunter 2007) PYTHON libraries, and the IPYTHON (Pérez & Granger 2007) command-line interface. This paper is supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871158, project AHEAD2020. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

Data availability

The data in the figures will be shared on reasonable request to the corresponding author. The EAGLE data are available at <http://icc.dur.ac.uk/Eagle/database.php>; the galaxy and halo catalogues are documented in McAlpine et al. (2016) and the full simulation data are documented in The EAGLE team (2017).

Appendix

3.A Measuring column densities and equivalent widths

In Table 3.4 we parametrised the relation between column densities and rest-frame equivalent widths measured in a specific velocity range around the maximum-optical-depth position. Here, we explore how the relation depends on the velocity range over which both are measured. We parametrise this relation with the best-fitting b parameter (eqs. 3.4, 3.5 and 3.6), assuming a single (or doublet, for O VIII) Voigt profile.

Fig. 3.15 shows the dependence of the best-fitting b parameter on the velocity window used. Since structures and correlations on scales approaching half the box size ($\Delta v \approx$




Figure 3.15: The best-fitting b parameter (eqs. 3.4, 3.5, and 3.6) for the different ions’ absorption lines (Table 4.1 different colors) in EAGLE at $z = 0.1$ as a function of the line-of-sight velocity range within which the column density and EW are measured. The indicated velocities are the maximum differences relative to the highest-optical-depth pixel, which is half the full velocity range. The filled points indicate the velocity windows used for each ion in Fig. 3.7 and throughout the paper, and the dashed lines indicate the thermal broadening for each ion at the temperature where its CIE ionization fraction is largest.

3200 km s^{-1}) cannot be reliably measured in a periodic box, we only show best-fitting values up to $\Delta v = 1500 \text{ km s}^{-1}$.

For very small velocity ranges, b parameters are lower than a single absorption component at a typical temperature for high column density absorbers would give. (Low column density absorbers are unsaturated and the fit is therefore not sensitive to their line widths.) However, the smallest ranges are close to or below the FWHM for such line profiles, so we cannot expect those ranges to be reliable. We interpret the initial rise with velocity range to be due to the velocity window encompassing more absorption from a single absorber or multiple correlated absorbers as the range is increased, and reaching a plateau when all the correlated absorption is included.

For the UV ions (O VI and Ne VIII), this plateau is reached roughly at the velocity cut $\Delta v = \pm 300 \text{ km s}^{-1}$ that we used (filled green and orange points in Fig. 3.15). This cut was motivated by the observations of Tumlinson et al. (2011) and choices by Johnson et al. (2015) and Burchett et al. (2019), and is apparently also reasonable for our simulated systems.

For the X-ray ions, we wanted to make sure the velocity range was not too small to probe with the Athena X-IFU, but there was otherwise no clear choice. Based on Fig. 3.15, we chose a velocity window $\Delta v = \pm 800 \text{ km s}^{-1}$ for these (filled red, blue, purple, and cyan points in Fig. 3.15). This is large enough to be in the plateau region for these ions, but stays clear of the half box size.

Note that Wijers et al. (2019) found a larger best-fitting b parameter for O VIII using the same set of sightlines, but measuring column densities and EWs over the full sightlines. The difference is not due to the inclusion of damping wings in the spectra in this work, which makes very little difference. Instead, it is driven by a subset of high column density sightlines that contain two strong absorbers, resulting from an alignment of two particular high-mass haloes along the Z -axis of the simulation (our line-of-sight direction). This affects

best-fitting b parameters at Δv approaching half the box size. These haloes also affect the large- Δv best-fitting b parameters for Fe xvii and Ne ix.

3.B How to split the CDDFs

When splitting the CDDFs into contributions from absorption by haloes of different masses, we mostly considered only SPH particles in haloes in each chosen mass range (halo-projection method).

For comparison, we also used a method where we attributed a pixel to a halo mass bin by checking whether any central galaxy in the same 6.25 cMpc slice along the line of sight was within R_{200c} (projected) for each absorption system, and if so, which one was closest to that system in normalised impact parameter units of r_{\perp} / R_{200c} . However, low-mass haloes show flat absorption profiles and, depending on the ion in question, have virial temperatures well below the temperatures where that ion exists in CIE. Some of the absorption attributed to low-mass haloes in this pixel-attribution method will therefore actually be due to higher-mass haloes nearby. We also refer to this as the mask-split method, since the column density maps are split into contributions by different halo masses based on True/False array masks.

We show a comparison of which absorption is attributed to which halo using the halo-projection and the pixel-attribution methods in Fig. 3.16, for O viii. The top left panel is the same as Fig. 3.6 but using the pixel-attribution method. We see that the modal column density increases with halo mass up to the ‘break’ in the CDDF, after which modal column densities decrease again and distributions become flatter.

In the top, middle panel (‘all halo gas’), the grey line shows absorption from haloes at any mass, while the brown lines show CDDFs from all-gas (dashed) and halo-only (solid) column density maps, with only the contributions from pixels within R_{200c} of an $M_{200c} > 10^{11} M_{\odot}$ halo. The differences between these three methods of measuring how much absorption is due to haloes are relatively small at high column densities, especially between the solid and dashed brown lines, which represent the two main ways to define the CDDF coming from roughly within the virial radii of haloes. This means that the split between halo and extra-halo absorption is relatively robust. (Note that gas outside R_{200c} or other virial radius definitions may still be associated with haloes, so where exactly the line between CGM and IGM lies is not generally agreed on.)

In the panels for specific halo masses, we can similarly compare the dashed and solid lines of the colour matching the halo mass in the panel to the grey line to get a sense of how robust the map-split and FoF-only methods are for determining which absorption comes from haloes. The agreement between these methods clearly depends on the halo mass. The higher column density absorption projected within R_{200c} of the lowest-mass haloes is clearly mostly due to gas outside those haloes, showing that much of the absorption attributed to these haloes using the mask-split method is simply due to these haloes being in roughly the same place as whatever structure is causing the absorption.

As the halo mass increases, the different methods agree well at column densities at or above the modal column density for that mass. The difference between the grey and coloured solid lines shows that this is largely due to lower-column-density absorption from gas in the FoF groups of haloes that is outside R_{200c} in projection. This will be due, in part, to the FoF groups not agreeing exactly with the overdensity definition of haloes, but also because SPH particles centred inside R_{200c} may extend beyond that radius. The very low column densities likely result, at least in part, from ‘edge effects’, where a particular sightline probes (the edge of) one or very few SPH particles at the outer extent of the FoF group.




Figure 3.16: A comparison between the column density distribution arising from haloes in different mass ranges by two methods. In the first, the CDDF is constructed in the same way as from all gas, but only considering absorption from gas within haloes in each mass range (halo-projection). By within haloes, we mean FoF particles and anything within R_{200c} of the centre of potential of each halo. For the second method, the column density distribution attributed to a halo mass bin is determined by, for each pixel in the column density map, checking which halo, if any, matches the line-of-sight slice of the absorber, is within R_{200c} , and is closest in units of r_{\perp} / R_{200c} (mask-split or pixel-attribution). This plot shows the distributions for O VIII in EAGLE at $z = 0.1$. The black lines show the total CDDF in all panels. In the top left panel, dashed coloured lines show the part of the total CDDF attributed to the halo mass range corresponding to that colour (colour bar) using the mask-split method. The ‘all halo gas’ panel uses masks for all haloes $M_{200c} > 10^{11} M_{\odot}$, cutting out parts attributed to haloes with $10^9 M_{\odot} < M_{200c} < 10^{11} M_{\odot}$. The masks are applied to an absorption map including only gas in haloes (of any mass, including $< 10^9 M_{\odot}$). The brown lines in this panel come from combining the masks for all the haloes with $M_{200c} \geq 10^{11} M_{\odot}$. In the rest of the panels, the range in the top left corner indicates the range of halo masses ($\log_{10} M_{200c} / M_{\odot}$) for which the gas is included in the column density maps the CDDFs are derived from. The dashed line of the colour matching the mass range shows the contribution according to the mask-split method applied to the column density maps for all gas, and the coloured solid lines show CDDFs from only the gas in the indicated range attributed to the same halo mass range by the mask-split method. The grey lines in the panels for specific halo masses show the halo-gas-only CDDFs used in the main text (Fig. 3.6).

Differences between the solid and dashed lines show absorption attributed to haloes (i.e. within R_{200c} in the plane of the sky), but not due to gas within the FoF group. This tends to lead to slightly more high-column-density absorbers, but the larger effect is typically at lower column densities, where gas outside the FoF groups is more important (Fig. 3.8).

For the other ions, the picture is very similar, except that the mass above which the methods agree at high column densities changes: the largest mass bin with disagreements at most about as bad as $M_{200c} = 10^{11.5} - 10^{12} M_{\odot}$ for O VIII is:

- $M_{200c} = 10^{10.5} - 10^{11} M_{\odot}$ for O VI (not shown here),
- $M_{200c} = 10^{11} - 10^{11.5} M_{\odot}$ for Ne VIII and O VII, and
- $M_{200c} = 10^{11.5} - 10^{12} M_{\odot}$ for Ne IX, Fe XVII (and O VIII).

Above those ranges, we consider the contributions of these ions to the CDDF to be fairly robust. The limiting mass for reliability increases with the CIE temperature range of the ions.

In summary, the attribution of absorption to haloes shown in Fig. 3.6 is fairly robust for higher-mass haloes ($> 10^{12} M_{\odot}$, and lower for some ions), for column densities above the peak of the CDDF for each M_{200c} bin. However, at lower column densities, the CDDF for a set of haloes depends quite strongly on how absorption and haloes are connected. The fraction of absorption beyond the CDDF breaks due to haloes does depend on these choices somewhat, but they do not change the qualitative conclusions.

3.C Redshift evolution

In Fig. 3.17 we investigate the redshift evolution of halo radial profiles. This is relevant for all the ions discussed in this work, but we focus on O VII and O VIII, which should be the most easily observed X-ray lines in this sample, and Ne VIII, which was observed at $z > 0.48$, but is not observable at $z = 0.1$ (Burchett et al. 2019). We compare column densities measured through the same comoving slice thickness (6.25 cMpc) at both redshifts.

The median changes are generally small ($\lesssim 0.2$ dex) at the halo masses where the column densities are largest, and within the range of the 80 per cent scatter. The redshift evolution is larger at larger and smaller halo masses, however. The changes are in line with how T_{200c} evolves: at the same M_{200c} , T_{200c} is somewhat larger at $z = 0.5$ than $z = 0.1$. Note that for a fixed density profile, M_{200c} and R_{200c} at the two redshifts will also differ. The differences do mean that any comparisons of absorbers to data should be done at matching redshifts.

Fig. 3.18 we similarly consider the evolution of the CDDFs of these ions from redshift 0.1 to 0.5. We see the CDDFs change in ways consistent with the radial profiles: the distributions look similar, but at $z = 0.5$, they look like those for somewhat higher-mass $z = 0.1$ haloes. Another difference is in the total gas CDDF. All haloes together contribute less to the high-column-density absorption at $z = 0.5$ than at $z = 0.1$. This might be related to how the FoF groups are defined: by particle separation relative to the average, which means that haloes will be ‘cut off’ at higher densities at the higher redshift. Overall evolution of the EAGLE O VII and O VIII CDDFs was discussed in more detail by Wijers et al. (2019), and Rahmati et al. (2016) discussed the Ne VIII CDDF evolution, as well as that of O VI.

Note that the highest mass bin only has 5 haloes at $z = 0.5$, and 9 at $z = 0.1$, so some changes here might be explained by selection effects (lower typical masses at higher redshift) and small sample sizes.




Figure 3.17: A comparison of the radial profiles for O VII, O VIII, and Ne VIII in EAGLE between redshift 0.1, as used throughout this work, and redshift 0.5. The mass bins are split over top and bottom panels for legibility. We see that at halo masses where the column densities peak, the median column density is similar at the two redshifts, but the peak range may be larger at $z = 0.5$, and the profiles at higher and lower halo masses change more. The central 80 per cent scatter is shown for the peak halo mass at $z = 0.1$ with the shaded bands.




Figure 3.18: A comparison between the column density distributions of Ne VIII, O VII, and O VIII in EAGLE at $z = 0.1$ (solid and dashed lines) and $z = 0.5$ (dotted and dash-dotted lines). The CDDFs for different halo mass ranges are derived from column density maps using only gas in FoF groups or within R_{200c} of haloes of the different masses. The colours match those of Fig. 3.6, but the lines for different halo masses are split over top and bottom panels for legibility.

4

The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray line emission

Nastasha A. Wijers, & Joop Schaye
2021, submitted to MNRAS

ABSTRACT

We estimate the detectability of X-ray metal-line emission from the circumgalactic medium (CGM) of galaxies over a large halo mass range ($M_{200c} = 10^{11.5} - 10^{14.5} M_{\odot}$) using the EAGLE simulations. With the XRISM Resolve instrument, a few bright (K- α or Fe L-shell) lines from $M_{200c} \gtrsim 10^{13} M_{\odot}$ haloes should be detectable. Using the Athena X-IFU or the Lynx Main Array, emission lines (especially from O VIII and O VII) from the inner CGM of $M_{200c} \gtrsim 10^{12.5} M_{\odot}$ haloes become detectable, and intragroup and intracluster gas will be detectable out to the virial radius. With the Lynx Ultra-high Resolution Array, the inner CGM of haloes hosting L_* galaxies is accessible. These estimates do assume long exposure times (~ 1 Ms) and large spatial bins ($\sim 1-10$ arcmin²). We also investigate the properties of the gas producing this emission. CGM emission is dominated by collisionally ionized (CI) gas, and tends to come from halo centres. The gas is typically close to the maximum emissivity temperature for CI gas (T_{peak}), and denser and more metal-rich than the bulk of the CGM at a given distance from the central galaxy. However, for the K- α lines, emission can come from hotter gas in haloes where the virialized, volume-filling gas is hotter than T_{peak} . Trends of emission with halo mass can largely be explained by differences in virial temperature. Differences between lines generally result from the different behaviour of the emissivity as a function of temperature of the K- α , He- α -like, and Fe L-shell lines. We conclude that upcoming X-ray missions will open up a new window onto the hot CGM.

4.1 Introduction

In order to understand the formation, evolution, and quenching of galaxies, we must understand the gas that surrounds them: the circumgalactic medium (CGM). This is because, first, accretion of gas from the CGM fuels star formation in star-forming galaxies. Without the additional gas supply, star-forming galaxies would deplete their (galactic) gas reservoirs on timescales too short to be consistent with their star formation histories (e.g., the CGM review by Tumlinson et al. 2017). In numerical simulations, cutting off this fuel supply is one way to quench star formation in galaxies (e.g., Oppenheimer et al. 2020a; Zinger et al. 2020). Secondly, galaxies inject mass and metals back into the CGM. Outflows from galaxies into the CGM have been observed (e.g., the review by Rupke 2018), and are required to explain the presence of metals in the diffuse intergalactic medium (IGM) (e.g., Aguirre et al. 2001; Booth et al. 2012). In numerical simulations, these outflows (driven by e.g., supernovae and AGN) are required to reproduce the galaxy stellar mass function (e.g., Cole 1991; Hopkins et al. 2014; Crain et al. 2015; Pillepich et al. 2018). Therefore, observing the CGM can teach us about the processes that drive, regulate, and quench star formation.

Around isolated galaxies at low redshift, especially at $\sim L_*$, much of what we know about the CGM comes from studies of UV absorption lines, often using the *Hubble Space Telescope's Cosmic Origins Spectrograph* (HST-COS) (e.g., Tumlinson et al. 2011; Johnson et al. 2015, 2017). These lines mainly trace cool to warm ($\sim 10^4$ – $10^{5.5}$ K) gas.

A few ions producing UV lines trace warmer gas: O VI is sensitive to $\sim 10^{5.5}$ K gas if it is collisionally ionized. However, simulations (e.g., Tepper-García et al. 2011; Rahmati et al. 2016; Oppenheimer et al. 2016, 2018; Roca-Fàbrega et al. 2019; Wijers et al. 2020) and observations (e.g., Carswell et al. 2002; Tripp et al. 2008; Werk et al. 2014, 2016) alike suggest both collisionally ionized and photo-ionized O VI is present in the CGM, and causes measurable absorption lines, complicating the interpretation of observations. The Ne VIII ion produces a doublet in the extreme UV (EUV) range, but at redshift $\gtrsim 0.5$, it redshifts into the far UV (FUV) energy band and can be observed. This ion has been used to study the hotter CGM ($\sim 10^5$ – 10^6 K), by e.g., Burchett et al. (2019) (observationally) and Tepper-García et al. (2013) (in a cosmological simulation).

However, much of the gas in the CGM of low-redshift galaxies is expected to be at higher temperatures ($> 10^{5.5}$ – 10^6 K, e.g., Wijers et al. 2020). We expect this hot gas to be present around $\sim L_*$ and more massive galaxies ($\log_{10} M_{200c} M_{\odot}^{-1} \gtrsim 11.5$ – 12.0), where a volume-filling, virialized gas phase has formed (e.g., Dekel & Birnboim 2006; Kereš et al. 2009; Van de Voort et al. 2011; Correa et al. 2018). Aside from some EUV lines, this gas produces most of its emission (e.g., Bertone et al. 2013) and absorption (e.g., Perna & Loeb 1998; Hellsten et al. 1998) lines in X rays.

Therefore, X-ray emission and absorption are one way we might detect the hot phase of the CGM. X-ray emission and absorption lines in this hot gas come from metals, as hydrogen and helium are fully ionized at such high temperatures. Besides these lines, the warm-hot gas produces X-ray continuum emission. Bertone et al. (2013) predicted that most X-ray emission from diffuse gas throughout the universe is in the form of continuum emission. However, around isolated ellipticals and in groups, X-ray emission is typically line dominated (e.g., the review by Werner & Mernier 2020).

Other observables of the warm/hot CGM include the dispersion measures of fast radio bursts (FRBs). These are sensitive to the free electron density along the line of sight, although they only measure the total electron column density (e.g., Prochaska & Zheng 2019), meaning they are equally sensitive to photo-ionized, cool gas and warm-hot, collisionally ionized gas.

They can be used to constrain the ionized gas content of haloes, but this requires a sufficiently large sample of FRBs with known redshifts (Ravi 2019).

The Sunyaev-Zel'dovich (SZ) effect is also sensitive to free electrons: the thermal SZ effect probes the electron pressure along the line of sight, and the kinetic SZ effect measures the electron bulk velocity. These effects have primarily been used to study clusters (e.g., the review by Mroczkowski et al. 2019). However, tSZ signals from massive filaments have also been detected by stacking pairs of massive galaxies (e.g., de Graaff et al. 2019; Tanimura et al. 2019), and studies of lower mass haloes have been done by fitting models, using the known positions of galaxy groups (e.g., Lim et al. 2018, 2020). Both kinetic and thermal SZ signals from low mass systems are difficult to study with current instruments due to their large beam size (spatial resolution; e.g., Mroczkowski et al. 2019).

Around massive galaxies and in groups and clusters, X-ray emission from the CGM, intra-group medium (IGrM) and intra-cluster medium (ICM) has been detected (e.g., the review by Werner & Mernier 2020). For isolated galaxies, these detections are mostly limited to massive (elliptical) galaxies. For lower-mass, spiral galaxies, studies have typically found upper limits or emission only in or close to galaxies (e.g., Bogdán et al. 2015). However, Das et al. (2020) found emission further from the galaxy, and measured a temperature profile out to ≈ 200 kpc.

Another exception is the Milky Way. The halo of our own Galaxy has been studied using X-ray line emission, often in combination with X-ray absorption lines (e.g., Bregman & Lloyd-Davies 2007; Gupta et al. 2014; Miller & Bregman 2015; Das et al. 2019). Other studies focussed on absorption lines (e.g., Kuntz & Snowden 2000; Hodges-Kluck et al. 2016; Gatuzz & Churazov 2018). These measurements have been used to constrain e.g., the hot phase temperature (e.g., Kuntz & Snowden 2000; Das et al. 2019) and halo rotation (Hodges-Kluck et al. 2016). It is not certain how extended the gas causing the absorption and emission is (e.g., Bregman & Lloyd-Davies 2007; Gatuzz & Churazov 2018), though Miller & Bregman (2015) placed some constraints on the density profile and metallicity using a combination of O VII and O VIII absorption and emission lines.

X-ray emission has been very useful in the study of the ICM. From spectra, the temperature, electron density, and element abundances (using the ratio of emission lines and continuum) of the X-ray emitting phase have been measured (e.g., the review by Werner & Mernier 2020). Turbulence has also been measured, using resonant scattering emission lines (e.g., the review by Werner & Mernier 2020), spatially resolved emission line profiles (e.g., Hitomi Collaboration et al. 2018), and surface brightness fluctuations (e.g., Zhuravleva et al. 2014).

Extending such studies toward lower halo masses would be very valuable. The mass of the CGM around e.g., L_* galaxies is very uncertain, especially the mass of the warm/hot gas (e.g., Werk et al. 2014 fig. 11). Theoretical predictions also differ hugely: for example, the EAGLE and IllustrisTNG simulations predict very different CGM gas masses in $\lesssim L_*$ haloes (Davies et al. 2020), even though both produce broadly realistic galaxy populations (e.g., Schaye et al. 2015; Pillepich et al. 2018). Oppenheimer et al. (2020b) have found that broad-band X-ray emission from the EAGLE and IllustrisTNG L_* inner CGM should be observable with eRosita, and that this instrument should be able to distinguish between the two models.

Many predictions of X-ray emission from hot haloes from numerical simulations have focussed on groups and clusters (e.g., Barnes et al. 2017; Truong et al. 2018; Cucchetti et al. 2018; Mernier et al. 2020), where much of the data is currently available and high photon counts will allow detailed information to be extracted. For EAGLE, Schaye et al. (2015) have studied X-ray emission from groups and clusters, and Davies et al. (2019) considered broad-band soft

X-ray emission over a large range of halo masses and found it to be a good diagnostic of the CGM gas mass at fixed halo mass. [Bogdán et al. \(2015\)](#) made predictions of broad-band X-ray emission from the Illustris simulation ([Vogelsberger et al. 2014](#)) and compared them to data, finding them to be broadly consistent. [Zhang et al. \(2020\)](#) made predictions of X-ray emission for HUBS ([Cui et al. 2020a,b](#)) observations across a large range of halo masses using IllustrisTNG ([Pillepich et al. 2018](#)), and [Truong et al. \(2020\)](#) related hot gas properties close to the central galaxy to the central galaxy properties in IllustrisTNG. [Van de Voort & Schaye \(2013\)](#) made predictions for X-ray line emission specifically, using the OWLS simulations ([Schaye et al. 2010b](#)). We update these predictions and extend them to a larger set of emission lines using EAGLE.

In this paper, we will study low-redshift ($z = 0.1$) X-ray emission lines as predicted using the EAGLE simulations. We describe the simulations and how we use them to predict line emission in §4.2. We select a number of the stronger emission lines we expect to find (§4.2.2), and compare them to estimated sensitivity limits of various planned and proposed X-ray telescopes (§4.4). We describe how we estimate those sensitivity limits in §4.3. We also investigate the gas responsible for the emission and how it compares to typical CGM gas in §4.4. In §4.5, we discuss our results, and we summarise them in §4.6. For a similar study of X-ray and highly ionized UV *absorption* lines (O VI–VIII, Ne VIII, Ne IX, and Fe XVII) in the CGM of EAGLE galaxies, see [Wijers et al. \(2020\)](#).

Note that we will often use ‘CGM’ or ‘halo’ as a catch-all term for what is typically called the CGM (gas around isolated galaxies), as well the IGrM and the ICM in the few clusters in the 100^3 cMpc³ EAGLE volume. We describe distances as comoving (e.g., ‘cMpc’) or proper/physical (e.g., ‘pkpc’), except for centimetres, which are always physical.

4.2 Methods

In this section, we will discuss the cosmological simulations we use to make our predictions, how we predict surface brightnesses from them, and the galaxy and halo information we use.

4.2.1 EAGLE

We study line emission using the EAGLE (‘Evolution and Assembly of GaLaxies and their Environments’; [Schaye et al. 2015](#); [Crain et al. 2015](#); [McAlpine et al. 2016](#)) cosmological, hydrodynamical simulations. Specifically, we use the Ref-L0100N1504 100^3 cMpc³ volume, with an initial gas mass resolution of $1.81 \times 10^6 M_\odot$ and a Plummer-equivalent gravitational softening length of 0.70 pkpc (at low redshift, like we study here). EAGLE uses a modified version of Gadget3 ([Springel 2005](#)); gravitational forces are calculated using the Tree-PM scheme, and hydrodynamical forces are calculated using a pressure-entropy formulation of SPH known as Anarchy ([Schaye et al. 2015](#) appendix A; [Schaller et al. 2015](#)).

Besides gravity and hydrodynamics, EAGLE also models the effects of processes that occur on scales below its resolution: so-called subgrid physics. Radiative cooling and heating is modelled as described by [Wiersma et al. \(2009a\)](#), including the effects of 9 metal abundances tracked in EAGLE. Because the resolution is too low to model the multi-phase ISM, molecular cooling and heating channels are not included, and artificial fragmentation of the interstellar medium (ISM) is prevented by setting a pressure floor in dense gas that ensures the Jeans mass remains marginally resolved ([Schaye et al. 2015](#)). This means the temperature

of star-forming gas is set by the pressure floor equation of state, and is generally not typical of what we expect for the ISM.

Stars form stochastically in sufficiently dense gas, with a threshold that depends on the gas metallicity (Schaye 2004). The star-formation rate depends on pressure in a way that, by design, reproduces the Kennicutt-Schmidt relation (Schaye & Dalla Vecchia 2008). Feedback from these stars is modelled as well. Core-collapse supernovae inject thermal energy stochastically into neighbouring gas particles (Dalla Vecchia & Schaye 2012). The thermal energy injection raises the gas temperature by $10^{7.5}$ K, with a probability set to match the (calibrated) supernova energy budget per unit stellar mass. Core-collapse supernovae, as well as AGB winds and type Ia supernovae, inject mass and metals into the surrounding gas, with metal yields from 9 individual elements following Wiersma et al. (2009b).

Black holes are seeded in sufficiently massive haloes that do not already contain them. They can merge and accrete gas following Rosas-Guevara et al. (2015). Black holes generate AGN feedback by thermal energy injection (Booth & Schaye 2009), like supernovae, but raise the gas temperature by $10^{8.5}$ K.

Because the way the feedback energy couples to gas on scales resolved in EAGLE is still uncertain, the feedback on resolved scales is calibrated to produce realistic galaxies. The supernova and black hole feedback is calibrated to match the $z = 0.1$ galaxy luminosity function and stellar-mass-black-hole-mass relation, and to produce reasonable galaxy sizes (Crain et al. 2015). The EAGLE simulation data has been publicly released, as described by McAlpine et al. (2016) and The EAGLE team (2017).

4.2.2 The emission lines

The default tables

We will describe the luminosity and surface brightness of the CGM in EAGLE for a set of soft X-ray emission lines. The basis for our selection of X-ray lines is a set of tables from Bertone et al. (2010a), and the lines studied in that work. These tables were calculated using CLOUDY v7.02 (last documented in Ferland et al. 1998), assuming a Haardt & Madau (2001) uniform, but redshift-dependent, UV/X-ray ionizing background. Note that this means that, when calculating the emission from a patch of gas, we ignore contributions to the incident radiation field from, e.g., nearby AGN or ICM. This is consistent with the radiative cooling used in the EAGLE simulation.

Following Charlotte Brand (private communication, 2017), our lines were selected to have peak emissivities in dense gas between 10^5 and 10^7 K, as this is the warm-hot gas phase we want to investigate. The lines have energies > 0.3 keV, based on absorption by our own Galaxy (see e.g., the left panel of Fig. 4.2).

The lower line energy limit constrains the emissivity peaks of these lines to be at $\gtrsim 10^6$ K (Fig. 4.1), while higher-energy lines such as the Si xiv Lyman- α -like (K- α) line and the S xv He- α -like recombination line are excluded based on their temperature peaks. In addition to the selection of bright lines from Bertone et al. (2010a), we also include more iron L-shell lines and the Mg xi He-like resonance line. Similarly, the Si ix- xii ions produce a few emission lines between 0.3 and 0.4 keV, the brightest of which matches the peak emissivity of the Si xiii He-like resonance line. However, this is still fainter than the carbon lines in this energy range. We choose to focus instead on the He- α -like and Lyman- α -like transitions, and a set of relatively bright iron L-shell lines. Following Bertone et al. (2010a), we mostly consider only the resonance line for He- α triplets, except for the brightest one, O vii. There,

we consider the forbidden (f) and intercombination (x and y) lines as well. We list the lines we study in Tables 4.1 and 4.2

We looked up the transitions for the bright iron L-shell lines from CLOUDY 7.02 in the Opacity Project database¹, described as the source of the L-shell lines (included in CLOUDY v7.02 via level2.dat). The exception is the Fe xvii 17.05Å transition, which we could not find in that database. We found the data for that transition by comparing the wavelength (and checking the weighted oscillator strength gf) to the line compilation of Mewe & Gronenschild (1981) and the lines in the NIST database² (Shorer 1979; Gordon et al. 1980).

The Lyman- α and He- α transitions are calculated internally in CLOUDY v7.02. We looked up the He-like transitions in the documentation (LineList_{HeLike}.dat). The Lyman α transitions include both doublet components (all $2p - 1s$ transitions). The $2s^3S - 1s^2^1S_0$ transition is the He-like forbidden line, $2p^3P - 1s^2^1S_0$ are the He- α -like intercombination lines (sum of the x and y lines), and $2p^1P - 1s^2^1S_0$ is the He- α -like resonance line. Porquet & Dubau (2000) give a helpful overview of these He-like transitions in their fig. 1.

The Fe L-shell lines

For the bright iron L-shell lines we investigated, the Bertone et al. (2010a) tables in the previous section contain an error: at densities $n_H \lesssim 10^{-4} \text{ cm}^{-3}$ (for Fe xvii 17.05Å) and $\lesssim 10^{-3} - 10^{-2} \text{ cm}^{-3}$ (for the other L-shell lines), there is no tabulated emission at all from these lines. The tabulated quantity is $\log_{10} L V^{-1} n_H^{-2}$, where L is luminosity, V is volume, and n_H is the hydrogen number density. This means that, in the collisionally ionised (CI) limit, the tabulated quantity should not depend on density. The drop in tabulated emissivity is a very sharp transition from emission in the CI limit to zero, so the drop is not a physically consistent decrease of emission to (floating point) zero at low densities. This behaviour is not the result of some physical process.

This bug also affects the Fe xvii 17.05Å analysis of Bertone et al. (2010a), as can be seen in e.g., their fig. 13. There is some slightly lower-density gas producing Fe xvii 17.05Å emission in that plot, at $z = 0.25$, than we find at $z = 0.1$. However, the lack of any emission at all at temperatures and densities where there is emission in the other metal lines indicates that there is a similar issue at this redshift.

In order to make predictions for these lines, we therefore use a different set of tables: those of Ploeckinger & Schaye (2020). We use their default tables, which include the effects of dust, self-shielding, cosmic rays, and local stellar radiation. However, these dense gas/ISM processes and effects are negligible for the X-ray lines we study here. The main differences for this work are that the Ploeckinger & Schaye (2020) tables were made with a newer CLOUDY version (v17.01, Ferland et al. 2017), including updated atomic data, and that these tables were calculated assuming a Faucher-Giguère (2020) UV/X-ray background³.

For the Ploeckinger & Schaye (2020) table lines, we look up the transitions for H-like and He-like species in the CLOUDY documentation like we did for the other tables; the wavelengths of these lines very closely match those of the earlier version. For the Fe L-shell lines, the transitions selected for the two table sets are different. The atomic data for these lines are more uncertain (e.g., Gu et al. 2007; de Plaa et al. 2012; Bernitt et al. 2012; Wu & Gao 2019; Gu et al. 2019), so slightly different wavelengths and other atomic data in different transition

¹<http://cdsweb.u-strasbg.fr/topbase/topbase.html>

²<https://physics.nist.gov/asd> accessed 2020-09-24

³Ploeckinger & Schaye (2020) modified this background somewhat for redshifts $z > 3$, but we do not consider such high redshifts in this work.

Table 4.1: Data for the lines from the Bertone et al. (2010a) tables we investigate in this work: the ion producing the line, its wavelength and energy (rest-frame), the peak line emissivity for collisional ionization equilibrium (CIE) conditions ($n_H = 10 \text{ cm}^{-3}$), using solar metal abundances (Table 4.3), the temperature at the CIE emission peak, and the temperature range in which the emissivity is at least 0.1 times the maximum value. The peak temperatures are the maxima obtained directly from the emissivity tables, which use a 0.05 dex temperature grid. The final three columns list data for comparing these lines to other works: the electron configurations of the upper and lower levels of the transition, and names of the lines in CLOUDY v7.02. We substitute underscores for spaces in the names. Note that the line energies we list here are derived from the CLOUDY v7.02 wavelengths; they are listed for convenience.

ion	λ Å	E keV	max $\Lambda n_H^{-2} V^{-1}$ $\log_{10} \text{ erg cm}^3 \text{ s}^{-1}$	T_{peak} $\log_{10} \text{ K}$	$T_{0.1 \times \text{peak}}$ $\log_{10} \text{ K}$	upper level	lower level	name
C V	40.27	0.3079	-24.4	5.95	5.7-6.3	1s 2p ¹ P	1s ² ¹ S ₀	C__5_40.27A
C VI	33.74	0.3675	-24.1	6.15	5.9-6.8	2p	1s	C__6_33.74A
N VI	28.79	0.4307	-24.7	6.15	5.9-6.5	1s 2p ¹ P	1s ² ¹ S ₀	N__6_28.79A
N VII	24.78	0.5003	-24.4	6.3	6.1-7.0	2p	1s	N__7_24.78A
O VII	21.60	0.5740	-23.9	6.3	6.0-6.7	1s 2p ¹ P	1s ² ¹ S ₀	O__7_21.60A
O VII	21.81	0.5685	-24.4	6.35	6.0-6.7	1s 2p ³ P	1s ² ¹ S ₀	O__7_21.81A
O VII	22.10	0.5610	-23.9	6.35	6.0-6.7	1s 2s ³ S	1s ² ¹ S ₀	O__7_22.10A
O VIII	18.97	0.6536	-23.6	6.5	6.2-7.2	2p	1s	O__8_18.97A
Ne IX	13.45	0.9218	-24.4	6.6	6.3-7.0	1s 2p ¹ P	1s ² ¹ S ₀	Ne_9_13.45A
Ne x	12.14	1.021	-24.2	6.8	6.5-7.5	2p	1s	Ne10_12.14A
Mg XI	9.169	1.352	-24.8	6.8	6.4-7.2	1s 2p ¹ P	1s ² ¹ S ₀	Mg11_9.169A
Mg XII	8.422	1.472	-24.6	7.0	6.7-7.8	2p	1s	Mg12_8.422A
Si XIII	6.648	1.865	-24.8	7.0	6.6-7.4	1s 2p ¹ P	1s ² ¹ S ₀	Si13_6.648A

Table 4.2: Data for the [Ploeckinger & Schaye \(2020\)](#) table lines we use in this work, analogous to the [Table 4.1](#) data. For the peak emissivity, the metallicity of the gas is scaled to the solar values of CLOUDY v7.02 ([Table 4.3](#)). The peak temperatures are the maxima directly from the emissivity tables, which use a 0.1 dex temperature grid. The line identifications/transitions come from the CHIANTI database (versions 7.0 and 10.0.1) used in CLOUDY v17.01 for these iron lines. Note that the line energies we list here are derived from the CLOUDY wavelengths; they are listed for convenience. The transition attributed to the Fe XVII 17.05Å line in the CLOUDY v7.02 data produces a line at 16.7757 Å in the CHIANTI database used to match the Fe L-shell lines in CLOUDY v17.01. However, we believe the comparison between both lines at 17.05 Å is more like-for-like (see the text for discussion).

ion	λ Å	E keV	max $\Delta n_{\text{H}}^{-2} V^{-1}$ \log_{10} erg cm ³ s ⁻¹	T _{peak} K	T _{0.1xpeak} K	upper level	lower level	name CLOUDY v17.01
Fe XVII	17.0960	0.7252	-24.1	6.7	6.3-7.0	2s ² 2p ⁵ 3s ³ P ₂	2s ² 2p ⁶ 1S ₀	Fe17 _____ 17.0960Å
Fe XVII	17.0510	0.7271	-24.0	6.7	6.3-7.0	2s ² 2p ⁵ 3s ¹ P ₁	2s ² 2p ⁶ 1S ₀	Fe17 _____ 17.0510Å
Fe XVII	16.7760	0.7391	-24.2	6.7	6.4-7.0	2s ² 2p ⁵ 3s ³ P ₁	2s ² 2p ⁶ 1S ₀	Fe17 _____ 16.7760Å
Fe XVII	15.2620	0.8124	-24.3	6.8	6.4-7.0	2s ² 2p ⁵ 3d ³ D ₁	2s ² 2p ⁶ 1S ₀	Fe17 _____ 15.2620Å
Fe XVIII	16.0720	0.7714	-24.9	6.8	6.5-7.1	2s ² 2p ⁴ (³ P) 3s ⁴ P _{5/2}	2s ² 2p ⁵ 2P _{3/2}	Fe18 _____ 16.0720Å

databases and CLOUDY versions are to be expected. We list the Fe L-shell lines we study in this work in Table 4.2.

For these L-shell lines, CLOUDY 17.01 uses the CHIANTI database⁴ data (version 7.1.4; Dere et al. 1997; Landi et al. 2013) by default. These default settings were used in Ploeckinger & Schaye (2020). We looked up the transitions using the version 7.0 line list provided on the CHIANTI website. We also checked the H- and He-like transition identifications against this list, and found they matched.

However, there was an issue with the Fe L-shell line identifications. The Mewe & Gronenschild (1981) wavelength and transition identification combinations do match quite well for the Fe xvii 17.10 Å and 15.26 Å lines, and the Fe xviii 16.07 Å line. The transitions causing the Fe xvii 16.78 Å and 17.05 Å lines are, however, reversed between the Mewe & Gronenschild (1981) and CHIANTI v7.0 line lists. The NIST database⁵ (Gordon et al. 1980; Hutcheon et al. 1976) seems to agree with the Mewe & Gronenschild (1981) classifications, while the latest CHIANTI database (version 10.0.1 Dere et al. 1997; Del Zanna et al. 2021) agrees with the earlier Chianti version. There may therefore be errors in the classification of the Fe L-shell lines, and it is not entirely clear whether the CLOUDY v7.02 Fe xvii 17.05 Å line is ‘the same’ as the CLOUDY v17.01 Fe xvii 17.0510 Å line.

The transition probabilities from the NIST and CHIANTI (version 7.1.4 or 10.0.1) databases do not match exactly, but they match the same-wavelength lines better than the ones identified with the same transitions. Therefore, when comparing the results from the two sets of tables we use, we will assume that the two 17.05 Å lines are ‘the same’.

Line luminosities

To calculate the line luminosity for each gas (SPH) particle, we use tables which tabulate gas luminosity. The Bertone et al. (2010a) tables list $\log_{10} L V^{-1} n_{\text{H}}^{-2}$, the luminosity per unit volume and squared hydrogen number density, as a function of $\log_{10} T$, $\log_{10} n_{\text{H}}$, and z , where T is the temperature and z the redshift. The n_{H}^2 factor accounts for the first-order dependence of emission on the collision (and therefore, excitation) rate, which scales as $L V^{-1} \propto n_{\text{ion}} n_{\text{e}}$ in collisional ionization equilibrium (CIE), where n_{ion} and n_{e} are the number densities of the ions producing the line and electrons, respectively. This emission is interpolated linearly, in log space where applicable. For each particle, we multiply the table values by n_{H}^2 and volume (mass divided by density) to get the line luminosity.

The Ploeckinger & Schaye (2020) tables list $\log_{10} L V^{-1}$, without the first-order dependence on hydrogen number density scaled out, and these values are a function of the total metallicity $\log_{10} Z/Z_{\odot}$. Again, we interpolate these tables linearly. We multiply by the particle volume (mass over SPH density) to obtain the SPH particle luminosity.

The line emission also depends on the abundance of the species producing the line. This dependence is linear to high accuracy: twice as many atoms of a given element mean double the number of the ions responsible for the emission, and each ion will experience the same number of excitations. (Metal ions and the electrons from these metals only make small contributions to the excitation rates.) Therefore, we scale the emission of each SPH particle by the ratio of the particle’s element abundance to the solar abundance that the tables were made for (tables from Bertone et al. 2010a) or by the ratio of the SPH particle’s element abundance to the assumed element abundance at the particle’s total metallicity (tables from Ploeckinger & Schaye 2020).

⁴<https://www.chiantidatabase.org/>

⁵accessed on 2021-03-25

Table 4.3: Solar abundances for the elements we use in this work. We give metallicities for the elements we generate emission lines for, in number density relative to hydrogen (column 2) and mass fraction (column 3). The corresponding total solar metallicity value is $Z_{\odot} = 0.0127$. The values were also used to generate the line emission tables (e.g., Bertone et al. (2010a) table 1) and the EAGLE cooling tables (Wiersma et al. (2009a) table 1). The values are the CLOUDY v7.02 (last documented in Ferland et al. (1998) defaults, from Holweger (2001) H01), Allende Prieto et al. (2001) AP01, and Allende Prieto et al. (2002) AP02).

element	metallicity		source
	$n_{\text{elt}} / n_{\text{H}}$	$\rho_{\text{elt}} / \rho_{\text{tot}}$	
C	2.45×10^{-4}	2.07×10^{-3}	AP02
N	8.51×10^{-5}	8.36×10^{-4}	H01
O	4.90×10^{-4}	5.49×10^{-3}	AP01
Ne	1.00×10^{-4}	1.41×10^{-3}	H01
Mg	3.47×10^{-5}	5.91×10^{-4}	H01
Si	3.47×10^{-5}	6.83×10^{-4}	H01
Fe	2.82×10^{-5}	1.10×10^{-3}	H01

Note that the solar and element abundances assumed in the two sets of tables are different. The element abundances in the EAGLE simulations do not depend on this choice, and we calculate the emission by scaling the emission values from each table by the appropriate table value. We list the solar abundances used in the Bertone et al. (2010a) table generation in Table 4.3. These are the CLOUDY v7.02 (last documented in Ferland et al. (1998) default abundances.

We show some of the properties of these lines in Fig. 4.1 and in Tables 4.1 and 4.2. Many of these lines have been explored in more detail in Bertone et al. (2013). Note that much of this applies to EAGLE, even though the paper uses the OWLS simulations, because the radiative cooling model is the same.

The line selection consists of Lyman α -like (K- α) lines from H-like ions, He- α -like (resonance) lines from He-like ions, and for iron, a number of lines from the L-shell ions. The differences between the different K- α and He- α resonance lines in Fig. 4.1 (top right and left panels, respectively) can largely be explained by the different solar element abundances (peak heights) and element numbers (peak temperatures).

The uncertainty in the atomic data for the Fe L-shell lines is illustrated by the comparison between the Fe xvii 17.05 Å lines from the two table sets. Note that the curves are scaled to the same metallicities and element abundances, so differences in assumed abundances do not explain the difference. For the K α en He- α -like lines, the differences between the curves from the different table sets are $\lesssim 0.1$ dex around the emissivity peaks.

4.2.3 Surface brightnesses

We calculate surface brightnesses similarly to the column densities in Wijers et al. (2019, 2020). First, we calculate the luminosity of each gas particle as described in §4.2.2 then we project the particles onto a grid, using the Wendland (1995) C2 kernel as the assumed shape of each gas particle, scaled by its smoothing length. We choose a pixel size of 3.125 ckpc, matching that used in Wijers et al. (2019, 2020).

Before projecting, we divide the particles into ‘slices’ along the Z-axis of the simulation (an arbitrary direction relative to haloes and galaxies). Each slice is 6.25 cMpc thick, again




Figure 4.1: Emissivity for CIE conditions ($n_{\text{H}} = 10 \text{ cm}^{-3}$, $z = 0.1$) assuming solar abundances as a function of temperature for the different lines we study in this work. Vertical, coloured ticks on the top axis indicate the temperature where the emissivity of each line peaks, and the top axis shows the halo mass whose virial temperature T_{200c} matches the temperature on the bottom axis for a first-order prediction of which halo masses are best probed with which emission lines. We group the lines into different panels for legibility: the different He- α resonance lines (top left) and Lyman- α lines (top right) show curves of the same shape. Note the relatively shallow decline of the Lyman- α emissivity towards high temperatures. The different O VII He- α lines (bottom left) have similar curve shapes, though they are not identical. The iron L-shell lines (bottom right) have more strongly peaked emissivities than the other lines. The curves for the Fe XVII 16.78, 17.05, and 17.10 Å lines largely overlap in the plot. We show the Bertone et al. (2010a) table Fe XVII 17.05 Å emissivity curve as a dashed, yellow line.

matching previous work. We then divide the luminosity in each pixel by its angular size and squared luminosity distance to get a surface brightness.

For surface brightness profiles, we take these maps and extract surface brightnesses at different distances to central galaxies. We average all the values in annular bins. For median profiles, we use the median, in each annular bin, of the annular averages around individual central galaxies. We calculated the mean profiles by similarly averaging the annular means around the different galaxies at each impact parameter. We use impact parameter bins of 0.1 dex for the medians and 0.25 dex for the averages, because the average profiles are quite noisy using 0.1 dex bins. The innermost bin covers the 0–10 pkpc range.

In our surface brightness profiles, we ignore any possible emission from the star-forming gas. Assuming the star-forming gas is at 10^4 K (warm ISM), emission from this gas is negligible. A hot ISM phase might generate more emission, but modelling this phase would come with large uncertainties because the phases of the ISM are poorly resolved in EAGLE. Similarly, emission from (central and satellite) galaxies, e.g., X-ray binaries or AGN, is not included in our profiles.

Some emission from galaxies is included, however: gas that has recently been heated by supernova or AGN feedback. This gas will eventually expand and cool, but just after a feedback event, it will be hot (heated to $10^{7.5}$ or $10^{8.5}$ K, respectively) and dense, and will therefore be relatively bright in X-ray emission. However, these temperatures are high for the lines we study (Fig. 4.1). We will later see that the line-luminosity-weighted temperature of these haloes is generally well below these high temperatures. Luminosity-weighted mean temperatures can reach $\gtrsim 10^7$ K in the centres of haloes with $M_{200c} \lesssim 10^{11.5}$ K, but we find that these haloes are too faint to observe even with emission from this gas included. This effect is present for the K- α lines and the Mg XI and Si XIII He- α -like resonance lines, which have relatively high emissivities at these high temperatures.

In Appendix 4.A, we discuss the effect of this recently feedback heated gas on the surface brightness. In short, the effects are larger for the mean surface brightnesses than for the medians. Within $0.1 R_{200c}$, the effect on the median profiles can be substantial if the halo virial temperature is small compared to the temperature at which the emissivity peaks. For the mean profiles, the effects can be substantial for haloes up to this virial temperature. At these small impact parameters, we expect that, in practice, emission from the central galaxy itself would make it difficult to detect any CGM emission. At impact parameters between 0.1 and $1 R_{200c}$, the effects are typically small where the surface brightness is high enough that we would expect emission to be detectable by upcoming missions. They can be large when a halo seems to be marginally detectable, or not detectable at all.

Our surface brightness profiles therefore reflect the emission from the gas surrounding galaxies, but not from the galaxies themselves. This is the gas we are aiming to characterize. Emission from the galaxies themselves may, in practice, make it more difficult to detect the surrounding CGM in emission.

In this work, we use SPH-smoothed element abundances (SMOOTHEDELEMENTABUNDANCE in EAGLE) to calculate luminosities for consistency with the radiative cooling in EAGLE. We also use these abundances to calculate metal mass fractions and hydrogen number densities.

4.2.4 Galaxies and haloes

We take galaxies and haloes from the EAGLE public data release (McAlpine et al. 2016). These are identified in EAGLE by first finding haloes using a friends-of-friends method, where any particles (resolution elements) separated by less than 0.2 times the mean inter-particle distance are connected, and all connected particles define a halo. This algorithm is applied to dark matter particles, and other particles (gas, stars, and black holes) are then classified in the same way as their nearest dark matter particle. The centre of potential of the halo is the particle with the minimum gravitational potential energy.

Galaxies were found within those haloes using subfind (Springel et al. 2001; Dolag et al. 2009), and the central galaxy is the one containing the halo centre of potential. The subfind code finds overdense regions within these haloes, and subhaloes are identified as the self-bound parts of these overdense regions. This binding factors in gravitational potential energy and kinetic energy, as well as thermal energy for gas.

The halo mass M_{200c} was determined from the centre of potential: a sphere was grown around this centre until its average internal density was 200 times the critical density. The radius of that sphere is R_{200c} and the enclosed mass is M_{200c} . When we calculate 2d (impact parameter) or 3d radial distances to halo centres, we use the centre of mass of the central galaxy (subhalo 0) instead of the centre of potential, to approximate the light-weighted centre

of the galaxy that might be used in observations.

The temperature T_{200c} of the hot CGM phase at the halo radius R_{200c} can be estimated from the virial theorem, assuming hydrostatic equilibrium, with pressure support coming from the thermal pressure of the hot phase:

$$T_{200c} = \frac{\mu m_H}{3k} G M_{200c}^{2/3} (200\rho_c)^{1/3}, \quad (4.1)$$

where μ is the mean particle mass in units of the hydrogen atom mass m_H , k is the Boltzmann constant, G is Newton's constant, and ρ_c is the critical density. We assume $\mu = 0.59$, for primordial gas with fully ionized hydrogen and helium. The assumption of hydrostatic equilibrium is not valid in especially the inner CGM of L_* galaxies in the EAGLE simulation (Oppenheimer 2018), but the volume-filling phase in the $\sim L_*$ CGM of EAGLE galaxies is still at X-ray producing temperatures, $\gtrsim 10^6$ K (e.g., Wijers et al. 2020).

4.2.5 CGM definitions

The CGM does not have one clear definition. It is roughly the gas surrounding a central galaxy, but it does not have a clear inner or outer boundary. Therefore, in our 3D profiles of gas properties, we only exclude star-forming gas (ISM), but otherwise show gas over a large radial range. When calculating total luminosities in halo mass ranges (Fig. 4.4), we use the 'FoF and gas within R_{200c} ' definition of haloes and CGM. For subhalo gas, we use the subfind definition, where subhalo gas is gravitationally bound to a subhalo other than subhalo 0, as indexed by subfind. The subhalo with index 0 is the central galaxy, and all gas bound to the halo, but no specific subhalo, is attributed to subhalo 0 by subfind. In other plots, where we use luminosities of individual haloes (Figs. 4.5 and 4.8), we define the CGM as all the non-star-forming gas within R_{200c} of (the center of mass of) the central galaxy. In Fig. 4.7 we do include star-forming gas in our radial profiles of gas properties. We note that the contribution of this star-forming gas to the halo luminosity of these lines is negligible.

4.3 Detectability

The detectability of emission is not easy to define. For example, Das et al. (2020) use two different tests to investigate how far from the galaxy they can detect emission. Detection might be limited by backgrounds and foregrounds (both astrophysical and instrumental), and non-CGM emission correlated with the source. Such correlated emission would be, for example, X-ray emission from X-ray binaries, hot ISM, or AGN in central and satellite galaxies. Since backgrounds and foregrounds (may) vary with time and position on the sky, these need to be fit to observations along with the object of interest, meaning that systematic errors in the background (models) are also relevant.

Here, we investigate detectability with a number of instruments under somewhat optimistic assumptions: we consider only statistical errors (Poisson noise in total counts), where backgrounds contribute to the noise, but not the signal, and we ignore the effects of other X-ray sources correlated with the CGM. Modelling galaxy X-ray emission would be difficult, since EAGLE does not resolve the multi-phase ISM or the time and spatial scales governing AGN variability, or assume binary fractions in its star formation and feedback model. However, it has been shown that e.g., aggregate AGN feedback and the resulting quenching of star formation anti-correlate with X-ray emission from the CGM at $z \approx 0$ in $\sim L_*$ galaxies

(Davies et al. 2019), meaning that just ‘painting in’ a galaxy of the right stellar mass might be insufficient. Modelling this galactic X-ray emission is beyond the scope of this paper.

We focus on the detection of individual emission lines and ignore the issue of detecting them over the continuum emission of the CGM. For high-mass systems, such as clusters, this might be an important limitation, but then possible observations of clusters in X rays have already been studied by e.g., Lotti et al. (2014), and can make use of present observations in modelling. The spectral resolution (full width at half maximum) of the X-IFU should be 2.5 eV up to 7 keV (Barret et al. 2018). For rest-frame energies between 0.3 and 2 keV, this corresponds to a velocity range between 2500 km s^{-1} and 375 km s^{-1} . These ranges are large enough that we assume the redshift of the emission line is known with sufficient precision relative to the central galaxy redshift, and that redshift trials are not an issue. (Note that this is not necessarily the case for high spectral resolution measurements with Lynx.) We ignore blending of different emission lines; in our line sample, the Fe xvii 17.05 and 17.10 Å lines would be blended. Finally, to get a single surface brightness limit, we assume uniform emission, at least within each region the surface brightness is measured in.

This brings us to the following equations, adapted from Takei et al. (2011) to include instrumental backgrounds:

$$r = \int_{\text{line}} dE \sum_j \text{SB}(E) A_{\text{eff}}(E) \text{LSF}(E, j) \quad (4.2)$$

where r is the count rate per unit angular size, E is the energy, j is the spectral channel, SB is the surface brightness, A_{eff} is the effective area, and LSF is the line spread function. The sum over the channels would be centred on the channel corresponding to the emission line energy; the range of extraction can be varied. This describes the conversion from photons to counts as encoded in the instrument response files. Then

$$N_{\sigma} = \frac{r_{\text{line}}}{\sqrt{r_{\text{line}} + r_{\text{bkg}}}} \sqrt{t_{\text{exp}} \Delta\Omega}, \quad (4.3)$$

where N_{σ} is the detection significance in units of σ , r_{line} and r_{bkg} are the count rates per unit observed solid angle for the line and total background, respectively, $\Delta\Omega$ is the angular size of the observed region, and t_{exp} is the exposure time. This assumes that errors in background modelling are negligible. Finally, we assume $N_{\sigma} \geq 5$ constitutes a detection.

Given the response functions and backgrounds, we can therefore estimate what line surface brightness would be detectable. Note that this surface brightness is not the intrinsic surface brightness of the source. Especially for the lower-energy lines we consider, Galactic absorption will reduce the amount of light that makes it to the instrument. Since this is a very simple correction, we give absorbed minimum detectable surface brightnesses and unabsorbed minima assuming the same Galactic absorption as in the X-IFU background model (McCammon et al. 2002): an xspec wabs model (Morrison & McCammon 1983) with a hydrogen column density $N_{\text{H}} = 1.8 \times 10^{20} \text{ cm}^{-2}$ (model parameter value 0.018).

Nominally, we assume an extraction area for surface brightnesses of $\Delta\Omega = 1 \text{ arcmin}^2$. This corresponds to a circle of radius 64 pkpc at redshift 0.1. We try exposure times of 0.1, 1, and 10 Ms. Because the extraction area and exposure time are degenerate for the purposes of minimum surface brightness estimates, we report their product instead of individual values. We generally find that 100 ks and 1 arcmin^2 would not be sufficient to detect line emission, so we focus on the larger values.

4.3.1 Overview of instruments

For the Athena X-IFU, we used the response matrices and backgrounds provided on the X-IFU website⁶. Barret et al. (2016, 2018), and Ravera et al. (2014) describe the production of the response matrices. The responses assume the cost-constrained configuration, in the baseline filter configuration (open filter wheel position). Lotti et al. (2012); Lotti et al. (2014) document the instrumental background estimates, and the McCammon et al. (2002) astrophysical background (applicable to sources at high galactic latitudes) is used. For the AGN contribution to the background, it is assumed that 80 per cent of the AGN can be (spatially) resolved and removed.

For the Science with the X-ray Imaging and Spectroscopy Mission (XRISM; XRISM Science Team 2020), we consider the Resolve instrument. We use the response files and instrumental backgrounds from the XRISM online database⁷. We use the models for a 5 eV spectral resolution (FWHM), with the gate-valve open. We use the .arf file for a constant surface brightness disk with a radius of 5 arcmin.

For the astrophysical background, we use a model from Simionescu et al. (2013), fit to Suzaku and ROSAT data taken around the Coma cluster (but in regions free from cluster emission), which was the target of the study. This includes the AGN background. Given the size of the PSF (1.2') compared to the field of view (2.9') of the Resolve instrument, we do not expect excising this point source background will be feasible. There might be similar issues separating galactic emission from CGM emission here, especially a possible hot ISM contribution, since this could have a very similar spectrum (collisionally ionized plasma) as the warm/hot CGM. Given a sufficiently deep galaxy survey, it might be possible to avoid this confusion by targeting a galaxy-free region of the CGM. At our nominal redshift (0.1), it will likely not be possible to avoid satellite galaxies given the extent of the point spread function (Table 4.4).

For Lynx, we use data provided by Alexey Vikhlinin (private communication, 2020). We investigate detectability with the .rmf response files provided by Alexey Vikhlinin and .arf files downloaded from the SOXS instrument simulator website⁸. We used astrophysical and instrumental background data provided by Alexey Vikhlinin, matching the specifications in The Lynx Team (2018). This means the Hickox & Markevitch (2007) astrophysical background model (derived from Chandra measurements) is used, and instrumental backgrounds for the Lynx X-ray Microcalorimeter (LXM) are based on Athena X-IFU predictions. Point sources are assumed to be fully resolved, and therefore subtractable from the data, in deep exposures.

For the LXM we consider two modes (The Lynx Team 2018): the Ultra-High Resolution Array (0.3 eV resolution, 1 arcmin FOV) and the Main Array (3 eV resolution over a 5 arcmin FOV). The UHRA is not sensitive at higher energies (> 0.95 keV), so it can only be used for the lower-energy lines we investigate.

We provide an overview of these instruments, which are potentially interesting for soft X-ray emission line detections, in Table 4.4. We list the point spread function (PSF) and field of view (FOV) sizes, and the spectral resolution. We show the effective area as a function of energy in the left-hand panel of Fig. 4.2. From eq. 4.2 the effective area is the main factor determining the emission line counts, while the spectral resolution determines how much of the background comes with it. The point spread function helps determine background

⁶<http://x-ifu.irap.omp.eu/resources-for-users-and-x-ifu-consortium-members/#accordion-item-latest-x-ifu-response-matrices> downloaded on September 28, 2020.

⁷<https://heasarc.gsfc.nasa.gov/FTP/xrism/prelaunch/simulation/sim2/> downloaded on October 20, 2020.

⁸<https://hea-www.cfa.harvard.edu/soxs/responses.html> downloaded November 21, 2018.

Table 4.4: Basic specifications for a number of instruments, and some relevant values for $z = 0.1$. We report the point spread function (PSF) and field of view (FOV) for the instruments, and translate those values to physical sizes at $z = 0.1$. The PSF types are the half-power diameter (HPD) and half-energy width (HEW). For the spectral resolution, we report the full width at half maximum (FWHM). The Lynx instruments are the Main Array (MA) and Ultra-High Resolution Array (UHRA) of the Lynx X-ray Microcalorimeter (LXM). For XRISM, we report the 1.2' angular resolution achieved by Hitomi, rather than the 1.7' requirement, and similarly, Hitomi's 5.0 eV spectral resolution rather than the 7 eV requirement. The Athena X-IFU detector is hexagonal, not square; the given field of view is an equivalent diameter. The data are taken from [Lumb et al. \(2017\)](#) L17), [XRISM Science Team \(2020\)](#) X20), and [The Lynx Team \(2018\)](#) L18).

instrument	PSF			FOV		ΔE (FWHM) eV	source
	arcsec	pkpc	type	arcmin	pkpc		
Athena X-IFU	5	10	HEW	5	570	2.5	L17
XRISM Resolve	72	140	HPD	2.9	330	5.0	X20
LXM (MA)	0.5	1.0	HPD	5	570	3	L18
LXM (UHRA)	0.5	1.0	HPD	1	110	0.3	L18

levels through the ability to resolve and remove individual background AGN, and the FOV determines how many pointings it would take to image a source. (We do not account for this in our exposure times; these are always for single pointings.)

The PSF also determines the level of galaxy ‘contamination’ in the CGM images. Because we do not model this galactic emission, we cannot make precise estimates, but the relative sizes of the PSFs of the instruments should at least give an idea of the relative effects.

Fig. 4.2 shows the effective area of these instruments (left panel), taken from the .arf files we described above. For XRISM, this includes a correction for the assumed sizes of the source and the instrument field of view. We also show the effect of Galactic absorption (transmitted fraction, scaled to 10^4 cm^2). The decreasing instrument sensitivities and strong Galactic absorption mean that lines at observed energies $\lesssim 0.3 \text{ keV}$ will be difficult to detect. The right panel of Fig. 4.2 shows the resulting minimum detectable source surface brightnesses as a function of line energy. Different colours indicate different instruments, line styles are for different exposure times, and faded lines indicate the minima after absorption by the Galaxy.

The Lynx ultra-high resolution array is the most sensitive instrument, but has a limited energy range (up to 0.95 keV). After that, the Lynx main array is most sensitive; it has a larger energy range, fully covering the 0.3–2 keV range we are interested in, and a larger field of view (a diameter of 5' instead of 1'). The Athena X-IFU will have a similar sensitivity at low energies, but the difference with the Lynx main array increases towards higher energies. The XRISM Resolve instrument has the lowest sensitivity, and the strongest sensitivity decrease towards lower energies.

The features in the sensitivity curves in the right panel of Fig. 4.2 are due to detector edges (jumps in the effective area curves; left panel) and features (emission lines) in the astrophysical background. Around these lines, systematic errors are likely to contribute significantly to the actual error budget, so our estimated surface brightness limits are likely underestimated between ≈ 0.3 and 1 keV. We have checked that the different astrophysical background models are similar, so differences between those models should not be driving the sensitivity differences between the instruments.

The features are stronger for larger exposure times and effective areas. This is because, for small exposure times ($t_{\text{exp}} \Delta\Omega \ll N_{\sigma}^2 / r_{\text{bkg}}$), the limiting factor for detecting an emission line is the number of detected source photons. When the exposure time increases, the as-

Table 4.5: Minimum source surface brightness (\log_{10} photons $\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$) detectable at a 5σ significance for the different lines and instruments at $z = 0.1$, given different exposure times and angular sizes. These include the effect of Galactic absorption according to a wabs model. The limits are averaged over 11 redshifts from $z = 0.095$ to 0.105 . Dashes indicate lines outside the sensitivity range of an instrument. The final column indicates the (log) difference between the minimum detectable source surface brightnesses including and excluding the effect of Milky Way absorption (wabs model with $N_{\text{H}} = 1.8 \times 10^{20} \text{ cm}^{-2}$).

instrument	XRISM Resolve			Athena X-IFU			LXM UJHR			LXM main			Δ_{wabs} [\log_{10} SB]	
	1e7	1e6	1e5	1e7	1e6	1e5	1e7	1e6	1e5	1e7	1e6	1e5		
C V	0.8	1.7	2.7	2.7	-0.7	-0.1	0.6	-1.2	-0.6	0.2	-0.8	-0.2	0.4	0.26
C VI	1.3	2.2	3.1	3.1	-1.0	-0.3	0.5	-1.5	-0.7	0.3	-1.0	-0.4	0.4	0.18
N VI	0.6	1.5	2.5	2.5	-1.2	-0.6	0.2	-1.8	-1.0	-0.1	-1.3	-0.7	0.0	0.11
N VII	0.2	1.0	2.0	2.0	-1.4	-0.8	0.1	-2.0	-1.2	-0.3	-1.6	-0.9	-0.2	0.08
O VII (f)	-0.1	0.7	1.7	1.7	-1.6	-0.9	-0.1	-2.1	-1.3	-0.4	-1.7	-1.1	-0.3	0.05
O VII (i)	-0.1	0.7	1.7	1.7	-1.6	-0.9	-0.1	-2.1	-1.4	-0.4	-1.7	-1.1	-0.3	0.05
O VII (r)	-0.2	0.7	1.7	1.7	-1.6	-0.9	-0.1	-2.1	-1.3	-0.4	-1.7	-1.1	-0.3	0.05
O VIII	-0.2	0.7	1.7	1.7	-1.5	-0.9	0.0	-2.1	-1.4	-0.4	-1.7	-1.1	-0.3	0.06
Fe XVII (17.05 A)	-0.4	0.5	1.5	1.5	-1.7	-1.0	-0.1	-2.3	-1.5	-0.5	-1.7	-1.2	-0.4	0.04
Fe XVII (15.26 A)	-0.5	0.3	1.3	1.3	-1.6	-0.9	-0.2	-2.2	-1.5	-0.6	-1.8	-1.2	-0.5	0.04
Fe XVII (16.78 A)	-0.4	0.5	1.4	1.4	-1.8	-1.1	-0.2	-2.3	-1.5	-0.5	-1.9	-1.3	-0.5	0.04
Fe XVII (17.10 A)	-0.3	0.5	1.5	1.5	-1.6	-1.0	-0.1	-2.3	-1.5	-0.5	-1.6	-1.0	-0.4	0.05
Fe XVIII	-0.5	0.4	1.4	1.4	-1.7	-1.0	-0.2	-2.3	-1.5	-0.5	-1.8	-1.2	-0.5	0.04
Ne IX	-0.7	0.2	1.2	1.2	-1.8	-1.1	-0.3	-2.4	-1.6	-0.6	-2.0	-1.4	-0.6	0.03
Ne X	-0.7	0.1	1.1	1.1	-1.9	-1.2	-0.3	-	-	-	-2.1	-1.4	-0.6	0.02
Mg XI	-0.9	0.0	0.9	0.9	-2.1	-1.3	-0.4	-	-	-	-2.3	-1.6	-0.7	0.01
Mg XII	-0.9	-0.1	0.9	0.9	-2.1	-1.3	-0.4	-	-	-	-2.3	-1.6	-0.7	0.01
Si XIII	-0.8	0.0	1.0	1.0	-2.1	-1.3	-0.3	-	-	-	-2.2	-1.5	-0.6	0.00




Figure 4.2: Effective area (left) and minimum detectable surface brightness (with 5σ confidence; right) as a function of observed line energy for the different instruments (different colours). We show the Athena X-IFU (X-IFU), the Main Array (main) and Ultra-High Resolution Array (UHR) of the Lynx X-ray Microcalorimeter (LXM), and the XRISM Resolve instrument (XRISM-R). Along with the effective area curves, we also show the fraction of transmitted radiation through the Milky Way according to the wabs model with a hydrogen column density $N_{\text{H}} = 1.8 \times 10^{20} \text{ cm}^{-2}$. We show minimum surface brightnesses for two values of exposure time and angular extent of the source: $t_{\text{exp}} \Delta\Omega = 10^6$ and $10^7 \text{ arcmin}^2 \text{ s}$, with solid and dotted lines, respectively. Saturated and faded colours show minimum surface brightnesses before and after absorption by the Milky Way halo, respectively, using the wabs model shown in the left-hand panel. For each instrument, we indicate the (full) energy range, centred on the line energy, over which the signal and noise were measured. We assume an intrinsic line width of 100 km s^{-1} .

trophysical and instrumental backgrounds become relatively more important, as shown by solving equation 4.3 for the minimum detectable emission line count rate r_{line} :

$$r_{\text{line}} = \frac{N_{\sigma}^2}{2t_{\text{exp}} \Delta\Omega} \left(1 + \sqrt{1 + \frac{4t_{\text{exp}} \Delta\Omega r_{\text{bkg}}}{N_{\sigma}^2}} \right). \quad (4.4)$$

This means that emission lines in the astrophysical background have a stronger effect on the minimum detectable surface brightness when the exposure time is larger.

The effect of the effective area is most clearly seen for the XRISM resolve instrument. Here, the background emission line features are mostly absent because the instrument has a much smaller effective area than the others (Fig. 4.2 left panel). Though the instrumental background count rate of XRISM Resolve is lower than for the other instruments, its smaller effective area means that the total background is nonetheless mostly dominated by the instrumental background rather than the astrophysical one, meaning the lines in the astrophysical background have a smaller effect.

Finally, we see that Galactic absorption makes a considerable difference in what can be detected at the lowest energies, but the effect is small ($\lesssim 0.1$ dex) at the higher energies ($\gtrsim 0.4 \text{ keV}$) we consider.

In Table 4.5 we show the minimum detectable surface brightnesses (SB) for the different lines and instruments that we investigate. For example, according to Table 4.5 for Athena X-IFU the O VIII detection limit is $\text{SB} = 10^{-0.9} \text{ photons cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ for $t_{\text{exp}} \Delta\Omega = 10^6 \text{ s arcmin}^2$. This means that for this surface brightness we require 10^6 s to detect a region

with angular size 1 arcmin^2 , or 10^5 s to detect a region of size 10 arcmin^2 . To convert these minimum detectable surface brightnesses to units of $\text{photons s}^{-1} \text{ cm}^{-1} \text{ arcmin}^{-2}$, multiply the values by 1.18×10^{-7} (or subtract 7.07 from the log values).

These minima include the effect of Milky Way absorption on the observed surface brightness. In the final column we show how much of a difference this absorption makes; there are small variations between the different redshifts used for one line, but these are $\lesssim 0.01 \text{ dex}$. We provide these differences to make it easier to calculate the minima with different absorption models or absorbing columns, at least to first order.

4.4 Results

We start by demonstrating how line emission overall relates to haloes in EAGLE (§4.4.1). We then examine this emission as a function of impact parameter with surface brightness profiles (§4.4.2), and compare these surface brightnesses to rough estimates of what could be detected with various instruments. Finally, we examine emission-weighted gas temperatures, densities, and metallicities to study which gas produces the emission, and how it relates to the overall gas content of haloes (§4.4.3).

4.4.1 Line emission in relation to haloes

Fig. 4.3 shows the emission from the different lines we study in a part of the 100^3 cMpc^3 volume selected to have haloes with a range of masses. The regions we choose are among the most overdense regions in the volume. We indicate the positions of the haloes for comparison; the circles are at R_{200c} in all panels but the top, middle panel, where we indicate haloes at $2R_{200c}$. We can see that the line emission is brightest in haloes. The emission from the cosmic web is weak, and will not be directly detectable.

In the top left of the smaller panels in Fig. 4.3, we see emission from various lines that does not seem to correspond to a halo, close to the most massive halo in the panel. It also does not correspond to a halo centred just outside the region along the Z-axis. However, the emission in a number of lines reveals it is connected to that massive halo ($M_{200c} = 10^{14.47} M_{\odot}$). In fact, the gas there is part of the same FoF group as the massive halo in the top left. Evidently, this is a merging system, and the top left halo is no longer separately identified. This is therefore halo gas/CGM emission, and not e.g., part of a bright filament.

Fig. 4.4 divides the total luminosity from different lines in the EAGLE volume into contributions from haloes of different masses and the IGM, with the mass fractions in these components shown at the bottom for comparison. Star-forming gas is included at 10^4 K , but its contribution to the total is negligible. Line emission is dominated by haloes, in contrast to the mass. The halo mass contributions to the emission differ between lines. The halo mass ranges contributing most to the total luminosity are mostly those for which T_{200c} corresponds to the CIE peak emission temperature shown in Fig. 4.1. In Wijers et al. (2020 fig. 2), we saw that the metals and ions producing a number of these lines are less concentrated in haloes than is the case for the emission. We expect the difference in results from the $\propto n_{\text{H}}^2$ dependence of line emission, compared to the $\propto n_{\text{H}}$ dependence of ion mass.

Note that the IGM contributions here can differ considerably, in relative terms, between the two emissivity tables we used in the calculations, since this low-density gas is photoionized, and the tables assume different UV/X-ray backgrounds. The contributions to the total are small in either case.




Figure 4.3: Surface brightness maps for a large set of X-ray emission lines. The O VIII line maps (top row) show different parts of the $100 \times 100 \times 6.25 \text{ cMpc}^3$ slice of the EAGLE RefL100N1504 volume at $z = 0.1$, centred on $Z = 21.875 \text{ cMpc}$. The projection is along the Z-axis. The top left panel shows a $25 \times 25 \text{ cMpc}^2$ region, centred on $X, Y = 57.5, 4.5 \text{ cMpc}$ at a resolution of 62.5 ckpc per pixel. The middle panel shows the full slice at a resolution of 250 ckpc per pixel. Relative to the simulation coordinates, the Y coordinates were shifted by 15 cMpc so the region on the left would not overlap the periodic boundary at the bottom of the image. The top right panel, and all the other panels, show a $12 \times 12 \text{ cMpc}^2$ region at 31.25 ckpc per pixel, centred on $X, Y = 64.5, 29.5 \text{ cMpc}$. Circles indicate the positions of haloes in the volume; these are centred on the center of mass of the central galaxy and have a radius of R_{200c} , except for the panel showing the whole slice; there the circles indicate $2R_{200c}$. The color additionally shows the halo mass. The surface brightness color bar transitions to colour in the range where emission roughly becomes directly observable ($10^{-1.5} \text{ photons s}^{-1} \text{ cm}^{-2} \text{ sr}^{-1}$). The brightest emission comes from within haloes, and the IGM emission is typically very weak.




Figure 4.4: Fraction of the total luminosity of the different lines originating in haloes of different mass (different colors) and the IGM (gray) in EAGLE at $z = 0.1$. Halo gas is anything in a FoF group or within R_{200c} , and the IGM is anything else. The numbers at the right of each bar indicate the volume-averaged luminosity density in each line ($\log_{10} \text{erg s}^{-1} \text{cMpc}^{-3}$). For comparison, we also show the fraction of the gas mass in these components.

Gas that has recently been directly heated by feedback can be responsible for a large fraction of the emission from haloes at masses where T_{200c} is too low for the CGM to produce detectable emission. This effect is substantial in roughly the same halo mass ranges where the effect of this gas on the surface brightness profiles of the haloes is substantial, described in Appendix 4.A. This means that at low halo masses, where the contribution of a given halo mass range to the total emission of a given line is already small, it would, in general, be even smaller if the recently, directly heated gas were excluded.

Comparing the halo emission contributions to the Fe L-shell, He- α -like, and K- α line emission, we see a secondary trend. The lines with the narrowest emissivity peaks (the iron lines, Fig. 4.1) have most of their halo emission coming from only two halo mass bins. The He- α -like lines, with wider emissivity peaks, have haloes over a wider mass range contributing to their total emission. The K- α lines tend to come from a wider range of halo masses, especially towards higher halo masses, reflecting their wider emissivity peaks, with relatively shallow

slopes towards high temperatures.

We will find that many trends of line emission with halo mass are driven primarily by these two characteristics of the emissivity curves of the lines: the temperature of the emissivity peak, compared to the halo virial temperature, and the width of the peak.

Fig. 4.5 shows the median line flux as a function of halo mass. Generally, we see that the oxygen lines are strongest. Though the luminosities generally increase with halo mass, we do see differences in these trends.

Fig. 4.5 also shows the scatter (central 80 per cent) in those luminosities at fixed halo mass. This scatter is generally quite large: at least ≈ 1 dex at $M_{200c} \lesssim 10^{13} M_{\odot}$, and ≈ 2 dex for Ne x, Mg xi, Mg xii, and Si xiii for L_* haloes. This large scatter implies that average and median surface brightnesses can differ considerably, and that the manner in which luminosity-weighted temperature, density, and metallicity distributions from different haloes are stacked can have a real impact on the resulting radial profiles.

For the He- α resonance lines in the top panel of Fig. 4.5 we see a trend where halo luminosities tend to flatten as a function of halo mass above the emission peak temperature (Fig. 4.1). The Fe L-shell lines even decrease in luminosity at the highest halo masses, when the haloes become hotter than their emissivity peaks. For the Lyman- α lines (second panel from the top), there appears to be a weak flattening above the emissivity peak temperature haloes. The flattening is probably less obvious than for the He- α lines because the emissivities decrease less strongly with temperature for the Lyman- α lines than for the He- α lines. In the third panel from the top, we see that the different O vii He- α lines follow very similar trends, including their scatter.

4.4.2 Surface brightness profiles

In Fig. 4.6 we show radial surface brightness profiles. Different panels show different emission lines, and different colours correspond to different halo mass bins. The solid (dotted) lines show median (mean) profiles. For these, the starting point is a set of surface brightness maps for each emission line. We first average the surface brightnesses in annular bins around the central galaxy of each halo in a halo mass bin. Then, we take the median of these profiles in each annular bin. For the mean profiles, we average the profiles in each annular bin instead. The median profiles show the surface brightness profile we can expect for a typical halo in each mass range, while the means show what we could expect from stacking the emission in each halo mass bin.

We use larger radial bin sizes for the means than for the medians. We choose these larger bins for legibility and to highlight trends with halo mass and radius. Light-coloured, vertical bands indicate the virial radii R_{200c} for each mass bin. Horizontal shaded regions indicate the estimated sensitivity limits for different instruments, for exposure times and sky areas of 10^6 – 10^7 arcmin² s.

Generally, the median surface brightnesses decline as a function of distance to the halo centre, and with the exception of the most massive haloes, typically drop off by $\gtrsim 2$ dex from their peak before R_{200c} . The different behaviour for the most massive haloes ($M_{200c} \geq 10^{14} M_{\odot}$) is most likely because their virial temperatures are $\gtrsim 10^7$ K, above the emission peaks for all these lines, and the centres of the haloes are hotter than the outskirts (Fig. 4.7).

In Appendix 4.A we investigate the effect of gas recently heated directly by feedback. In short, the effects of this gas tend to be limited to halo centres (impact parameters $\lesssim 0.1 R_{200c}$) and regions where the surface brightnesses including the feedback-heated gas are too low to be detected ($\lesssim 10^{-2}$ photons cm⁻²s⁻¹sr⁻¹). The effects are larger for the mean profiles than




Figure 4.5: Median flux from $z = 0.1$ EAGLE haloes (all gas within R_{200c}) as a function of halo mass for the different lines. We show the flux in units suited for estimating photon statistics for observations: photons per 100 ks per m² of effective area. We calculate the photon fluxes using a luminosity distance of 1.48×10^{27} cm. Error bars, shown in alternating halo mass bins for legibility, show the central 80 per cent of luminosities at each halo mass. (They represent scatter in the population, not noise.) Panels show, from top to bottom, the He- α -like resonance lines, K- α lines, O VII He- α -like lines, and the iron L-shell lines. The O VII He- α resonance line is shown in both the first (top) and third panels. The curves for the O VII resonance and forbidden lines largely overlap, as do the Fe XVII 16.78, 17.05, and 17.10 Å curves. Luminosities almost always increase with halo mass, and luminosity scatter is typically largest at low halo masses and low luminosities.




Figure 4.6: Median and mean surface brightness as a function of impact parameter r_{\perp} and halo mass. Means (dotted) are calculated using all haloes in each mass bin, or a subsample of 1000 randomly chosen haloes if the bin contains more than that. For the means, we use 0.25 dex bins, starting after the first, 0–10 pkpc, bin. For the medians, we use 0.1 dex bins. The medians are per-mass-bin medians of the annular average profiles of individual haloes. We also show minimum observable surface brightnesses for long integrations with different instruments. The radial range of the indicated limits is set by the field of view and point spread function (half the diameters in Table 4.4) of each instrument. The emissivity curves (Fig. 4.1) predict the halo mass for which the median SB peaks reasonably well for the He- α triplets and Fe L-shell lines, but for Lyman- α lines, the emission just peaks at the highest halo mass for all lines. Comparing these lines to detection limits gives an impression of whether we could detect typical emission in these haloes. The right axis shows surface brightnesses in units useful for estimating photon numbers; for reference, 10 arcmin^2 is the area of a circle with radius 204 pkpc at $z = 0.1$.

for the medians. Where differences between mean and median profiles are relatively large, the detectability of emission in stacks (mean profiles) may hinge on the inclusion of the gas recently heated by feedback.

Even without stacking, Athena and Lynx should be able to detect many emission lines from groups and clusters ($M_{200c} \gtrsim 10^{13} M_{\odot}$), and some lines from L_* and local group mass haloes ($M_{200c} \approx 10^{12}-10^{13} M_{\odot}$) out to distances far into the CGM. With XRISM, we should be able to detect a few emission lines from groups and clusters. Given the declining surface brightness with impact parameter, mapping the line emission of a typical halo less massive than clusters ($M_{200c} \lesssim 10^{14} M_{\odot}$) out to R_{200c} will, however, be very difficult. However, it may be feasible to detect the emission statistically for a large sample of objects.

The mean profiles (dotted lines) broadly follow the trends of the medians (solid lines), with a few key differences. First, the mean surface brightnesses are larger than the medians. At halo masses close to that where the surface brightness peaks, and close to halo centres, the mean surface brightness tends to lie around the 90th percentile of the surface brightness distribution at each radius. However, at larger radii, differences between mean and median values become larger, and the mean is more clearly dominated by extremes. Generally, the differences are larger at lower median surface brightnesses, and the mean profiles often flatten out at large radii, unlike the medians.

Additionally, there is often a stronger central peak in the mean surface brightnesses; this is typically at $\lesssim 0.1 R_{200c}$, in a region that may still be associated with the central galaxy. Often, this bright emission is associated with recent feedback events (Appendix 4.A).

At large radii, where median surface brightnesses are low, the mean profiles become noisy (especially clear when using smaller radial bins), likely dominated by bright emission in one or a few galaxies, even in mass bins with many haloes.

The trends of the median profiles with halo mass differ between the different sets of lines (K- α , He- α -like, and Fe L-shell) we investigate. The K- α lines have brightnesses that consistently increase with halo mass, except in the centres of the highest mass haloes. On the other hand, the He- α -like lines have surface brightnesses that more clearly peak with halo mass. The iron L-shell lines have even more extreme surface brightness peaks with halo mass. This follows the trends we saw for halo luminosities (Fig. 4.5), except that the luminosity is spread over a larger area in more massive haloes, so roughly constant luminosities with halo mass lead to surface brightness peaks with halo mass.

In general, we find that most of these emission lines should be detectable in haloes of some mass with the Athena X-IFU. The brightest emission line is O VIII K- α , and we expect it to be detectable in haloes of $M_{200c} \gtrsim 10^{12} M_{\odot}$ with the X-IFU. In groups and clusters ($M_{200c} \gtrsim 10^{13} M_{\odot}$), this line may even be detectable out to R_{200c} . The other K- α lines are typically detectable in haloes $\gtrsim 10^{12.5}$ or $10^{13} M_{\odot}$ with this instrument.

Emission from the brightest He- α -like species, O VII, may also be detectable in $M_{200c} \gtrsim 10^{12} M_{\odot}$ haloes with the X-IFU. Emission lines from C v and N VI will be difficult to detect at all due to the high sensitivity limits at these low energies. However, higher-energy He- α -like lines will likely be detectable with the X-IFU in groups ($M_{200c} \approx 10^{13}-10^{14} M_{\odot}$), and emission from some of these lines may be detectable in $M_{200c} \gtrsim 10^{14} M_{\odot}$ or $M_{200c} \approx 10^{12.5}-10^{13} M_{\odot}$ haloes.

The O VII forbidden line is generally about as bright as the resonance line, and its detection prospects are similar. The intercombination line is somewhat weaker, and would therefore be more difficult to detect. With the X-IFU and a long integration time, detecting all three lines may be possible in the inner CGM of $M_{200c} \approx 10^{12.5}-10^{14} M_{\odot}$ haloes, making the diagnostic information of the combination available.

The Fe L-shell lines should be clearly detectable in groups ($M_{200c} \approx 10^{13} - 10^{14} M_{\odot}$) using the X-IFU; a few of these lines may be detectable out to almost R_{200c} . Some Fe L-shell emission lines from the centres of local group mass systems ($M_{200c} \approx 10^{12.5} - 10^{13} M_{\odot}$) should also be detectable with the X-IFU.

Overall, using the Athena X-IFU, we expect to be able to detect many different emission lines from the CGM of galaxy groups and clusters (IGrM/ICM; $M_{200c} \gtrsim 10^{13} M_{\odot}$), and a few of these lines may be detectable out to R_{200c} . Emission from the inner CGM of $M_{200c} \approx 10^{12.5} - 10^{13} M_{\odot}$ haloes should be clearly detectable for some of the emission lines, and marginally detectable in most of the others we study. For $M_{200c} \approx 10^{12} - 10^{12.5} M_{\odot}$ haloes, detections will be difficult for most emission lines we study, but detection of O VII and O VIII emission is likely possible with large exposure times ($t_{\text{exp}} \Delta\Omega \approx 10^7 \text{ s arcmin}^2$).

XRISM Resolve is clearly less sensitive than the X-IFU, but emission from groups and clusters ($M_{200c} \gtrsim 10^{13} M_{\odot}$) may be detectable in a few bright lines. Fe XVII emission may be marginally detectable with XRISM, especially if the different lines are taken together. (At the spectral resolution of XRISM Resolve, the 17.05 and 17.10 Å Fe lines will be blended.)

The Lynx MA generally has sensitivity limits similar to those of the X-IFU, but it is a bit more sensitive to line emission, especially at higher energies. The UHRA is clearly more sensitive. In addition to what is possible with the X-IFU, this instrument will enable clear detections of emission from the centres of $M_{200c} \approx 10^{12.5} - 10^{13} M_{\odot}$ haloes in more lines and out to larger impact parameters, and it will increase the number of emission lines we can detect from $M_{200c} \approx 10^{12} - 10^{12.5} M_{\odot}$ haloes.

Note that due to the relatively small field of view of the Lynx MA, for many lines and halo masses, multiple pointings would be needed to cover the area where we expect emission to be detectable. The high spectral resolution of this instrument also likely means we underestimate some of the uncertainties involved in line detections with the UHRA. For example, we might have to account for at least a few different possible line centres when defining the detection significance. This can raise the 5σ surface brightness limit above the values we report here for a single redshift trial.

Despite reasonable detection prospects for a number of these haloes with the different instruments, detailed imaging will be very difficult except for the brightest lines and most massive ($\gg L_*$) haloes at $z = 0.1$. This is because these detections would require combining large areas (at least a square arcminute), together with long exposure times (at least 1 Ms). Examining haloes at lower redshifts (< 0.1) might be helpful here, though too low redshifts would make the emission difficult to distinguish from local and Milky Way halo line emission.

We do not expect line emission from $M_{200c} < 10^{12} M_{\odot}$ haloes to be detectable with these instruments. Though nominally, it seems like this can be overcome by stacking in halo centres, emission here will be difficult to attribute unambiguously to the CGM (as opposed to e.g., hot ISM). Moreover, in EAGLE this emission is largely due to gas that is at potentially unphysical temperatures and densities as a result of direct heating by the subgrid model for supernova or AGN feedback (Appendix [4.A](#)).

4.4.3 Origin of the emission

To investigate which gas in haloes is responsible for the line emission, we examine emission-weighted temperatures, densities, and metallicities as a function of distance to the central galaxy. For each emission line, we first make a histogram of the SPH particles around each individual central galaxy, binning them by distance to the central galaxy (normalized by R_{200c}) and e.g., temperature, weighted by luminosity. To combine these into a profile for the

whole halo mass bin, we first normalize the individual halo profiles by the total luminosity of the emission line within R_{200c} . By doing so, we ensure that stacked profiles are representative of the haloes in each mass bin, rather than being dominated by the most luminous haloes in the bin. We then add the normalized, luminosity-weighted temperature distributions in each radial bin for each halo in a mass bin. This gives us a luminosity distribution in radius and temperature that is representative for each halo mass bin. Fig. 4.7 shows the median of these distributions at each halo mass and radius. We show the median of the stacked distribution in temperature, hydrogen number density, and metallicity (oxygen mass fraction) in each radial bin, for each halo mass bin. The enclosed luminosities $L(< r) / L_{200c}$ are averaged over the haloes in each mass bin.

The middle and right columns of Fig. 4.7 show the emission-weighted profiles for O VII (middle column) and O VIII (right column). These can be compared to the similarly obtained mass- and volume-weighted gas properties shown in the left column. Star-forming gas is included at 10^4 K. The data in the first column of Fig. 4.7 are the same as the first column of Wijers et al. (2020, figs. 12 and 13)⁹

In the top row of Fig. 4.7 we see the enclosed mass and luminosity as a function of distance to the central galaxy in different halo mass bins. We see that within R_{200c} , the emission tends to be more concentrated in halo centres than the mass. The upturns outside R_{200c} for some lower-mass haloes are stronger in these stacks than in a few individual example haloes we examined, and are typically absent or much less prominent if we omit the normalization by the luminosity enclosed in R_{200c} . This points towards larger upturns being mostly a result of stacks including haloes with low overall luminosities, possibly combined with their positions being correlated with higher-mass haloes producing more emission.

The emission-weighted metallicities are similar to those found by Van de Voort & Schaye (2013) using the OWLS simulations. The trends are also similar to those for the parent ion-weighted metallicities of Wijers et al. (2020); indeed, emission and absorption share the same $\propto Z$ scaling, so a similar bias would be expected. Emission is biased towards high-metallicity gas. The emission-weighted median metallicity is higher than the mass- or volume-weighted value, and declines less strongly with radius. The difference with the mass- and volume-weighted metallicities indicates that there is substantial scatter in gas metallicity at large radii. Depending on the emission line, the emission-weighted metallicity varies between declining and mostly flat with radius.

The emission-weighted densities tend to broadly follow the trend of the volume-weighted profiles, but show a bias towards higher densities. Again, this is consistent with the findings of Van de Voort & Schaye (2013). The bias is particularly large at high halo masses and outside R_{200c} , where the mass- and volume-weighted densities are lowest. Such a bias is expected given that emission scales with the density squared. However, we do notice that where the mass- and volume-weighted densities differ, at small radii for lower-mass haloes, the emission-weighted densities are lower than the mass-weighted densities. This is because those high densities coincide with too low temperatures (the cool ($\sim 10^4$ K), dense gas phase) to produce significant emission in these soft X-ray lines (Fig. 4.1).

Indeed, the emission-weighted temperatures are consistently high, and are not very sensitive to the overall gas temperature. Rather, this emission traces whatever gas is present around its emissivity peak temperature. However, we do see some emission from below the emissivity peak and at lower densities in the lowest-mass haloes we consider, suggesting

⁹There is only a very minor technical difference in that here, we use the actual hydrogen number density, calculated from SmoothedElementAbundance/Hydrogen, rather than the mass density and an assumed (primordial) hydrogen mass fraction of 0.752.




Figure 4.7: Radial profiles of enclosed mass and line luminosity (top row), and mass-, volume-, and luminosity-weighted gas temperature (second row), density (third row), and metallicity (fourth row). Mass- and volume-weighted median quantities are shown with solid and dashed lines, respectively. The columns corresponds mass/volume (left), O VII (r) (middle) and O VIII (right). The median profiles are obtained by stacking the profiles of individual haloes after normalizing such that each halo carries equal weight.




Figure 4.8: Mass-, volume-, and luminosity-weighted gas temperature (top row), density (middle row), and metallicity (bottom row) for all non-star-forming gas within R_{200c} as a function of halo mass. Dark lines and shading show the median and scatter (80 per cent) of the individual halo averages, respectively, while lighter lines show weighted averages over the individual halo values at each halo mass. From left to right, different columns show quantities weighted by mass/volume, O VII (r) luminosity and O VIII luminosity. Mass- and volume-weighted quantities are shown in solid, green and black, dashed lines, respectively. Red lines in the temperature plots (top row) show T_{200c} in the left panel, while in the remaining panels of the top row they indicate the peak CIE temperature (solid) and the range where the emissivity is at least 0.1 times the maximum in CIE (dashed).

photo-ionization is a factor there. We note that most of the emission within R_{200c} comes from radii where collisional processes dominate. This also agrees with the findings of [Van de Voort & Schaye \(2013\)](#).

In [Fig. 4.8](#) we explore these trends as a function of halo mass. Here, the temperature, density, and metallicity are averaged within R_{200c} for each halo, and we show the trends of these weighted averages with halo mass. As [Fig. 4.7](#) would suggest, the mass- and volume-weighted gas temperatures roughly follow T_{200c} , but are somewhat larger.

For the emission lines, we find trends for He- α -like resonance lines and Lyman- α -like lines that are illustrated by the O VII and O VIII lines we show. The Fe L-shell lines follow trends similar to the He- α -like lines. The emission-weighted temperature lies close to the CIE emissivity peak. The emission-weighted temperature does follow the halo virial temperature trend over a limited mass range, but within the constraints of the line emissivity peak.

The volume-weighted density reflects the halo (non-star-forming) gas fraction. We see that this increases with halo mass. Interestingly, the scatter decreases around the same halo mass where the scatter in line luminosity decreases (Fig. 4.5). This is in line with the trend Davies et al. (2019) found with total soft X-ray luminosity (their fig. 4), where the scatter at fixed halo mass was driven by the halo gas fraction, especially for haloes hosting galaxies with stellar luminosity $\gtrsim L_*$. The mass-weighted density remains high relative to the volume-filling density even at the largest halo masses.

We see a generally rising trend of emission-weighted median density with halo mass, with emission tracing higher densities than mass at high halo masses, where the virial temperature exceeds the emissivity peak temperature. For the different lines, the emission-weighted density becomes roughly constant with halo mass in the regime where the emission-weighted temperature is similar to the emissivity peak temperature. For the K- α lines, and some other lines with relatively high peak emissivity temperatures, the emission-weighted mean temperature falls above the 80 per cent halo-to-halo scatter range here, which is typically true for the density in a much larger halo mass range. This indicates that the brightest haloes here differ considerably from the typical haloes in the gas that causes their emission. Comparing the radial profiles in Fig. 4.7 to similar profiles obtained without the individual halo normalization step (not shown), suggests that this difference is driven by emission-weighted temperatures in the halo centres. The emission in the brightest haloes at these low masses is often driven by direct heating of gas by feedback, meaning the luminosity predictions in these brightest, low-mass haloes are less reliable (Appendix 4.A).

The mass-weighted metallicities are likely higher than the volume-weighted ones simply because dense gas tends to be closer to the galaxies where the metals are produced. The emission-weighted average halo metallicities are of order Z_\odot , which is well above the mass- and volume-weighted ones; this is expected, since metal-line emission is inevitably biased towards metals.

For the lines with narrower emissivity peaks (He- α -like and Fe L-shell), the emission-weighted metallicity tends to increase with halo mass starting roughly where the halo gas temperature (mass or volume weighted) exceeds the peak emissivity temperature of the line. For the K- α lines, the metallicity tends to flatten out at this mass instead. At similar halo masses, the emission-weighted densities rise. Hence, the larger the factor by which the typical temperature exceeds the value for which the emissivity peaks, the more highly biased the emission-weighted density and metallicity tend to become.

4.5 Discussion

4.5.1 The EAGLE simulations

Current hydrodynamical simulations lack the resolution to model feedback processes from first principles and must hence make use of subgrid models that are calibrated to some observables. In the case of large-volume simulations such as EAGLE, the model is calibrated to the observed low-redshift galaxy mass function and sizes. However, models with widely varying in- and outflow rates can result in the same galaxy masses (e.g., Mitchell et al. 2020; Mitchell & Schaye 2021). Indeed, CGM predictions can vary dramatically between simulations that reproduce the galaxy mass function.

For example, Davies et al. (2020) compared the gas mass content of the CGM in EAGLE and IllustrisTNG 100-1. They found that at halo masses $M_{200} \gg 10^{12} M_\odot$, the EAGLE CGM contains a somewhat higher gas fraction, while the IllustrisTNG CGM contains much more

gas at masses $\lesssim 10^{12} M_{\odot}$. In fact, the IllustrisTNG gas fractions have a minimum (as a function of halo mass) at $M_{200} \sim 10^{12.5} M_{\odot}$. In EAGLE, the gas fraction consistently increases with halo mass, although the slope does change at $M_{200} \sim 10^{12.5} M_{\odot}$. This difference in low-mass halo gas fraction, and corresponding soft X-ray luminosity (Davies et al. 2019), likely drives differences in O VII emission from $M_{200c} \approx 10^{11.5} M_{\odot}$ haloes found by Simionescu et al. (2021).

Schaye et al. (2015) found that the gas fractions in high-mass EAGLE halos ($M_{500, \text{hse}} > 10^{13.5} M_{\odot}$) are too high at fixed halo masses. The soft X-ray luminosity is ≈ 0.3 dex too high for fixed spectroscopic temperatures $\gtrsim 1$ keV. Barnes et al. (2017) investigated the X-ray properties in more detail and found that these most massive haloes in the EAGLE 100^3 cMpc^3 Reference model contain too much gas at fixed M_{500c} , and are a bit too cool. The soft X-ray luminosities (0.5–2 keV) are about right though, as are the metallicities (iron), so the metal emission line predictions might be realistic despite the simulation’s flaws. For halo masses $\lesssim 10^{13} M_{\odot}$, X-ray observations are currently insufficiently sensitive to test the predictions.

The sensitivity of the CGM to in- and outflows (e.g., Mitchell & Schaye 2021) makes it a useful testbed for models of galaxy formation, which motivates studies like ours. Oppenheimer et al. (2020b) predict that, with eRosita stacking, the difference between the IllustrisTNG 100-1 and EAGLE CGM soft X-ray emission predictions for nearby $\sim L_*$ galaxies should be detectable, as well as the connection between quenching and halo gas fraction (via the central galaxy star formation rate and soft X-ray surface brightness).

If the numerical resolution is changed in a large-volume galaxy formation simulation like EAGLE, the subgrid prescription effectively changes since it moves to a different scale and will generally result in different CGM gas flows (see the discussion in §2 of Schaye et al. 2015). Hence, we expect the predictions for CGM emission to also change with the resolution of the simulation. This will remain true even if the subgrid parameters are recalibrated to match the galaxy mass function, since we know that calibration on galaxy properties leaves room for a wide range of CGM predictions.

We test the effect of simulation resolution on the surface brightness profiles shown in Fig. 4.6 in Appendix 4.B. We test for this using a recalibrated, higher-resolution version of the EAGLE simulation, run in a 25^3 cMpc^3 volume: Recal-L025N0752 (Schaye et al. 2015). This simulation has 8 (2) times better mass (spatial) resolution than our fiducial simulation Ref-L100N1504. Because we are testing the resolution dependence in a smaller volume, our sample of high-mass haloes is very small. There are no haloes with $M_{200c} > 10^{13.5} M_{\odot}$, and only one with $M_{200c} > 10^{13} M_{\odot}$.

For haloes with $M_{200c} \approx 10^{11.5} - 10^{13} M_{\odot}$, the properties of the CGM depend somewhat on the resolution and its implications for feedback, but these effects are relatively small. For those haloes, the median and mean surface brightness profiles typically differ by < 0.5 dex between the simulations with these different resolutions, across the different emission lines. This difference is small compared to the range of surface brightness values in the 0.1–1 R_{200c} impact parameter range. The high-resolution median surface brightnesses tend to be higher than the Ref-L100N1504 values (those in Fig. 4.6), meaning the predictions in Fig. 4.6 for the detectability of individual haloes are, in this sense, conservative.

At lower halo masses ($M_{200c} \lesssim 10^{11.5} M_{\odot}$), the intrinsic properties of the haloes (CGM gas fraction and temperature) differ more between haloes at different resolutions, and the convergence of the mean surface brightness profiles becomes poorer, particularly in the central regions. This motivates the range of halo masses we show throughout this work.

4.5.2 Literature comparison

Other predictions of CGM soft X-ray emission lines have been made. [Van de Voort & Schaye \(2013\)](#) used the $100h^{-1} \text{ cMpc}^3$ OWLS simulations ([Schaye et al. 2010b](#)) to predict the CGM emission from a number of soft X-ray emission lines (C IV, N VII, O VII, O VIII, and Ne X), and compared these to estimated detection limits of a set of X-ray instruments. They used different halo mass bins and have more high-mass haloes due to their larger simulation volume. We note that their mass resolution is nearly two orders of magnitude lower than for EAGLE and that their fiducial model does not include AGN feedback.

Although these differences make direct comparison difficult, a few trends are clearly similar. The hierarchy of line brightnesses for the five soft X-ray lines is similar, and the brightnesses fall in a similar range. The shapes of the profiles are, however, different. The [Van de Voort & Schaye \(2013\)](#) profiles show a central core at $M_{200c} = 10^{12} - 10^{13} M_{\odot}$, while the surface brightness continues to rise towards the smallest radii at $M_{200c} = 10^{13} - 10^{14} M_{\odot}$, and there is a central peak in surface brightness at $M_{200c} = 10^{14} - 10^{15} M_{\odot}$. We see a trend in the opposite direction: the lowest-mass haloes have the most centrally peaked emission, while at $M_{200c} = 10^{13} - 10^{14} M_{\odot}$, there is more of a core within $\sim 0.1 R_{200c}$. Though our $M_{200c} > 10^{14} M_{\odot}$ sample is small (9 haloes), we see a clear dip in surface brightness in the centres of these most massive haloes. This is physically plausible because the centres of these EAGLE haloes are their hottest parts (Fig. 4.7), and these haloes are hotter than ideal for producing these lines overall (Fig. 4.1). We saw a similar effect in soft X-ray absorption in [Wijers et al. \(2020\)](#).

[Simionescu et al. \(2021\)](#) compare predictions for O VII CGM emission in IllustrisTNG 100-1 ([Pillepich et al. 2018](#)) and EAGLE in their fig. 3. The EAGLE profiles were calculated with a different set of line emission tables than we use. This should, however, not make a big difference for the predictions, because the emission mostly comes from CIE gas, so the UV/X-ray background is not important¹⁰. The EAGLE and IllustrisTNG predictions are similar at a halo mass of $10^{12.5} M_{\odot}$, but differ substantially at $10^{11.5} M_{\odot}$: the IllustrisTNG haloes are much brighter in their centres, but the emission drops off more rapidly with impact parameter, leaving the EAGLE haloes brighter at the virial radius. Note that these low-mass EAGLE haloes are not detectable in O VII emission at any radius, at least with the instruments considered in this work. The predictions from the Illustris simulation (the predecessor of IllustrisTNG; [Vogelsberger et al. 2014](#)) differ substantially from the EAGLE and IllustrisTNG predictions at both halo masses.

In agreement with [Van de Voort & Schaye \(2013\)](#), we find that metal emission-line-weighted metallicities are biased high relative to mass- and volume-weighted metallicities, across the halo masses we consider (Fig. 4.7). The bias increases with distance from the central galaxy, as the line-weighted values are a roughly constant $Z \approx 0.3 - 1 Z_{\odot}$ outside $\approx 0.3 R_{200c}$, while the mass- and volume-weighted metallicities decrease with distance to the central galaxy out to at least $\approx 3 R_{200c}$, reaching $Z \approx 0.03 - 0.1 Z_{\odot}$ at R_{200c} . These emission-line-weighted metallicities are similar to the metallicities [Barnes et al. \(2017\)](#) found from mock, (broadband) X-ray observations of their C-EAGLE clusters ($M_{500c} = 10^{13.9} - 10^{15.1} M_{\odot}$). These are a set of simulated clusters, which use a variation of the EAGLE code similar to the Reference model we used in this work: AGNdT9 ([Schaye et al. 2015](#)). The values they find from their mock observations are consistent with metallicities measured from observations.

Various metallicity measurements from ICM emission spectra indeed indicate that the

¹⁰Median profiles for the K- α and He- α -like lines computed from the two sets of tables we use in this paper differ by ≈ 0.1 dex in the potentially observable (surface brightness $> 10^{-2} \text{ photons s}^{-1} \text{ cm}^{-2} \text{ sr}^{-2}$) regime. This is consistent with the differences we find between the emissivities as a function of temperature in CIE.

metallicity (iron) of the ICM is roughly constant from $\approx 0.3 R_{200c}$ to the largest radii where there are measurements, $\sim R_{200c}$ (e.g., fig. 3 of the review by [Mernier et al. \(2018\)](#)). [Martizzi et al. \(2016\)](#) studied cluster (halo mass $\sim 10^{15} M_{\odot}$) metallicities using a different set of simulations. Their X-ray-emissivity-weighted metallicities were lower than metallicities measured from observations, but the metallicity bias is similar to what we find in lower-mass haloes in EAGLE: it increases with distance to the central galaxy. The emissivity-weighted values are roughly constant, while the mass-weighted metallicity decreases with distance to the central galaxy. They attribute this difference to the fact that their metals are concentrated in dense gas clumps at large distances. (Their emissivity weighting is based on bremsstrahlung density and temperature scalings, and does not depend on metallicity.) [Biffi et al. \(2018\)](#) similarly found flatter metallicity profiles when weighting by emission instead of mass; they additionally used 3-dimensional distances for their mass-weighted profiles and impact parameters for the emission-weighted profiles. Their emission-weighted profiles matched observations.

In the CGM of simulated $\sim L_*$ galaxies, [Crain et al. \(2013\)](#) found a similar metallicity bias in broadband X-ray emission. Like what we find for X-ray emission lines, this bias increases with distance to the central galaxy, as luminosity-weighted metallicities remain constant around R_{200c} , or decrease less strongly with distance than the mass-weighted values.

Our results suggest that the biases in metallicity measurements from ICM X-ray emission extend to the CGM of haloes that are three orders of magnitude less massive than those clusters. This highlights the value of numerical simulations in the interpretation of observational findings. We do note that we compare mass- and volume-weighted metallicities to values weighted by metal line emission, not total X-ray emission. This likely results in at least somewhat larger differences (biases) than would result from observations.

4.6 Conclusions

We have investigated soft X-ray metal-line emission from the CGM and the IGM in the 100^3 cMpc^3 EAGLE simulation, for a sample of bright lines including the brightest ones we expect. We investigated K- α and He- α -like emission lines, and a few iron L-shell lines, with rest-frame energies between 0.3 and 2 keV and emissivity peaks in CIE between $\approx 10^6$ and 10^7 K. Our main conclusions about the line emission are:

- Line emission is dominated by haloes, i.e., CGM, rather than by the interhalo IGM (Figs. [4.3](#) and [4.4](#)). The emission is more concentrated in haloes than the ions producing this emission, where we have data for both (O VII, O VIII, Ne IX, and Fe XVII from [Wijers et al. \(2020\)](#)). The difference is most likely due to the stronger density dependence of emission compared to ion density.
- The brightest emission comes from the O VIII K α doublet, and the other K α lines have bright peak surface brightnesses as well. The brightest He- α resonance lines come from O VII. The Fe L-shell lines reach peak surface brightnesses similar to or somewhat larger than that of the O VII resonance line, in a narrower range of halo masses (Fig. [4.6](#)).
- There is large scatter in line luminosity at fixed halo mass. The scatter decreases towards higher halo masses and median luminosities (Fig. [4.5](#)).
- Line emission originates mainly from gas at CIE temperatures, even far from the central galaxy. For K- α lines, emission can originate from hotter gas in high-mass haloes

(where T_{200c} is above the emissivity peak temperature). Photo-ionization by the UV/X-ray background may be important in some low-mass haloes, but it does not seem to matter for emission from haloes we might be able to detect in line emission (Fig. 4.8).

- Line emission is biased to high-metallicity gas (Figs. 4.7 and 4.8), in agreement with the findings of Van de Voort & Schaye (2013). This bias is similar to what we found for metal absorbers in Wijers et al. (2020). Others have found similar metallicity biases for broadband X-ray emission in clusters (e.g., Martizzi et al. 2016; Barnes et al. 2017; Biffi et al. 2018) and the CGM of $\sim L_*$ galaxies (Crain et al. 2013).
- We have also examined trends of the halo luminosity and surface brightness in various emission lines with halo mass. The primary driver of these trends is how close the temperature of the halo ($\sim T_{200c}$) is to the temperature where the emissivity of the line peaks. This is the ‘virial temperature thermometer’ effect discussed by Oppenheimer et al. (2016) in the context of O VI absorption.
- Secondly, the shape of the emissivity curve (as a function of temperature in CIE) matters. For the Fe L-shell lines, the emissivity peaks are narrow (Fig. 4.1), and surface brightnesses depend strongly on halo mass. For the He- α -like lines, the emissivity peaks are less narrow, and the dependence of surface brightness on halo mass is less strong. The K- α lines have the widest peaks, with emissivity decreasing slowly towards high temperatures. The surface brightnesses of these lines generally keep increasing with halo mass, and only start to peak or plateau at $\sim 10^{14} M_\odot$, where the sample size is severely limited by the volume of the simulation.

We also assessed the prospects for detecting line emission from the CGM with different instruments. We did this by calculating simplified and generally optimistic estimates of minimum observable surface brightnesses (§4.3). We ignore any systematic errors and define emission as detectable if it would constitute a 5σ detection as determined from the signal to noise ratio. For the noise, we include estimates of instrumental and astrophysical backgrounds. We use the limits for exposure times and spatial binning $\Delta t \Delta \Omega = 1$ and 10 Ms arcmin^2 . We compare these detection limits to the surface brightness profiles of typical CGM emission and stacked CGM emission in Fig. 4.6

- With the XRISM Resolve instrument, we will likely be able to observe some CGM emission in the brightest lines, from haloes with $M_{200c} \gtrsim 10^{13} - 10^{13.5} M_\odot$.
- With the Athena X-IFU and the Lynx Main Array, it will be possible to detect line emission from haloes down to $M_{200c} \approx 10^{12} - 10^{12.5} M_\odot$. For haloes with $M_{200c} \gtrsim 10^{13} M_\odot$ it may even be possible to detect the outer CGM of the haloes in O VIII and O VII emission lines, with very long exposure times (1–10 Ms) or stacking. The inner CGM of $\sim L_*$ galaxies may also be accessible with these two ions, long exposures, and stacking.
- For emission lines below $\approx 1 \text{ keV}$, the Lynx Ultra-High Resolution Array will provide increased sensitivity. With this instrument, imaging the inner CGM of galaxies down to $\sim L_*$ masses will be possible, without stacking but with long exposure times, in N VII and C VI K- α emission. In O VIII and O VII line emission, less extreme exposure times or spatial binning are expected to be sufficient.

Acknowledgements

We thank Luigi Piro for help with the X-IFU detection limits, Aurora Simionescu for help with the XRISM Resolve limits, and Alexey Vikhlinin for help with the Lynx limits. We thank Ben Oppenheimer for useful discussions.

We used the PYTHON SHERPA package (Freeman et al. 2001; Doe et al. 2007) and ASTROPY¹¹ (Astropy Collaboration et al. 2013, 2018) to handle the response matrices and backgrounds for detection limits. Other PYTHON packages we used include NUMPY (Oliphant 2006), SCIPY (Jones et al. 2001), H5PY (Collette 2013), and MATPLOTLIB (Hunter 2007), and the IPYTHON (Pérez & Granger 2007) command-line interface. We thank Paul Tol for making his colour blind friendly colour schemes publicly available.

This work is partly funded by Vici grant 639.043.409 and research programme Athena 184.034.002 from the Dutch Research Council (NWO). This paper is supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871158, project AHEAD2020. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

Data availability

The EAGLE halo and galaxy catalogues (McAlpine et al. 2016) and the complete simulation outputs (snapshots; The EAGLE team 2017) are publicly available at <http://icc.dur.ac.uk/Eagle/database.php>. The response files and backgrounds for the different instruments are available as indicated in the text. Data and plots from this work are available from the corresponding author on reasonable request.

Appendix

4.A Gas directly heated by feedback

In this appendix, we investigate the effect of gas that has been directly heated by stellar or AGN feedback on the surface brightness profiles. Feedback in EAGLE is implemented by a stochastic energy injection, causing a fixed temperature increase of $10^{7.5}$ or $10^{8.5}$ K in the directly heated gas particles for supernovae and AGN, respectively. These values are motivated by numerical considerations and calibration of galaxy population properties, not by expected physical temperatures of e.g., supernova bubbles, which remain unresolved. Therefore, if the surface brightness profiles we find were dominated by this directly heated gas, then the profiles may not be a realistic prediction of the EAGLE simulation.

We test the effect of this directly heated gas by making profiles excluding it. For this, we use the maximum past temperature of each gas particle, and the redshift at which that maximum was achieved, which are tracked by the simulation. We refine our selection by inspecting phase diagrams: the distribution of gas mass in density-temperature space. We compare all gas in the simulation to the distribution of gas that has maximum temperatures

¹¹<http://www.astropy.org>




Figure 4.9: Examples of the effect that gas particles that have recently been directly injected with feedback energy have on emission line surface brightness profiles. The dark blue lines (‘all gas’) match those of Fig. 4.6. Solid lines are medians of the annular average profiles around individual haloes, dotted lines show the means of these profiles. The lines in other colours show the profiles obtained by excluding gas directly heated by feedback, less than 3, 10, or 30 Myr ago. We show the profiles for two emission lines: O VIII in the left two panels and Si XIII in the right two panels. For each emission line, we show two halo mass ranges: $M_{200c} = 10^{12} - 10^{12.5} M_{\odot}$ in the leftmost and centre-right panels, and $M_{200c} = 10^{13} - 10^{13.5} M_{\odot}$ in the centre-left and rightmost panels. The emission lines and halo mass ranges (in $\log_{10} M_{\odot}$) are indicated in the panels. The top panels show the surface brightness profiles, the bottom panels show the differences with the respective mean and median ‘all gas’ profiles. The directly heated gas can have a substantial effect on the emission in the halo centre ($\lesssim 0.1 R_{200c}$), and at halo masses where the virial temperature is too low for the emission line to be strong. While the effect on the medians is modest, it can be large for the mean profiles. At halo masses sufficiently high for the virial temperature to be \gtrsim the peak emissivity temperature, the effect of directly heated gas is small.

$\log_{10} T - \log_{10} T + \Delta$ corresponding to each type of feedback as a function of the time since the maximum temperature was reached. Much more gas is directly heated by supernovae than by AGN, and its temperatures are closer to the emissivity peak temperatures of our emission lines, so the details of the AGN-heated gas selection are less important than those of the supernova-heated gas. Using the phase diagrams, we estimate which maximum temperatures and time lags include the high-density and high-temperature gas that has just been heated, and not too much of the gas that forms the bulk of the mass distribution in EAGLE. This is because, after enough time has passed, the predictions for the temperature of the gas reflect the properties of the bulk outflows and are less sensitive to the precise manner in which the energy was injected into individual particles. We assume that gas that reached a maximum temperature between $10^{7.5}$ and $10^{7.7}$ K was heated by supernova feedback, and that a maximum between $10^{8.5}$ and $10^{8.7}$ K means the gas was heated by AGN feedback. However, the supernova feedback temperatures can also be reached by virialized gas at high halo masses (Fig. 4.7). We estimate that gas at densities $n_{\text{H}} \lesssim 10^{-2} \text{ cm}^{-3}$, and temperatures $\gtrsim 10^{7.4}$ K, is part of a continuous distribution of gas, heated by e.g., virial shocks instead of supernovae. Therefore, we do not exclude this diffuse gas from the ‘no direct heating’ profiles.

We show the resulting profiles, excluding gas that was heated less than 3, 10, or 30 Myr ago, in Fig. 4.9. To illustrate the general trends, we show profiles for two emission lines and two halo masses. Firstly, in the halo centres (impact parameters $\lesssim 0.1 R_{200c}$) the effects of the directly heated gas can be quite large for both the mean and median profiles. However,

this is the region where, in observations, the emission from the CGM would be difficult to distinguish from that of the central galaxy (e.g., the hot ISM). Secondly, although at larger impact parameters (up to R_{200c}) the effects of direct heating can be severe for mean profiles, this is limited to halo masses which produce little emission overall in that emission line.

We also looked at other lines and halo masses than plotted in Fig. 4.9. For impact parameters $\approx 0.1-1 R_{200c}$, we find that the differences in the mean and median profiles are typically not worse than those in the leftmost panel of Fig. 4.9 at all halo masses we investigate for the carbon, nitrogen, and oxygen lines. For the neon and iron L-shell lines, and the Mg XI (r) line, this difference threshold is met in $M_{200c} \gtrsim 10^{12.5} M_{\odot}$ haloes. For the Mg XII K- α and Si XIII (r) lines, the threshold lies at $M_{200c} \gtrsim 10^{13} M_{\odot}$.

This means that, at halo masses for which we predict the CGM to be observable (median profiles in Fig. 4.6), our predictions are not very sensitive to the direct heating of gas by feedback. Where haloes only seem to be observable within $0.1 R_{200c}$ (typically marginally), the surface brightnesses might however be artificially high due to the way feedback is implemented in EAGLE. The same is true for halo masses that seem observable only in mean stacks, especially in halo centres, but where the stacked mean surface brightness is much higher than the median surface brightness.

4.B Numerical convergence

In this section, we discuss the convergence of the surface brightness profiles with the resolution of the simulation. In order to test this, we compare surface brightness profiles from two EAGLE volumes: Ref-L025N376, and Recal-L025N0752 (Schaye et al. 2015). Both have a volume of 25^3 cMpc^3 , which is smaller than the 100^3 cMpc^3 of the main simulation we use throughout this work (Ref-L100N1504). The mass (spatial) resolution of the Recal-L025N0752 simulation is $8 (2) \times$ higher than that of Ref-L100N1504. Its feedback parameters were calibrated in the same way as those of Ref-L100N1504, but at its higher resolution. The Ref-L025N0376 uses the same resolution and feedback prescription as Ref-L100N1504, but in the same volume and using the same initial conditions as Recal-L025N0752.

The comparison between the Ref-L025N0376 and Recal-L025N0752 models tests the ‘weak convergence’ of the emission profiles, in the terminology of (Schaye et al. 2015). This is based on the idea that, even at fixed parameters, the effect of subgrid feedback will typically depend on the scale at which it is injected, and therefore on the resolution of the simulation. In that context, a resolution test cannot be seen independently of the feedback model, and a simulation using a similar calibration at higher resolution provides a fair test of resolution convergence.

We illustrate the level of convergence in Fig. 4.10 where we compare the profiles for the O VIII K α doublet as an example, which is representative of the level of convergence at a given halo mass for other potentially observable emission lines. In short, the profiles are well-converged in haloes of mass $M_{200c} \approx 10^{12.5}-10^{13} M_{\odot}$. In haloes with $M_{200c} \lesssim 10^{11.5} M_{\odot}$ convergence is however poor for the mean profiles in the central regions. In these low-mass haloes the CGM has not developed a hot, virialized phase (e.g., Dekel & Birnboim 2006; Kereš et al. 2009; Van de Voort et al. 2011; Correa et al. 2018), leading to very low surface brightnesses, below the predicted detection limits. For haloes with $M_{200c} \approx 10^{11.5}-10^{12.5} M_{\odot}$, convergence is reasonable given the range of surface brightnesses within R_{200c} . Differences of ≈ 0.5 dex remain, but these are small compared to the decline in surface brightness with radius and compared with the differences between the mean and median profiles.

Across halo masses and emission lines, the Recal-L025N0752 median surface brightness




Figure 4.10: A comparison of the surface brightness profiles from the Ref-L025N0376 and the Recal-L025N0752 simulations. The Ref-L025N0376 simulation uses the same model and resolution as the Ref-L100N1504 simulation used throughout this work, but it has a volume of 25^3 cMpc^3 instead of 100^3 cMpc^3 . Recal-L025N0752 is an EAGLE simulation with its feedback parameters recalibrated at its $8\times$ higher mass resolution, and the same 25^3 cMpc^3 volume as Ref-L025N0376. The number of objects in each M_{200c} bin is shown at the top right of each panel. We show profiles for O VIII in three halo mass bins. The mass ranges are indicated in the bottom left of the panels, in $\log_{10} M_{\odot}$. These O VIII profiles are representative of the level of convergence for other emission lines. Except for the mean profile at small radii and low halo masses, the results are reasonably converged.

predictions tend to be higher than the Ref-L025N0376 values. In this sense, the Fig. 4.6 predictions for detectability of individual haloes in soft X-ray line surface brightness are conservative.

The halo temperature, density, and metallicity, and their emission-line-weighted values as shown in Fig. 4.8 are reasonably converged at $M_{200c} \gtrsim 10^{11.5} M_{\odot}$. Some small differences remain: the Recal-L025N0752 haloes are typically slightly cooler and contain slightly more gas, especially at $M_{200c} \lesssim 10^{12.5} M_{\odot}$.

5

Comparing EAGLE simulation predictions to X-ray absorption line observations

Based on contributions to:

F. Nicastro, J. Kaastra, Y. Krongold, S. Borgani, E. Branchini, R. Cen, M. Dadina, C. W. Danforth, M. Elvis, F. Fiore, A. Gupta, S. Mathur, D. Mayya, F. Paerels, L. Piro, D. Rosa-Gonzalez, J. Schaye, J. M. Shull, J. Torres-Zafra, N. Wijers, & L. Zappacosta. *Observations of the missing baryons in the warm-hot intergalactic medium*. 2018, *Nature*, 558, 406–409

Sean D. Johnson, John S. Mulchaey, Hsiao-Wen Chen, Nastasha A. Wijers, Thomas Connor, Sowgat Muzahid, Joop Schaye, Renyue Cen, Scott G. Carlsten, Jane Charlton, Maria R. Drout, Andy D. Goulding, Terese T. Hansen, & Gregory L. Walth. *The Physical Origins of the Identified and Still Missing Components of the Warm-Hot Intergalactic Medium: Insights from Deep Surveys in the Field of Blazar 1ES1553+113*. 2019, *ApJ*, 884, L31

Jussi Ahoranta, Jukka Nevalainen, Nastasha Wijers, Alexis Finoguenov, Massimiliano Bonamente, Elmo Tempel, Evan Tilton, Joop Schaye, Jelle Kaastra, & Ghassem Gozaliasl. *Hot WHIM counterparts of FUV O VI absorbers: Evidence in the line-of-sight towards quasar 3C 273*. 2020, *A&A*, 634, A106

Jussi Ahoranta, Alexis Finoguenov, Massimiliano Bonamente, Evan Tilton, Nastasha Wijers, Sowgat Muzahid, & Joop Schaye. *Discovery of a multi-phase O VI and O VII absorber in the circumgalactic/intergalactic transition region*. 2021, submitted to *A&A*

ABSTRACT

We examine a number of recent comparisons between claimed detections of extragalactic soft X-ray line absorbers and predictions from the EAGLE simulations. These lines are expected to be sensitive to the warm-hot phases of the intergalactic medium and the gaseous haloes of galaxies. Generally, EAGLE seems to be consistent with these observations. However, the comparisons remain difficult. This is due to the small number of detected systems, and the strong selection effects imposed by the limits of currently available instruments. We find a hint that the strong soft X-ray absorbers studied may be too concentrated around galaxies in EAGLE, but this is very uncertain. The simulation has proven useful in verifying the reasonability of assumptions made in the modelling and interpretation of the absorption data, such as whether different ions are expected to trace the same gas phase, or how rare we think strong absorbers would be, far from any detected galaxies. The simulation can also be used to test whether a claimed detection is physically reasonable.

5.1 Introduction

Much is still unknown about the warm-hot intergalactic medium (WHIM) and the warm/hot phase of the circumgalactic medium (CGM). The intra-group medium (IGrM) and intra-cluster medium (ICM) have been detected in X-ray emission (see, e.g., the review by [Werner & Mernier 2020](#)). This has been used to constrain a number of ICM properties, such as the electron density, temperature, and metallicity of the hot gas. Additionally, ICM turbulence has been constrained in different ways (e.g., [Hitomi Collaboration et al. 2018](#); [Zhuravleva et al. 2014](#)). Some observations of X-ray emission around lower-mass ($M_{\star} \approx 10^{11} M_{\odot}$), spiral galaxies exist. Many of these observations are limited to the close environment of the central galaxy (e.g., [Bogdán et al. 2015](#)), though observations of the more extended CGM exist ([Das et al. 2020](#)).

For the CGM, around lower-mass galaxies, and at low redshift, $z < 1$, most of what we know comes from far ultraviolet (FUV) absorption lines, especially from surveys with the *Cosmic Origins Spectrograph* on the *Hubble Space Telescope* (HST-COS) (e.g., [Tumlinson et al. 2011](#); [Johnson et al. 2015](#), [2017](#)). The ions producing these absorption lines mainly arise in $\sim 10^4$ – $10^{5.5}$ K gas (e.g., [Tumlinson et al. 2017](#), fig. 6). However, at halo masses $M_{200c} \gtrsim 10^{11.5}$ – $10^{12} M_{\odot}$, we expect much of the CGM to consist of a virialized, warm-hot, volume-filling phase (e.g., [Dekel & Birnboim 2006](#); [Kereš et al. 2009](#); [Van de Voort et al. 2011](#); [Correa et al. 2018](#)), with much of its gas at temperatures $\gtrsim 10^{5.5}$ – 10^6 K (e.g., Chapter 3; [Wijers et al. 2020](#)).

The warm-hot phase, at $\sim 10^{5.5}$ – 10^7 K, has been more difficult to detect. In the observational census of baryons at low redshift by [Shull et al. \(2012\)](#), some mass was still missing. The census included galaxies, cool and warm ($\lesssim 10^6$ K) IGM, ICM, IGrM, and an uncertain CGM component. The expected baryon density comes from measurements of the cosmic microwave background (CMB; e.g., [Planck Collaboration et al. 2014](#)) and big bang nucleosynthesis (e.g., the review by [Cyburt et al. 2016](#)). In cosmological, hydrodynamical simulations, the low-redshift ‘missing baryons’ are in the warm-hot, diffuse phase, forming the WHIM (e.g., [Hellsten et al. 1998](#); [Cen & Ostriker 1999](#)) and the warm-hot phase of the CGM.

Recently, [Macquart et al. \(2020\)](#) did detect a baryon density consistent with CMB measurements using the dispersion measures of fast radio bursts (FRBs). However, due to the small number of sightlines, combined with uncertainties about the dispersion measure contributions of the Milky Way (halo) and the FRB host galaxies, the range of allowed baryon density values is still large. FRB dispersion measures are only sensitive to the (total) column density of electrons along a line of sight. This means that they can be useful in constraining the distribution of gas in the CGM and IGM (e.g., [McQuinn 2014](#); [Ravi 2019](#); [Walters et al. 2019](#)), but these measurements are not sensitive to e.g., the metallicity of the gas, or the temperature of diffuse gas. This is because hydrogen in the IGM is almost fully (photo-)ionized, and dispersion measures are not sensitive to the difference between cool, photo-ionized and warm/hot, collisionally ionized gas.

The Sunyaev-Zel’dovich (SZ) effect has also been used to measure warm-hot gas, in massive filaments, CGM, IGrM and ICM. Studies of clusters (e.g., the review by [Mroczkowski et al. 2019](#)) and massive filaments (e.g., [de Graaff et al. 2019](#); [Tanimura et al. 2019](#)) have used the thermal SZ (tSZ) effect, which is determined by the total pressure along a line of sight. To find the massive filaments, pairs of massive galaxies were stacked. This yielded detections of the massive filaments between the pairs of galaxies. The haloes of lower-mass galaxies (CGM and IGrM) were detected using a combination of data. [Lim et al. \(2018\)](#) studied the IGrM ($M_{500c} \geq 10^{12} M_{\odot}$) using the tSZ effect, and known positions of galaxy groups. In ad-

dition to tSZ data, [Lim et al. \(2020\)](#) used measurements of the kinetic SZ (kSZ) effect, as well as the positions, masses, and peculiar velocities of haloes, to constrain the temperatures and gas masses of CGM, IGrM, and ICM. These peculiar velocities are important because the kSZ effect measures the electron momentum along the line of sight, meaning that with a known velocity, the kSZ effect is a measure of the (ionized) gas mass of the halo. It has been difficult to measure the SZ effect in the CGM of galaxies roughly the mass of the Milky Way due to the large beam size (spatial resolution) of the instruments (e.g., [Mroczkowski et al. \(2019\)](#)).

Soft X-ray lines are also expected to be good probes of warm-hot ($T \sim 10^{5.5} - 10^7$ K) gas around or between galaxies (e.g., [Perna & Loeb \(1998\)](#), [Hellsten et al. \(1998\)](#), [Chen et al. \(2003\)](#), [Cen & Fang \(2006\)](#), [Branchini et al. \(2009\)](#)). There have been various attempts to detect soft X-ray absorption lines against bright X-ray sources. Some of these have been blind searches, while others searched at specific, promising redshifts. This is a way to detect the warm-hot gas, and absorption lines also provide an avenue to measure, e.g., its temperature and density (e.g., [Branchini et al. \(2009\)](#)).

These attempts have typically yielded low-significance measurements, or claimed discoveries that were later disputed. For example, in the spectrum of PKS 2155–304, [Fang et al. \(2002b\)](#) found an O VIII absorber using *Chandra* low energy transmission grating (LETG) data, at the same redshift as a small galaxy group and H I absorbers. The absorber had not been detected by [Nicastrò et al. \(2002\)](#), who analysed *Chandra* LETG data of the same source. [Rasmussen et al. \(2003\)](#) and [Cagnoni et al. \(2004\)](#) also observed this source, with the *XMM-Newton* reflection grating spectrometer (RGS), and did not find evidence for this line. [Nicastrò et al. \(2005\)](#) similarly found two combined O VII and N VII absorbers in the spectrum of Mrk 421, using ≈ 200 ks of *Chandra* LETG data. However, [Kaastra et al. \(2006\)](#) concluded that these detections were not statistically significant, due to the large number of redshift trials used in the blind search, and did not find evidence for these absorbers in *XMM-Newton* data of the same source.

[Mathur et al. \(2003\)](#) found three soft X-ray absorbers at $2-3\sigma$ significance, by searching for X-ray counterparts to six known O VI absorbers in the sightline to H1821+643. More recently, [Bonamente et al. \(2016\)](#) searched for soft X-ray absorption at specific redshifts where FUV absorbers had been found. They found likely O VIII absorption at the redshift of a broad Lyman α absorber using *Chandra* LETG observations. [Kovács et al. \(2019\)](#) found O VII absorption, attributed to the WHIM, by stacking data from different background sources and redshifts. They included 17 systems and redshifts where Lyman α had been detected near massive galaxies. [Nicastrò et al. \(2017\)](#) review these searches for the WHIM.

Around the Milky Way, warm-hot gas has been detected in soft X-ray absorption and emission lines (e.g., [Bregman & Lloyd-Davies \(2007\)](#), [Gupta et al. \(2014\)](#), [Miller & Bregman \(2015\)](#), [Das et al. \(2019\)](#)). The spatial extent of this gas is still uncertain (e.g., [Bregman & Lloyd-Davies \(2007\)](#), [Gatuzz & Churazov \(2018\)](#)), although [Miller & Bregman \(2015\)](#) did constrain this by fitting a model for the halo radial density profile to O VII and O VIII absorption and emission line measurements along many lines of sight. The detected emission and absorption lines have taught us much about the Milky Way CGM. From these lines, e.g., [Kuntz & Snowden \(2000\)](#) and [Das et al. \(2019\)](#) measured the temperature(s) of the warm/hot Milky Way CGM, e.g., [Miller & Bregman \(2015\)](#) measured its density and metallicity, and [Hodges-Kluck et al. \(2016\)](#) measured its rotation. This illustrates the potential that extragalactic absorption line measurements have to constrain the properties of a larger diversity of galaxy haloes, and of the WHIM. However, the limited number of X-ray sightlines through these extragalactic systems will make it more difficult to constrain, e.g., their rotation, especially in individual systems.

We hope to learn more about the extragalactic warm/hot, virialized phase of the CGM, IGrM, and WHIM using soft X-ray absorption lines. However, with current instrumentation, only the strongest absorbers have been detectable, the number of observations is small, and many detections are of low statistical significance or disputed. Therefore, the extant comparisons to simulations are, for the most part, a sanity check of theoretical predictions, rather than a very stringent test.

On the other hand, the comparisons can provide a sanity check for uncertain detections. The theoretical predictions are valuable for checking the reasonability of assumptions made in the modelling of absorbers, such as whether different ions are expected to trace the same gas phase. With the increased sensitivity and spectral resolution of future missions, we hope to be able to test such assumptions instead, e.g., by comparing data on more transitions (absorption lines) than we can currently detect. This can provide stronger constraints on the conditions of the gas, such as its temperature. If sufficiently sensitive instruments are available, the combination of absorption and emission is also promising.

In this chapter, we will discuss comparisons of predictions from the EAGLE cosmological simulation to a few recently claimed detections of soft X-ray absorption lines tracing the WHIM or CGM. These are summaries of the EAGLE comparisons I contributed to different papers. First, we discuss blindly detected absorbers. In §5.3.1, we discuss a claim of two blind detections of O VII absorption in the spectrum of a very bright blazar (Nicastrò et al. 2018). In §5.3.2 we discuss the further investigation of those absorbers by Johnson et al. (2019), which changed the interpretation of one of those absorbers, and cast doubt on the detection of the other. Next, we discuss soft X-ray absorbers found at the redshifts of two O VI absorbers. In §5.4.1 we discuss an O VIII and Ne IX absorber at one out of two investigated O VI redshifts (Ahoranta et al. 2020), and in §5.4.1 we discuss an O VII counterpart to the single O VI absorber in a sightline (Ahoranta et al. 2021).

When discussing distances, we indicate whether they are physical/proper (e.g., ‘pkpc’) or comoving (e.g., ‘cMpc’). The exception are centimetres, which are always physical.

5.2 Methods

We used various results from EAGLE to compare its predictions to the different observational datasets. However, the basis for all the comparisons was the same: maps of ion column densities in EAGLE. In this section, we will give an overview of the EAGLE simulations, and explain how we produce these column density maps.

5.2.1 EAGLE

The EAGLE (‘Evolution and Assembly of GaLaxies and their Environments’; Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016) simulations are cosmological, hydrodynamical simulations. Gravitational forces are calculated using the GADGET-3 Tree-PM method (Springel 2005), and hydrodynamical forces are calculated using Smoothed Particle Hydrodynamics (SPH). Specifically, EAGLE uses the Anarchy implementation of SPH (Schaye et al. 2015; appendix A; Schaller et al. 2015). For the predictions in this chapter, we used the Ref-L100N1504 simulation. This is a simulated 100^3 cMpc³ volume, with an (initial) gas mass resolution of $1.81 \times 10^6 M_{\odot}$ (Schaye et al. 2015). A Λ CDM cosmogony with cosmological parameters $\Omega_m = 0.307$, $\Omega_{\Lambda} = 0.693$, $\Omega_b = 0.04825$, $h = 0.6777$, $\sigma_8 = 0.8288$, $n_s = 0.9611$, and $Y = 0.248$ (Planck Collaboration et al. 2014) is assumed in EAGLE, and in the predictions we make from it.

Besides gravity and hydrodynamics, a number of unresolved processes need to be modelled in order to produce (realistic) galaxies. These are called ‘subgrid’ processes. Firstly, radiative cooling and heating of gas is modelled following [Wiersma et al. \(2009a\)](#). This model includes metal line cooling, using the abundances of 9 metal species tracked in EAGLE. It assumes the gas is in ionization equilibrium, including the effect of photo-ionization by a [Haardt & Madau \(2001\)](#) UV/X-ray background.

Secondly, star formation occurs in sufficiently dense ($n_{\text{H}} \gtrsim 10^{-1} \text{ cm}^{-3}$) gas, where the exact threshold depends on the metallicity of the gas ([Schaye 2004](#)). The star formation rate of this gas depends on the gas pressure, in such a way that the Kennicutt-Schmidt relation is reproduced by design ([Schaye & Dalla Vecchia 2008](#)). Stars, through AGB winds and type Ia and core-collapse supernovae, return metals to their neighbouring gas following [Wiersma et al. \(2009b\)](#). Feedback from star formation is implemented by stochastically injecting thermal energy into a neighbouring SPH particle, heating it by $10^{7.5} \text{ K}$. The probability of such a heating event is such that the average energy injected per unit stellar mass matches a set budget ([Dalla Vecchia & Schaye 2012](#)).

Finally, in sufficiently massive haloes, a black hole is formed if the halo does not already contain one. The black holes grow by accreting gas and by merging, as described by [Rosas-Guevara et al. \(2015\)](#); [Schaye et al. \(2015\)](#). AGN feedback is implemented by stochastic heating of neighbouring SPH particles, like the supernova feedback, but gas is heated by $10^{8.5} \text{ K}$.

The parameters of the feedback processes were calibrated to observations ([Crain et al. 2015](#)). Specifically, the simulations were calibrated to reproduce the redshift 0.1 galaxy stellar mass function and stellar-mass-black-hole-mass relation, and to produce reasonable galaxy sizes. The EAGLE galaxy and halo data has been made public by [McAlpine et al. \(2016\)](#), and [The EAGLE team \(2017\)](#) describes the public data release of the full simulation data (‘snapshots’).

5.2.2 Column density calculations

From this simulation, we calculate column densities by first calculating the ion content of the individual SPH particles (gas resolution elements), then dividing these ions over a dense grid of pixels based on the position and size of each SPH particle.

The ion content of an SPH particle is calculated as

$$M_{\text{ion}} = M_{\text{g}} Z_{\text{elt}} f_{\text{ion}}, \quad (5.1)$$

where M_{ion} is the total ion mass in the SPH particle. The number of ions in the SPH particle is obtained by dividing this mass by the mass of a single particle of the parent element of the ion. The mass of the SPH particle M_{g} and the mass fraction in the parent element Z_{elt} (ELEMENTABUNDANCE) are taken directly from the EAGLE output data. The ion fraction f_{ion} , which is the fraction of the parent element nuclei in the desired ionization state, is calculated from the temperature and density of the gas, and the redshift of the EAGLE snapshot. The redshift-dependence comes from the changing UV/X-ray background.

We calculate the ion fractions by linearly interpolating the ion fractions tabulated by [Bertone et al. \(2010a\)](#), as a function of log temperature, log density, and redshift. They calculated the ion fractions using CLOUDY v7.02 (last documented in [Ferland et al. 1998](#)), assuming ionization equilibrium, with photo-ionization by a uniform, redshift-dependent [Haardt & Madau \(2001\)](#) UV/X-ray background. These assumptions are consistent with those made in the EAGLE simulation for the radiative cooling ([Wiersma et al. 2009a](#)).

To calculate column densities, we first select a rectangular prism (box) within the EAGLE volume. We choose volumes that span the whole simulated volume in two dimensions ($100 \times 100 \text{ cMpc}^2$), and divide the volume into an integer number of ‘slices’ along the third axis. This third axis represents the line of sight direction, and for simplicity, we choose it to be parallel to one of the coordinate axes of the simulation (the Z-axis). Note that, relative to simulated structures, this is a random direction. We use different slice depths (numbers of slices), depending on the requirements for the comparison, but our default is 6.25 cMpc . Perpendicular to the line of sight, we divide the box into pixels of $3.125 \times 3.125 \text{ ckpc}^2$.

Given this box and grid, we first select all particles with positions (centres) between the slice boundaries along the line of sight. We assume that their gas mass is distributed according to the Wendland (1995) C2 kernel, with its size set by the smoothing length of the SPH particle. We evaluate the kernel at each pixel centre, then normalize the sum of the values to one. Each SPH particle contributes this normalized kernel value times its total ion content to each pixel in the grid. The normalization, and a minimum kernel size (radius) of half the diagonal of a pixel, ensure that mass is conserved in the projection step.

After calculating the number of ions in each rectangular prism in the grid, we simply divide by the area of the pixels to get column densities. The result of this calculation is a map of the column densities in a particular slice of the simulation.

Note that these column densities are an approximation of column densities measured in observations. There, column densities are not measured in a pre-defined spatial grid, but from measured absorption lines, in spectra taken of some background source. The predictions we make are best compared to ‘absorption systems’: clusters of absorption lines, interpreted as originating from a single object, such as the CGM of a single central galaxy. We have confirmed that the pixels we use are small enough that they match the ‘pencil-beam’ measurements of column density that would be obtained against a background point source, such as the sources we compare to here (Wijers et al. 2019 Chapter 2). The depth of the slices has some effect on the column density distribution, but this is limited to at most ≈ 0.2 dex at the higher column densities that are potentially observable (Wijers et al. 2019 Chapter 2).

5.3 A blind detection of two O VII absorbers?

5.3.1 The initial detections

Nicastro et al. (2018) used a detailed analysis of X-ray observations of the brightest X-ray blazar, 1ES 1553+113, with *XMM-Newton* to search for soft X-ray absorption along that sightline. The analysis used a total of 1.85 Ms of observations, combining different time periods. This was a blind search, meaning that X-ray absorbers were sought out at any redshift, without any prior expectations. Two absorption systems were found, at redshifts 0.36 and 0.43. They were used to constrain the density of O VII absorbers in the universe. This was also used to estimate the density of the warm/hot ‘missing baryons’, although the low number of measured data points, and necessary assumptions about e.g., the gas metallicity, make this estimate quite uncertain.

In Fig. 5.1 we show fig. 3 of Nicastro et al. (2018), comparing the number and strength of the absorbers they found to predictions from cosmological, hydrodynamical simulations. Here, we will focus on the comparison of these data to the EAGLE predictions.

The EAGLE predictions for the equivalent width distribution in Fig. 5.1 come from column density distribution predictions. To predict the column density distribution used here,

we used a column density map (§5.2.2) from EAGLE at redshift $z = 0.1$, with a slice depth of 100 cMpc. This redshift does not quite match that of the absorbers, but the difference between redshifts 0.1 and 0.5 in the abundances of O VII absorbers is only ≈ 0.2 dex at column densities of 10^{15} – 10^{16} cm^{-2} (Fig. 2.3). The column density distribution is simply a histogram of the column density values in this map. The redshift interval corresponding to this column density map is simply the redshift path covered by each pixel (from the slice depth) times the number of pixels. To convert column densities to equivalent widths, a line width (Doppler parameter) of 100 km s^{-1} was assumed.

The equivalent widths of the measured absorbers are consistent with the EAGLE predictions, although the uncertainties in the measurements are large. This is primarily due to the small number of measured absorbers. Considering the various predictions, this small number is the result of the limited sensitivity (effective area) and spectral resolution of the instruments currently available, making the minimum measurable equivalent width high, and therefore the expected number of detectable absorbers small. This is illustrated by Fig. 2.22; note that an equivalent width of 10 mÅ corresponds to a column density of $\approx 10^{15.5}$ – 10^{16} cm^{-2} . Note that this nearly 2 Ms observation of the brightest X-ray blazar represents a very optimistic case for a currently feasible blind survey.

In Fig. 2.6 of Chapter 2 (fig. 5 of Wijers et al. 2019), we showed a similar comparison. There, we also give a few more details on the predictions of Cen & Ostriker (2006) and Branchini et al. (2009). There are, however, a few differences between the two figures.

Firstly, the data points from in Fig. 5.1 taken from Nicastro et al. (2018), are different from the measurements shown in Fig. 2.6. This is because the Fig. 2.6 data points include a revision of one of the measured equivalent widths by Nicastro (2018). This revision was based on revised atomic data for the N II $K\alpha$ triplet, meaning a Galactic absorption line would be blended with the higher-redshift ($z = 0.4339$), higher-equivalent-width absorber of Nicastro et al. (2018). The absorber was still detected, but at lower significance (2.9 – 3.3σ instead of 4.1 – 4.7σ for the single line) and with a lower equivalent width (10 ± 3 mÅ instead of 14.7 ± 3.1 mÅ). These revised data points are consistent with all the simulation predictions shown in Fig. 5.1.

Secondly, we refined the EAGLE prediction for Fig. 2.6 (Chapter 2, Wijers et al. 2019). We still use column density distributions from column density maps as a starting point. However, we use two different distributions in Fig. 2.6 and compare the results: we start from maps with a 100 cMpc slice depth, as shown in Fig. 5.1 and from maps with 6.25 cMpc slice depths, dividing the simulation volume into 16 slices. Fig. 2.6 shows that this makes relatively little difference at the scale of the figure.

We also refined the way in which we convert column densities to equivalent widths. Instead of assuming a line width, we used a distribution of equivalent widths in bins of column density, calculated from synthetic spectra through the EAGLE volume (for more details, see Chapter 2, Wijers et al. 2019). The column densities and equivalent widths were summed over the full 100 cMpc sightlines we calculated the spectra for. Further improvements, resulting in the column-density-equivalent-width relations in Fig. 3.7 (Chapter 3, Wijers et al. 2020) were not included in Fig. 2.6.

In Fig. 2.21 we showed that scatter in the relation used to convert the column density distribution to an equivalent width distribution has relatively little impact on the resulting equivalent width distribution in the range shown in Fig. 5.1 relevant to this comparison. The line width (Doppler parameter) of 90 km s^{-1} , derived from the virtual spectra and used in that comparison, is similar to the value of 100 km s^{-1} used in Fig. 5.1 and the maximum effect of this difference is a small: a 10 per cent horizontal shift of the curve. All in all, the

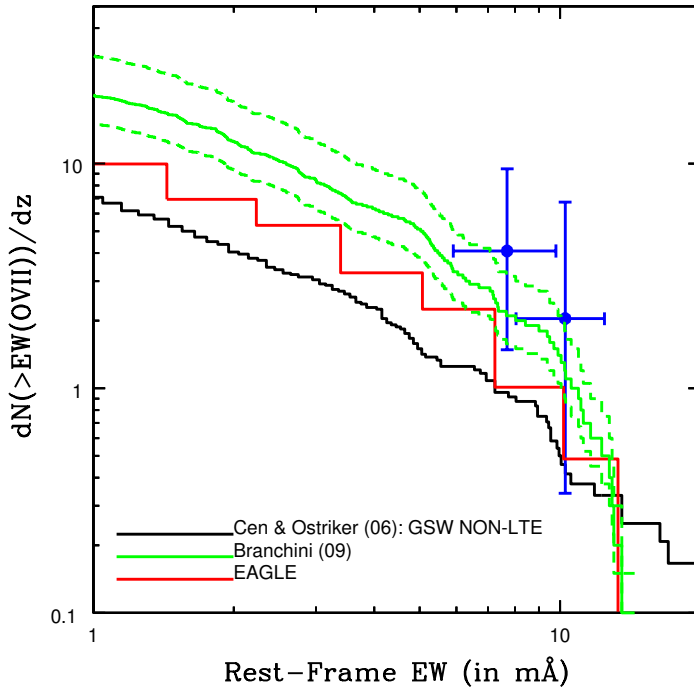


Figure 5.1: Comparison of the O VII (He- α recombination line) rest-frame equivalent width distribution measured by [Nicastrò et al. \(2018\)](#) to predictions from cosmological, hydrodynamical simulations; the figure is fig. 3 from [Nicastrò et al. \(2018\)](#). The curves and data points show the cumulative number of absorbers (i.e., the number above the equivalent width on the x-axis) per unit redshift. The blue data points show the equivalent widths of the two [Nicastrò et al. \(2018\)](#) absorbers. The black curve shows the prediction of [Cen & Ostriker \(2006\)](#), using their model with galactic stellar winds, and with O VII to total oxygen ratios calculated without assuming ionization equilibrium. The green lines show the predictions of [Branchini et al. \(2009\)](#). The red line shows the predictions from EAGLE ([Schaye et al. 2015](#)), calculated from the 100^3 cMpc^3 reference EAGLE volume at redshift 0.1. For EAGLE, column densities were calculated along 100 cMpc sightlines, and converted to equivalent widths assuming a fixed line width (Doppler parameter) of 100 km s^{-1} .

effects of these refinements are small.

Third, the selection of data from [Branchini et al. \(2009\)](#) and [Cen & Ostriker \(2006\)](#) is different. Fig. 5.1 shows an additional [Cen & Ostriker \(2006\)](#) model, where ionization fractions were computed assuming ionization equilibrium. In contrast, we only show one of the [Branchini et al. \(2009\)](#) models. This model uses their more realistic model for the gas metallicity (including scatter in the metallicity-density relation they use), which also predicts the highest absorber densities of their models in the equivalent width range shown.

All in all, with or without these refinements to the measurements and predictions, the EAGLE predictions agree with the absorber densities measured here. The large uncertainties in the measurements mean that this is mostly a sanity check, however, as the measurements are consistent with the predictions of the other simulations as well.

5.3.2 A further investigation of the absorber environments

Johnson et al. (2019) investigated these absorbers further, using a follow-up galaxy redshift survey along the line of sight to the blazar. One absorber ($z = 0.4339$) is suspected to come from the IGrM of the blazar host group. This is based on a determination of the blazar redshift, from the redshifts of Lyman- α absorbers in its spectrum (the Lyman- α forest), and a distribution of redshift differences between quasars and their highest-redshift Lyman- α absorbers. This redshift is consistent with a galaxy group found in the redshift survey (light-weighted $z = 0.433$), indicating that this group is most likely the blazar host. The group is also the only group found along the sightline, and a galaxy group is the expected environment for a blazar like 1ES 1553+113. The redshift of this group and the blazar match that of the high-redshift absorber of Nicastró et al. (2018), which therefore most likely is IGrM absorption from the blazar host. Non-detections of O VIII or O VI absorption at the same redshift indicate that the gas is collisionally ionized rather than photo-ionized, meaning absorption must come from relatively dense gas. This supports the idea that the absorber is part of the IGrM, not nearby WHIM.

Nicastró et al. (2018) did not identify this group because they used photometric galaxy redshifts in their analysis, with much larger redshift errors. They did argue the O VII absorption was not intrinsic to blazar outflows, based on its lack of time variability, or to the blazar host galaxy ISM, based on the lack of UV absorber counterparts.

The absorber intrinsic to the blazar host group cannot be used for a determination of the column density distribution of O VII and estimates of the hot gas of the universe overall, at least without additional assumptions. This is because it is not a *random* absorber along the line of sight, as it is connected to the blazar used as a backlight.

The estimated blazar host group dynamical mass is $2\text{--}5 \times 10^{13} M_{\odot}$; at those masses, an O VII absorber of $10^{15.8} \text{ cm}^{-2}$ (Nicastró 2018) is consistent with EAGLE expectations, shown as the yellow and orange lines in Fig. 3.8. Note that we expect about half the column density in the group to lie behind the blazar. In EAGLE, such an absorber is more likely in groups at the lower end of the estimated mass range, assuming the dynamical mass is similar to M_{200c} .

The EAGLE analysis of Johnson et al. (2019) focussed on the second absorber, at $z = 0.355$. The galaxy redshift survey yielded no galaxies within 630 pkpc of the absorber, and within 1000 km s^{-1} . We predicted how rare or common it would be for an absorber as strong as the measured one ($10^{15.6} \text{ cm}^{-2}$; Nicastró et al. 2018) to be found so far from any galaxies. This is generally unexpected, as the detected oxygen must have been produced in a galaxy.

In order to do this, we used *conditional* column density distributions. We used the catalogues of McAlpine et al. (2016), listing galaxy positions and stellar masses in EAGLE to create ‘masks’ for the EAGLE column density maps. A mask is a selection of pixels in the column density maps meeting particular criteria. In this case, we selected pixels at a minimum distance (630 pkpc) from any galaxy above a given stellar mass ($10^{9.7} M_{\odot}$). These are the impact parameter and stellar mass for the closest galaxy Johnson et al. (2019) found to the $z = 0.355$ absorber. We then created a column density distribution using only these pixels, and compared it to the distribution for all pixels. The ratio of the two gives the probability that an absorber of a given strength is found to be as isolated as the one claimed by Nicastró et al. (2018), once galaxy information is included.

We also tested the probability for distances $> 2273 \text{ pkpc}$ from $M_{\star} > 10^{11.2} M_{\odot}$ galaxies, based on a more distant galaxy with a larger halo virial radius, but we found the distance to the less massive galaxy to give stronger constraints. Both galaxies are estimated to lie about five virial radii from the absorber.

We used column density maps with a slice thickness of 33.3 cMpc (3 slices), at a redshift




Figure 5.2: The EAGLE O VII column density distribution function (CDDF) at $z = 0.37$, measured from column density maps with a slice thickness of 33.3 cMpc, corresponding to a Hubble flow velocity of $\approx 2000 \text{ km s}^{-1}$. The top panel shows the absorber incidence rate, with respect to the absorption length dX (eq. 2.2). In the top panel, the black line shows the total column density distribution, while the coloured lines show the density of absorbers that are *not* as isolated as the $z = 0.36$ absorber of Nicastro et al. (2018). We measure the level of isolation using a minimum distance to any galaxy above a given stellar mass. Based on two galaxies Johnson et al. (2019) found in the vicinity of the absorber, we use a minimum distance (impact parameter) of 630 (2273) pkpc from any galaxy with a stellar mass greater than $10^{9.7}$ ($10^{11.2}$) M_{\odot} . Along the line of sight, we only consider galaxies that lie in the same slice as each absorber for these distance determinations ('pos.': solid curves). However, in reality, only the redshifts of the absorbers and galaxies are known. These depend on their peculiar velocities as well as their positions along the line of sight. Therefore, we also select galaxies in the same slice in a different way to produce the dashed curves ('vel'). First, we adjust the galaxy line of sight positions according to their peculiar velocities. Then, we determine the distances (impact parameters) between absorbers and galaxies that lie in the same slice according to their modified positions. Note that we do not account for the peculiar velocities of the absorbers. The bottom panel shows the fraction of the total number of absorbers at each column density contributed by each absorber category. Thin, horizontal lines show the fraction of the sky meeting each of the non-isolation criteria. In EAGLE, O VII absorbers as strong as the $z = 0.355$ system claimed by Nicastro et al. (2018) are rarely as isolated as it is found to be by Johnson et al. (2019).

of 0.37, close to that of the absorber. The slice thickness was chosen to match the velocity range over which [Johnson et al. \(2019\)](#) searched for galaxy neighbours to the absorbers; the Hubble flow over 33.3 cMpc at $z = 0.37$ is $\approx 2000 \text{ km s}^{-1}$.

To match the galaxies to these absorbers, we primarily consider the (impact parameter) distances to galaxies that lie (spatially) between the boundaries of each simulation slice along the line of sight. Our column density maps divide the simulation only spatially, without considering the peculiar velocities of the absorbing gas.

To roughly estimate the impact of peculiar velocities, we also consider distances to galaxies selected for each slice in a different way. Here, we adjust the galaxy positions (in physical coordinates) along the line of sight by the ratio of the galaxy peculiar velocity to the Hubble constant $v_{\text{pec}} / H(z)$. We then determine whether a galaxy should be cross-matched with a particular slice by comparing the slice edges to these adjusted galaxy line of sight positions. Note that this is a conservative estimate of how bad the effect of ignoring peculiar velocities is, since they are not considered for the gas in the column density maps, which means there is a mismatch between the gas and galaxy redshifts here that would not be present in the observations, aside from errors.

We find that the probability that an absorber with an O VII column density $\geq 10^{15.6} \text{ cm}^{-2}$ would be located $\geq 630 \text{ pkpc}$ from any galaxy within $\pm 1000 \text{ km s}^{-1}$ and with $M_{\star} \geq 10^{9.7} M_{\odot}$ is 1 per cent using the position-based galaxy-absorber matching. Using the velocity-shifted galaxy matching, the probability is 3 per cent, with the difference most likely the result of galaxies being velocity-shifted away from their associated absorbers. The number density of such isolated, high-column-density absorbers in EAGLE $\text{dn}(> 10^{15.6} \text{ cm}^{-2}, \text{iso.}) / dz$ is 0.17, using the position-based galaxy matching. This means that over the $\Delta z \lesssim 0.43$ redshift path probed along this sightline, the absolute probability of detecting such an absorber is low ($\lesssim 0.07$), according to EAGLE.

For comparison, [Bonamente \(2019\)](#) estimated the chance that a noise fluctuation of at least the size corresponding to the measured absorber would arise somewhere along a redshift path of this length is 4 per cent. This is for a single line, excluding the higher-redshift line from the re-analysis on the grounds that it is not a ‘random’ absorber.

Therefore, the second absorber may be real, albeit detected at low significance, which would be in tension with EAGLE predictions because of its isolated nature. This may be due to low-mass (in EAGLE, unresolved, and in the survey, undetectable) galaxies producing metals, and possibly hot gas in outflows, that are absent in EAGLE. Alternatively, outflows from the detected galaxies may reach larger distances from the galaxies than predicted by EAGLE, or the temperature of the metal-enriched outflows may be underestimated in EAGLE.

On the other hand, the spectral feature identified as the second absorber may in fact just be noise. The absence of such an absorber would be consistent with EAGLE predictions. Note that [Johnson et al. \(2019\)](#) also analyse the distances to galaxies of broad H I, O VI, and narrow H I absorbers, and conclude that metals likely do not typically reach the distances from galaxies where the $z = 0.355$ O VII absorber of [Nicastrò et al. \(2018\)](#) was found.

5.4 Soft X-ray counterparts to FUV absorbers

Next, we consider absorbers that were found using targeted searches. The goal of [Ahoranta et al. \(2020\)](#) and [Ahoranta et al. \(2021\)](#) was to detect X-ray absorption by searching at the redshifts of known FUV absorbers (H I and O VI). They started from a sample of 16 sightlines with UV and X-ray spectra with sufficiently high signal-to-noise ratios for their analysis (J. Ahoranta, priv. comm.). In the UV spectra, there were 11 O VI absorbers, and 12 broad H I

absorbers (Doppler b parameters $> 70 \text{ km s}^{-1}$). The line width criterion for H I was used to select warm gas; though the broadening may not be thermal, an absorber with a narrow line is definitely not warm. They searched for a variety of X-ray absorption lines at the redshifts of these absorbers, including O VII, O VIII, Ne IX, and Ne X.

They investigated all 23 redshifts, and found X-ray counterparts to two of the O VI absorbers. The number of absorbers studied here is therefore small, and the selection effects are strong, as current X-ray instruments are only sensitive to the strongest absorption lines. No counterparts to broad H I absorbers were found. Because we only have two detected counterparts, we do not know whether this is a coincidence, or a result of different X-ray counterpart populations to broad H I and O VI. Therefore, it is uncertain whether the strength of the selection effect should be calculated relative to the number of O VI redshifts investigated or the total number of redshifts. However, given the uncertainties that come with the small numbers of detected absorbers, this difference is relatively small. We estimate that the detected absorbers represent the $\approx 10\text{--}20$ per cent strongest counterparts to O VI absorbers in the population.

In these works, the absorption systems have measured column densities for at least two ions (O VI and the X-ray absorber), and upper limits on others. Because the O VI redshifts were used as priors in the X-ray absorber search, we are not looking at random X-ray absorbers. We include this selection effect in the comparison to EAGLE by examining the properties of absorption systems with O VI column densities that match the measured values. We note that ‘counterparts’ in EAGLE may have arbitrarily low column densities, including zero, so the EAGLE predictions properly account for both undetectable and absent X-ray counterparts.

Our starting point for these comparisons is a set of EAGLE column density maps for each ion that was measured or constrained observationally. Here, an absorption system is defined by a pixel in a column density map of a given ‘slice’ of the simulation volume (see §5.2.2). We ‘match’ column densities for the different ions simply by comparing column densities in matching pixels. For these works, we used slices with a depth of 6.25 cMpc, and a simulation ‘snapshot’ at redshift $z = 0.1$.

The redshift is close to the redshifts of the X-ray absorbers that were found, as well that of an O VI absorber in the 3C 273 spectrum for which no X-ray counterpart was found. The path length for the initial 16 sightlines was $\Delta z \approx 2$, and for the nine sightlines where the UV absorbers were found, $\Delta z \approx 1$, so the redshift $z = 0.1$ EAGLE snapshot is reasonably representative of the redshifts searched.

The slice depth corresponds to a Hubble flow velocity $\Delta v = 404 \text{ km s}^{-1}$, though we note that our slicing only accounts for the positions of SPH particles along the line of sight, and not their peculiar velocities. Peculiar velocities can be similar to the Hubble flow over this slice depth, so this velocity range does not necessarily correspond to the velocity range an absorption system would have.

We also compare temperatures for the different absorbers. For this, we make maps of ion-weighted temperatures (for the different ions), alongside the column density maps. For these maps, we take the temperature of each SPH particle contributing ions to a given pixel, and weight the temperatures by each particle’s ion contribution. We compare temperatures for each gas phase to the ion-weighted temperature for a representative ion for that phase.

5.4.1 System 1: 3C 273, redshift 0.09

The first absorption system that we will discuss is from [Ahoranta et al. \(2020\)](#). They searched for X-ray absorption lines in the spectrum of quasar 3C 273, at the redshifts of two O VI absorbers from [Tilton et al. \(2012\)](#), and found an X-ray counterpart to one of them. They found O VIII and Ne IX absorption lines at redshift $z = 0.09$. For O VII, an upper limit was obtained. Its He- α line would be blended with a Galactic O I line at this redshift, meaning the constraints come from weaker O VII lines like He- β .

These lines were found by combining archival measurements taken at different times, with different instruments: the *XMM-Newton* RGS and *Chandra* LETG (ACIS and HRC) and HETG (MEG), for a total of ~ 1 Ms of data. To avoid artefacts, they stacked spectra from the same instrument (but taken in different observations), and fitted the stacked data from different instruments jointly. In these fits, the continuum (fitted with a spline function) was allowed to vary between the stacked spectra for the different instruments, but the absorption model (Galactic and extragalactic) was not. This is reasonable because the observations were taken at different times and 3C 273 is a variable source. The data analysis was done carefully, to avoid issues like artefacts in the stacked spectra that could be confused with absorption lines.

The data were fit with two different models for the extragalactic absorption system: a ‘slab’ model and a ‘CIE’ model. The CIE (collisional ionisation equilibrium) model assumes the X-ray absorption comes from one gas cloud in collisional ionization equilibrium, with solar relative abundances. Note that the metallicity is degenerate with the hydrogen column in these fits, since only metal lines are measured (outside our Galaxy). The slab model fits the absorption lines from the different transitions of each ion, to measure each ion’s column density independently. Because the absorption line shapes are unresolved, a line width (Doppler b parameter) $b = 100 \text{ km s}^{-1}$ was assumed in these fits. In our comparisons, we generally use the slab model constraints on the column densities of each ion, and the CIE model for the temperature of the absorption system.

The slab model yielded a 3.9σ joint detection of O VIII and Ne IX, which trace gas at similar temperatures in CIE, and are thus assumed to represent a single gas phase. The assumption of collisional ionization is motivated by our predictions in Chapter 2 ([Wijers et al. 2019](#)), where we predict that the high column density absorbers detectable with current instruments should be collisionally ionized. This model adequately fits the data. Its temperature is constrained from above by the non-detection of Fe XVII absorption, and from below by non-detections of O VII and N VII.

The X-ray CIE model predicts an O VI column density of $10^{12.2 \pm 0.2} \text{ cm}^{-2}$, well below the $10^{13.26} \text{ cm}^{-2}$ column density measured in the FUV spectrum ([Tilton et al. 2012](#)). The observed O VI lines are also narrower than predicted by thermal broadening at the best-fit CIE temperature for the X-ray lines. This is evidence that the measured X-ray absorbers are tracing hot gas, and the O VI measured in the FUV spectrum is tracing a different, warm gas phase.

In Fig. 5.3, we compare the O VIII and Ne IX column densities measured in the X-ray spectra to the predicted values for counterparts to an O VI absorber with the same column density as the $z = 0.09$ system (right panel). The X-ray ion column densities are higher than is typical in EAGLE. However, they seem consistent with the strongest ≈ 10 –20 per cent of predicted X-ray counterparts, which is the X-ray absorber selection effect we estimated from UV counterpart non-detections. The measured values are consistent with EAGLE expectations at the 1 – 2σ level.

In the right panel of Fig. 5.3, we modify the absorption system sample by adding a selec-




Figure 5.3: The measured O VIII and Ne IX column densities at $z = 0.09$, compared to EAGLE predictions for counterparts to the $z = 0.09$ O VI absorber. This is fig. 9 from [Ahoranta et al. \(2020\)](#). The measured column densities are indicated in blue and black, with 1σ error bars. The ‘slab’ column densities were measured for each ion individually, by fitting gaussian absorption lines for each transition of the ion. The ‘CIE’ column densities were instead fit to the X-ray data assuming the O VIII and Ne IX absorption arises from the same gas, in collisional ionization equilibrium (CIE) and with solar metal abundance ratios. For the EAGLE predictions, we show a histogram of individual absorption systems and contours enclosing the fractions of absorption systems indicated in the legend. We only show absorption systems with ion column densities ($\log_{10} \text{cm}^{-2}$) in the ranges indicated in the bottom right of each panel. These constraints match the measured O VI column density (1σ range) and the ‘slab’ model upper limit on the O VII column density, considering the $\approx 80\text{--}90$ per cent of O VI (UV) absorbers for which no X-ray counterpart was found.

tion criterion: the O VII column density must agree with the measured constraint. We find this has fairly little effect on the predicted O VIII and Ne IX column densities.

In Fig. 5.4 we compare the EAGLE predictions to the measured temperature of the absorber. A single CIE model was found to be inconsistent with the combined UV and X-ray data, so the FUV O VI absorber is assumed to represent a different, warm phase. The temperature of this phase is most strongly constrained by the X-ray data. With the O VI column density fixed to the best-fit value from the FUV data, X-ray upper limits on the O IV, O V, and O VII column densities constrain the temperature of the absorber.

In Fig. 5.4 we plot the warm and hot phase temperatures predicted by EAGLE, measured as the O VI and O VIII-weighted temperatures, respectively. We plot these temperatures for absorption systems meeting the various column density constraints listed in the panels. The O VI constraint is the 1σ range of [Tilton et al. \(2012\)](#), the other constraints are the upper limit and 1σ ranges from the slab model X-ray measurements.

When all the measured column densities are considered, the EAGLE predictions agree well with the measured temperatures. EAGLE also clearly predicts that, at these column densities, the gas should consist of two different phases. Without the O VII upper limit, a single-phase solution is still allowed; the constraint on this ‘intermediate’ ion is needed to exclude such a single-phase solution in EAGLE. The strongest single constraint is provided




Figure 5.4: A comparison of the hot and warm phase temperatures predicted from EAGLE to those measured using a two-component CIE fit to the measured X-ray and UV data towards the quasar 3C 273 at $z = 0.09$; this is fig. 10 from [Ahoranta et al. \(2020\)](#). The measurement is indicated in blue, with 1σ error bars. For the warm and hot phase temperatures in EAGLE, we use the O VI and O VIII-weighted temperatures, respectively. The grey histogram and coloured contours indicate the distribution of ion-weighted temperatures found in EAGLE. The contours show different enclosed fractions of EAGLE absorption systems, as indicated in the legend. We only show absorbers with ion column densities ($\log_{10} \text{cm}^{-2}$) in the ranges indicated in the bottom right of each panel. These constraints match the measured column densities and upper limits of [Tilton et al. \(2012\)](#) (O VI) and [Ahoranta et al. \(2020\)](#) (other ions). Specifically, we use the column densities and upper limits that [Ahoranta et al. \(2020\)](#) obtained using a ‘slab’ model, where the column densities of each ion are measured individually, by fitting gaussian line models to each ion’s absorption line(s) independently.

by Ne IX. Considering the relative frequency of absorbers of these four ions at their measured column densities, the Ne IX absorber is also the rarest in EAGLE (Fig. 3.5).

All in all, the $z = 0.09$ X-ray absorber in the 3C 273 quasar spectrum seems to match the EAGLE predictions, considering the observational selection effects on the strength of the X-ray absorption lines. The agreement between the temperatures, given the absorption system column densities, will largely be a consequence of the CIE conditions in high-column-density absorbers, which means they will not be a very sensitive test of the EAGLE model. The comparison between the counterpart column densities is more sensitive, although we note that the correlation between the O VIII and Ne IX column densities will largely be driven by the similar gas in which these ions arise (Fig. 3.1).

Additional analysis of galaxy distributions from SDSS around the sightline to quasar 3C 273 revealed two filaments crossing the sightline around the redshift of the absorption system. Given the uncertainties, both filaments are consistent with the redshift of the X-ray absorbers, but only one is consistent with the O VI redshift. There is also a Milky-Way-like galaxy within the filament at a redshift consistent with the absorption system, at an impact

parameter of 500 kpc.

As the origin of the absorption system, this is consistent with the (large) range of column densities predicted from such a galaxy and impact parameter in the IllustrisTNG simulation (D. Nelson, priv. comm. with Ahoranta et al. 2020). This also seems consistent with EAGLE predictions for O VI (Fig. 3.10). The O VIII and Ne IX column densities are relatively high for those stellar masses and distances, though the O VII upper limit is consistent with expectations. In EAGLE, absorbers with the measured O VIII column density could come from the CGM or filament gas (Fig. 3.6). The O VII upper limit is similarly consistent with both interpretations, but the Ne IX column density would most likely come from the CGM in EAGLE, albeit that of haloes somewhat more massive than that hosting the Milky Way.

5.4.2 System 2: Ton S 180, redshift 0.046

Ahoranta et al. (2021) found the second absorption system we discuss here in the spectrum of the Seyfert galaxy Ton S 180, at $z = 0.046$. For the UV analysis, data from *FUSE* (Far-Ultraviolet Spectroscopic Explorer) and *STIS* (Space Telescope Imaging Spectrograph) on the Hubble Space Telescope were used. The spectra were stacked for analysis, but were also separately fit with a joint model in the same way as the X-ray spectra of §5.4.1 (Ahoranta et al. 2020) to verify the results.

The starting point for this analysis was a known feature in the Ton S 180 spectrum, at $z = 0.0456$ (e.g., Danforth et al. 2006; Tilton et al. 2012). Ahoranta et al. (2021) analysed the FUV spectra and found an O VI absorber at this redshift, with another O VI absorption component, H I (Lyman α and β), and C III nearby.

This redshift was the starting point in the search for an X-ray absorber in the Ton S 180 spectrum. For this system, 200 ks of data were available for both RGS instruments on *XMM-Newton*. Some *Chandra* data were also available, but the exposure time was too short, meaning the constraining power of those observations would be very poor. Therefore, only *XMM-Newton* data were used in this analysis. Like in §5.4.1 the spectra from the different instruments and observing times were processed separately, and fitted jointly. The absorption components were fixed between the models for the different spectra, but the intrinsic spectrum of Ton S 180 was allowed to vary. The X-ray spectra were fitted in the same way as in §5.4.1 using slab models to measure individual column densities from each ion's transitions, and a CIE model to fit the temperature of the gas.

Using the slab model, Ahoranta et al. (2021) found an O VII absorber at 5σ significance. An O VIII absorber is not detected in a statistically significant way, but Ahoranta et al. (2021) did fit its column density. As can be expected, the uncertainties on this column density are large.

In Fig. 5.5, we compare the measured O VII column density to what would be expected from EAGLE, given the column density and temperature of the O VI absorption system. The distribution of EAGLE O VI and O VII column densities is shown in the histogram, while the lines show percentiles of the O VII column density distribution at fixed O VI column density. For these distributions, we split the sample by O VI-weighted temperature, at the 3σ upper limit on the temperature of the FUV absorption system. This limit was calculated from the upper limit on the width of the O VI line in the FUV spectrum, assuming pure thermal broadening.

The measured O VI column density from the FUV spectrum is shown with the vertical band. The red cross shows the measured column density of O VII, and O VI column density inferred from the FUV spectra and the CIE model fit to the X-ray data. The column density




Figure 5.5: The predicted correlation between O VI and O VII column densities from EAGLE. This is fig. 8 from [Ahoranta et al. \(2021\)](#). The column densities are from the $z = 0.1$ EAGLE snapshot, using a slice depth of 6.25 cMpc. The histogram (grey) shows the distribution of all absorbers. The lines show median O VII column densities, and enclosed fractions of O VII absorption systems, at fixed O VI column density. The lines show these percentiles for O VI absorption systems with O VI-weighted temperatures as indicated in the legend; the limit corresponds to the measured 3σ upper limit on this temperature. The lines show estimated detection limits for different instruments. The O VI limit is for *HST-COS*, and the O VII limits are for the *XMM-Newton* RGS, and the future Athena X-IFU. The vertical band (horizontally dashed) shows the measured O VI column density from the FUV spectrum (1σ range), and the red cross shows the measured O VII column density and the total O VI column density, inferred from the FUV spectrum and ionization modelling of the X-ray data, with 1σ error bars.

of O VII was measured using the CIE model at the best-fit redshift of the X-ray absorption system. For the O VI column density, we add the two FUV component column densities from the vertical band to the O VI column density inferred from the CIE model.

At first glance, the large O VI column density in gas at $1.7 \pm 0.2 \times 10^6$ K, inferred from the CIE model X-ray fit, is inconsistent with the measured upper limit of 1.4×10^6 K for the O VI column density in the FUV spectrum. [Ahoranta et al. \(2021\)](#) found that these FUV and X-ray O VI predictions can nonetheless be reconciled. From the X-ray data, no O VI transitions can be measured; the column density for this ion is merely inferred. In the FUV data, a very broad O VI absorber, as predicted by the CIE model, is undetectable, due in part to degeneracies with the continuum model, even with the inferred column density of $10^{13.9 \pm 0.2} \text{ cm}^{-2}$.

In EAGLE, absorption systems with at least the O VI column density measured from the FUV data, and with a temperature below the measured 3σ upper limit, usually have lower O VII column densities than the one measured here. This is true even considering that only the strongest ≈ 10 – 20 per cent of X-ray counterparts were found. However, the O VI inferred to be present in this absorption system includes a hot component as well as the FUV-measured warm phase. In the EAGLE predictions, those phases would both contribute to the ion-weighted temperature. Therefore, the observed system could end up with a temperature




Figure 5.6: A comparison of the galaxy environment of the measured absorption system to predictions from EAGLE. The left panel is fig. 9 from [Ahoranta et al. \(2021\)](#). We search for the nearest neighbour galaxies in the same slice of the EAGLE simulation as absorbers with column densities like the measured ones. For the O VII and O VIII column densities, we use the 1σ ranges of slab column densities, measured at the redshift of the O VI absorber. For O VI, we require a column density at least as large as the best-fit (total) column density measured in the FUV spectrum. We also allow higher column densities because the CIE model predicts a broad O VI absorber in the same gas phase as the X-ray lines, with similar column density to the FUV system. We search for galaxies with stellar masses $M_{\star} \geq 10^9 M_{\odot}$, which is roughly the detection limit in the survey. The left panel shows the distances (impact parameters) of these galaxies to the absorbers. Vertical lines show the measurements, the curves show the cumulative distributions from EAGLE. The neighbours are numbered from closest (nb. 1) to third closest (nb. 3). The measured absorber is relatively far from galaxies compared to EAGLE, and especially far from the closest galaxy. The right panel shows the distribution of stellar masses, as the fraction of absorbers per neighbour stellar mass bin ($\Delta \log_{10} M_{\star} = 0.1$). The measured masses and 1σ errors are shown by the vertical lines and bands. This seems consistent between EAGLE and the measured system, but as the neighbour stellar mass distribution is quite flat, this is not a very informative comparison.

above the warm-phase upper limit in the EAGLE data, depending on the actual temperature of the warm phase. The measurements are consistent with an O VI temperature determined by the hot phase rather than the warm phase.

Due to these complications, it is difficult to account for the prior information on the temperature of FUV O VI absorber in the EAGLE analysis. Because the majority of the O VI systems in EAGLE are below the temperature limit, ignoring temperature constraints would yield predictions very similar to those for $T < 10^{6.1}$ K O VI absorption systems. Therefore, the column density of the O VII absorber is higher than than expected from the column density of the FUV O VI absorber, even considering the non-detections of X-ray counterparts to ≈ 10 –20 per cent of the O VI (UV) absorbers. However, given the temperatures inferred for the two phases, the combination of inferred total O VI and O VII column densities is plausible.

We note that the O VII slab model column densities, inferred from only the O VII transitions, are higher than the CIE model values, but the 1σ lower limits are similar. The level of (in)consistency between the measured data and EAGLE therefore does not depend too strongly on how the O VII column density is measured.

To investigate the source of the absorption, [Ahoranta et al. \(2021\)](#) searched for galaxies in

a 9×9 pMpc region around the absorber on the sky, with a depth of 4 pMpc. They estimated the stellar masses of the galaxies based on their B, R, and K-band photometry, and their halo masses from a stellar-mass-halo-mass relation. The absorber was found to be outside their virial radii, though the closest three are within $\approx 2-3 R_{200c}$.

We compare the distances and masses of these galaxies to those of the closest neighbour galaxies to EAGLE absorbers with column densities similar to those measured in the Ton S 180 sightline. We match galaxies to absorbers by considering galaxies with the centres of mass within the same slice as the absorber. We do not include any peculiar velocity adjustments like those in §5.3.2. We used 6.25 cMpc slices, which is smaller than the line of sight search region used in the observations. However, the observed galaxies are the closest even in a wider range of relative velocities (Ahoranta et al. 2021, fig. 9). Additionally, including galaxies up to 3.125 cMpc from the slice edges ($\pm \Delta z = 6.25$ cMpc) in the neighbour search makes little difference for the level of agreement between EAGLE and the observations. This means the larger search region in EAGLE should not matter very much for the comparison.

When comparing EAGLE O VI column densities to the observations, we consider the column density measured in the FUV spectrum to be a lower limit. This is because the CIE model fitted to the X-ray data predicts an O VI absorber produced by the same warm-hot gas producing the X-ray lines, with a column density similar to that of the FUV absorber. We note that the EAGLE column densities measure all O VI, and are insensitive to the shapes of any absorption lines. For O VII and O VIII, we use the slab model column densities, measured at the redshift of the FUV O VI absorber, with their 1σ error ranges. We only consider EAGLE galaxies with stellar masses $\geq 10^9 M_\odot$ in the neighbour search, matching the estimated limit of the survey used in the observations.

We find the nearest neighbours in EAGLE, and compare the distances (impact parameters) between these neighbours and the absorber, and the neighbour’s stellar masses, to the measured values in Fig. 5.6. EAGLE predicts absorbers closer to galaxies than found here, though only the nearest neighbour distance represents a clear tension. The neighbour stellar mass distributions are fairly flat up to $\approx 10^{11} M_\odot$, so ‘extreme’ values in the distribution are not really less likely than more ‘central’ values, and we plot the differential distribution rather than the cumulative distribution. The stellar masses of the neighbours are consistent with the measured values. However, the broad distribution of neighbour stellar masses in EAGLE, combined with the single absorber we compare to, mean that this comparison is mostly a sanity check.

Ahoranta et al. (2021) also made rough estimates of the metallicity of the hot gas, using upper limits on the H I column density associated with this phase. The Lyman α line arising from the hot phase would be blended with two nearby Lyman α lines, and parameters such as the line width are uncertain, so these are somewhat rough estimates.

The estimated 3σ (5σ) lower limit on the metallicity is $0.8 Z_\odot$ ($0.5 Z_\odot$). This would be higher than expected for typical IGM gas (e.g., Martizzi et al. 2019), though consistent with O VI–VIII-weighted metallicities up to at least $\sim 3 R_{200c}$ from galaxies with a wide range of halo masses (e.g., Chapter 3, Wijers et al. 2020). The high metallicity may therefore indicate a particular connection to a nearby galaxy, such as a far-reaching outflow, but it may also be a common value for absorbers discovered through metal line absorption.

In short, like the previous absorption system, this one has higher X-ray column densities than would be expected based on the column density of the O VI absorber. Unlike for the 3C 273 absorber, this seems true even considering that $\approx 80-90$ per cent of O VI (UV) absorbers had no observed counterparts. However, given 11 (23) independent trials, there is a probability of $1 - 0.99^{11} = 0.1$ (0.2) of finding at least one absorber stronger than the

99th percentile of the X-ray column density distribution, given the O VI column density. This 99th percentile is roughly the 1σ lower limit for the discrepancy with the EAGLE predictions (Fig. 5.5). Therefore, the high O VII column density may not represent a strong tension between the observations and the predictions of EAGLE.

The difference in predicted neighbour distances is less likely to be a selection effect, as we constrained the column densities of all the measured ions to create the sample of EAGLE absorbers, and matched the expected minimum stellar mass for the observed galaxies. The measured absorber is further from its closest neighbours than EAGLE would predict, though only the closest neighbour distance is unlikely enough in EAGLE that it may represent a tension. Ultimately, it is difficult to draw strong conclusions about the validity of the EAGLE model from a single absorption system.

5.5 Discussion

From two of the claimed absorption systems, it seems that EAGLE may under-predict the number of strong O VII absorbers well outside R_{200c} of galaxies. We stress that this is not a certain conclusion, given that we only have two absorption systems for which this seems to be the case, and the detection of one of those is particularly uncertain. However, if EAGLE indeed under-predicts the occurrence of such absorbers, there could be a number of issues at play.

If the absorption systems probe the outflows of galaxies detected in the surveys, the EAGLE outflows may not carry enough metals to distances significantly larger than R_{200c} , or the outflows may not be at the right temperature to produce O VII. Alternatively, these absorption systems may probe the CGM or outflows from smaller galaxies, undetected in the surveys and missing in EAGLE due to its limited resolution.

A lack of metals at sufficient densities outside the CGM is, however, somewhat in tension with the reasonable agreement of EAGLE UV metal ion column density distributions (Schaye et al. 2015; Rahmati et al. 2016) to observations. The EAGLE column density distributions represent lower-density structures as well as the CGM (Rahmati et al. 2016; Wijers et al. 2020). In addition, Johnson et al. (2019) searched the 1ES1553+113 FUV spectrum for O VI and H I absorbers, and searched for galaxies in the vicinity of these UV absorbers. Their analysis found that metals are not typically found as far from galaxies as the $z = 0.355$ system claimed by Nicastro et al. (2018). Therefore, FUV observations point to a sufficient amount of dense, metal-enriched gas outside the CGM in EAGLE. However, a dearth of hot, dense, metal-enriched gas cannot be excluded based on these UV ion measurements, and the spatial distribution of these metal absorbers may not be correctly predicted by EAGLE.

The number the X-ray absorber(s) found by Nicastro et al. (2018) is consistent with EAGLE predictions. However, the large uncertainties in the inferred absorber density along the line of sight, mostly due to the small number of absorbers, mean that this is more of a sanity check than a stringent test of the EAGLE model.

The X-ray absorbers found as counterparts to O VI absorbers have higher column densities than is typical for such counterparts in EAGLE. However, the relatively high column densities are reasonably consistent with the EAGLE predictions if we also account for the non-detections: the probability of at least one absorber as discrepant as the Ton S 180 absorber is ≈ 0.1 – 0.2 . The probability that at least two of the counterparts have column densities above the 90th (95th) percentiles of the probability distribution for a single UV absorber are 0.3–0.7 (0.1–0.3). The lower (upper) ends of the ranges we give are for 11 (23) trials, reflecting the number of redshifts of O VI absorbers (O VI and broad H I absorbers) searched

for X-ray absorption.

5.6 Summary

Overall, EAGLE seems to hold up reasonably to some of the scarce data available on soft X-ray absorption. There is some indication that the strongest absorbers are clustered too close to galaxies in EAGLE, but this is based on two observations, so it is not a strong tension. These studies and comparisons do illustrate the sorts of information that will become available when instruments with higher sensitivities and spectral resolution increase the number of absorption systems we can measure, and the information we can get for each system. For example, measuring more systems and more ions in each system, will mean we can better measure the temperature, spatial distribution, and mass of the warm-hot CGM and the WHIM.

The measurements we have also make clear that we can learn much more from X-ray absorption spectra when we combine them with other data. For example, the UV spectra provide information on the phase structure of the gas (through absorption from various lower ionization states) and the metallicity of the gas (through H I absorption). Galaxy redshift surveys can tell us whether any absorber represents CGM or filament gas, and more information on those galaxies will be very helpful in understanding the baryon cycle.

Bibliography

- Abazajian K. N., et al., 2016, arXiv e-prints, [p. arXiv:1610.02743](#)
- Aguirre A., Hernquist L., Schaye J., Katz N., Weinberg D. H., Gardner J., 2001, [ApJ](#), [561](#), [521](#)
- Ahoranta J., et al., 2020, [A&A](#), [634](#), [A106](#)
- Ahoranta J., Finoguenov A., Bonamente M., Tilton E., Wijers N., Muzahid S., Schaye J., 2021, submitted to A&A
- Allende Prieto C., Lambert D. L., Asplund M., 2001, [ApJ](#), [556](#), [L63](#)
- Allende Prieto C., Lambert D. L., Asplund M., 2002, [ApJ](#), [573](#), [L137](#)
- Anderson M. E., Bregman J. N., Dai X., 2013, [ApJ](#), [762](#), [106](#)
- Anderson M. E., Churazov E., Bregman J. N., 2016, [MNRAS](#), [455](#), [227](#)
- Astropy Collaboration et al., 2013, [A&A](#), [558](#), [A33](#)
- Astropy Collaboration et al., 2018, [aj](#), [156](#), [123](#)
- Bahcall J. N., Peebles P. J. E., 1969, [ApJ](#), [156](#), [L7](#)
- Barnes D. J., et al., 2017, [MNRAS](#), [471](#), [1088](#)
- Barret D., et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M., eds, Proc. SPIE Vol. 9905, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 99052F ([arXiv:1608.08105](#)), [doi:10.1117/12.2232432](#)
- Barret D., et al., 2018, in Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray. p. 106991G ([arXiv:1807.06092](#)), [doi:10.1117/12.2312409](#)
- Bernitt S., et al., 2012, [Nature](#), [492](#), [225](#)
- Bertone S., Schaye J., Dalla Vecchia C., Booth C. M., Theuns T., Wiersma R. P. C., 2010a, [MNRAS](#), [407](#), [544](#)
- Bertone S., Schaye J., Booth C. M., Dalla Vecchia C., Theuns T., Wiersma R. P. C., 2010b, [MNRAS](#), [408](#), [1120](#)
- Bertone S., Aguirre A., Schaye J., 2013, [MNRAS](#), [430](#), [3292](#)
- Biffi V., Planelles S., Borgani S., Rasia E., Murante G., Fabjan D., Gaspari M., 2018, [MNRAS](#), [476](#), [2689](#)

- Bogdán Á., et al., 2015, [ApJ](#), 804, 72
- Bonamente M., 2019, [Journal of Applied Statistics](#), 46, 1129
- Bonamente M., Nevalainen J., Tilton E., Liivamägi J., Tempel E., Heinämäki P., Fang T., 2016, [MNRAS](#), 457, 4236
- Booth C. M., Schaye J., 2009, [MNRAS](#), 398, 53
- Booth C. M., Schaye J., Delgado J. D., Dalla Vecchia C., 2012, [MNRAS](#), 420, 1053
- Borgani S., et al., 2004, [MNRAS](#), 348, 1078
- Bower R. G., Schaye J., Frenk C. S., Theuns T., Schaller M., Crain R. A., McAlpine S., 2017, [MNRAS](#), 465, 32
- Branchini E., et al., 2009, [ApJ](#), 697, 328
- Bregman J. N., 2007, [Annual Review of Astronomy and Astrophysics](#), 45, 221
- Bregman J. N., Lloyd-Davies E. J., 2007, [ApJ](#), 669, 990
- Brenneman L. W., et al., 2016, in Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray. Society of Photo-Optical Instrumentation Engineers (SPIE), Bellingham, Washington USA, p. 99054P, [doi:10.1117/12.2231193](https://doi.org/10.1117/12.2231193)
- Burchett J. N., et al., 2019, [ApJ](#), 877, L20
- Cagnoni I., Nicastro F., Maraschi L., Treves A., Tavecchio F., 2004, [ApJ](#), 603, 449
- Carswell B., Schaye J., Kim T.-S., 2002, *The Astrophysical Journal*, 578, 43
- Cen R., 2012, [ApJ](#), 753, 17
- Cen R., Fang T., 2006, [ApJ](#), 650, 573
- Cen R., Ostriker J. P., 1999, [ApJ](#), 514, 1
- Cen R., Ostriker J. P., 2006, [ApJ](#), 650, 560
- Chabrier G., 2003, [ApJ](#), 586, L133
- Chandra X-ray Center Chandra Project Science at MSFC Chandra IPI Teams 2020, The Chandra Proposers' Observatory Guide, <https://cxc.harvard.edu/proposer/POG/html/index.html>
- Chen X., Weinberg D. H., Katz N., Davé R., 2003, [ApJ](#), 594, 42
- Cole S., 1991, [ApJ](#), 367, 45
- Collette A., 2013, Python and HDF5. O'Reilly, Sebastopol CA, USA, <http://www.h5py.org/>
- Connor L., Ravi V., 2021, arXiv e-prints, p. [arXiv:2107.13692](https://arxiv.org/abs/2107.13692)
- Correa C. A., Schaye J., Wyithe J. S. B., Duffy A. R., Theuns T., Crain R. A., Bower R. G., 2018, [MNRAS](#), 473, 538

- Crain R. A., McCarthy I. G., Schaye J., Theuns T., Frenk C. S., 2013, [MNRAS](#), [432](#), 3005
- Crain R. A., et al., 2015, [MNRAS](#), [450](#), 1937
- Cucchetti E., et al., 2018, in Den Herder J.-W. A., Nikzad S., Nakazawa K., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 10699, Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray. p. 106994L, [doi:10.1117/12.2311957](#)
- Cui W., et al., 2020a, [Journal of Low Temperature Physics](#), [199](#), 502
- Cui W., et al., 2020b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 114442S ([arXiv: 2101.05587](#)), [doi:10.1117/12.2560871](#)
- Cyburt R. H., Fields B. D., Olive K. A., Yeh T.-H., 2016, [Rev. Mod. Phys.](#) [88](#), 015004
- Dalla Vecchia C., Schaye J., 2012, [MNRAS](#), [426](#), 140
- Danforth C. W., Shull J. M., Rosenberg J. L., Stocke J. T., 2006, [ApJ](#), [640](#), 716
- Das S., Mathur S., Gupta A., Nicastro F., Krongold Y., 2019, [ApJ](#), [887](#), 257
- Das S., Mathur S., Gupta A., 2020, [ApJ](#), [897](#), 63
- Davies J. J., Crain R. A., McCarthy I. G., Oppenheimer B. D., Schaye J., Schaller M., McAlpine S., 2019, [MNRAS](#)
- Davies J. J., Crain R. A., Oppenheimer B. D., Schaye J., 2020, Monthly Notices of the Royal Astronomical Society, [491](#), 4462
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, [ApJ](#), [292](#), 371
- Dekel A., Birnboim Y., 2006, Monthly notices of the royal astronomical society, [368](#), 2
- Del Zanna G., Dere K. P., Young P. R., Landi E., 2021, [ApJ](#), [909](#), 38
- Dere K. P., Landi E., Mason H. E., Monsignori Fossi B. C., Young P. R., 1997, [A&AS](#), [125](#), 149
- Doe S., et al., 2007, in Shaw R. A., Hill F., Bell D. J., eds, Astronomical Society of the Pacific Conference Series Vol. 376, Astronomical Data Analysis Software and Systems XVI. p. 543
- Dolag K., Borgani S., Murante G., Springel V., 2009, [MNRAS](#), [399](#), 497
- Dubois Y., et al., 2014, [MNRAS](#), [444](#), 1453
- Fabian A. C., 2012, [ARA&A](#), [50](#), 455
- Faerman Y., Sternberg A., McKee C. F., 2017, [ApJ](#), [835](#), 52
- Faerman Y., Sternberg A., McKee C. F., 2020, [ApJ](#), [893](#), 82
- Fang T., Bryan G. L., Canizares C. R., 2002a, [ApJ](#), [564](#), 604
- Fang T., Marshall H. L., Lee J. C., Davis D. S., Canizares C. R., 2002b, [ApJ](#), [572](#), L127
- Faucher-Giguère C.-A., 2020, [MNRAS](#), [493](#), 1614

- Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, [Publications of the Astronomical Society of the Pacific](#), **110**, 761
- Ferland G. J., et al., 2013, *Rev. Mexicana Astron. Astrofis.*, **49**, 137
- Ferland G. J., et al., 2017, *Rev. Mexicana Astron. Astrofis.*, **53**, 385
- Freeman P., Doe S., Siemiginowska A., 2001, in Starck J.-L., Murtagh F. D., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 4477, Astronomical Data Analysis. pp 76–87 ([arXiv:astro-ph/0108426](#)), [doi:10.1117/12.447161](#)
- Furlanetto S. R., Phillips L. A., Kamionkowski M., 2005, *MNRAS*, **359**, 295
- Gatuzz E., Churazov E., 2018, *MNRAS*, **474**, 696
- Gordon H., Hobby M. G., Peacock N. J., 1980, *Journal of Physics B: Atomic and Molecular Physics*, **13**, 1985
- Gu M. F., Chen H., Brown G. V., Beiersdorfer P., Kahn S. M., 2007, *ApJ*, **670**, 1504
- Gu L., et al., 2019, *A&A*, **627**, A51
- Guainazzi M., Tashiro M. S., 2018, arXiv e-prints, [p. arXiv:1807.06903](#)
- Gupta A., Mathur S., Krongold Y., Nicastro F., Galeazzi M., 2012, *ApJ*, **756**, L8
- Gupta A., Mathur S., Galeazzi M., Krongold Y., 2014, *AP&SS*, **352**, 775
- Gupta A., Mathur S., Krongold Y., 2017, *ApJ*, **836**, 243
- Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays. Commissariat à l'énergie atomique (CEA), France ([arXiv:astro-ph/0106018](#))
- Haardt F., Madau P., 2012, *ApJ*, **746**, 125
- Hellsten U., Gnedin N. Y., Miralda-Escudé J., 1998, *ApJ*, **509**, 56
- Hickox R. C., Markevitch M., 2007, *ApJ*, **661**, L117
- Hitomi Collaboration et al., 2018, [Publications of the Astronomical Society of Japan](#), **70**
- Hodges-Kluck E. J., Miller M. J., Bregman J. N., 2016, *ApJ*, **822**, 21
- Holweger H., 2001, in Wimmer-Schweingruber R. F., ed., American Institute of Physics Conference Series Vol. 598, Joint SOHO/ACE workshop “Solar and Galactic Composition”. pp 23–30 ([arXiv:astro-ph/0107426](#)), [doi:10.1063/1.1433974](#)
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, **445**, 581
- Hopkins P. F., et al., 2018, *MNRAS*, **480**, 800
- Hunter J. D., 2007, [Computing in Science & Engineering](#), **9**, 90
- Hutcheon R. J., Pye J. P., Evans K. D., 1976, *MNRAS*, **175**, 489

- Johnson S. D., Chen H.-W., Mulchaey J. S., 2013, [MNRAS](#), [434](#), 1765
- Johnson S. D., Chen H.-W., Mulchaey J. S., 2015, [MNRAS](#), [449](#), 3263
- Johnson S. D., Chen H.-W., Mulchaey J. S., Schaye J., Straka L. A., 2017, [ApJ](#), [850](#), L10
- Johnson S. D., et al., 2019, [ApJ](#), [884](#), L31
- Jones E., Oliphant T., Peterson P., et al., 2001, SciPy: Open source scientific tools for Python, <http://www.scipy.org/>
- Kaastra J., 2018, private communication
- Kaastra J. S., Werner N., Herder J. W. A. d., Paerels F. B. S., de Plaa J., Rasmussen A. P., de Vries C. P., 2006, [ApJ](#), [652](#), 189
- Kereš D., Katz N., Fardal M., Davé R., Weinberg D. H., 2009, Monthly Notices of the Royal Astronomical Society, 395, 160
- Khaire V., Srianand R., 2019, [MNRAS](#), [484](#), 4174
- Kovács O. E., Bogdán Á., Smith R. a. K., Kraft R. P., Forman W. R., 2019, [ApJ](#), [872](#), 83
- Kuntz K. D., Snowden S. L., 2000, [ApJ](#), [543](#), 195
- Landi E., Young P. R., Dere K. P., Del Zanna G., Mason H. E., 2013, [ApJ](#), [763](#), 86
- Li Y.-S., White S. D. M., 2008, [MNRAS](#), [384](#), 1459
- Li Z.-Z., Qian Y.-Z., Han J., Li T. S., Wang W., Jing Y. P., 2020, [ApJ](#), [894](#), 10
- Lide D. R., ed. 2003, CRC Handbook of Chemistry and Physics, 84 edn. CRC Press LLC, Boca Raton
- Lim S. H., Mo H. J., Li R., Liu Y., Ma Y.-Z., Wang H., Yang X., 2018, [The Astrophysical Journal](#), [854](#), 181
- Lim S. H., Mo H. J., Wang H., Yang X., 2020, [The Astrophysical Journal](#), [889](#), 48
- Lotti S., Perinati E., Natalucci L., Piro L., Mineo T., Colasanti L., Macculi C., 2012, [Nuclear Instruments and Methods in Physics Research A](#), [686](#), 31
- Lotti S., et al., 2014, [Astronomy & Astrophysics](#), 569
- Luisi M., et al., 2021, [Science Advances](#), 7
- Lumb D., den Herder J.-W., the Athena Science Team 2017, Issue/Revision 2.01, Athena Science Requirements Document. European Space Agency, European Space Research and Technology Centre, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
- Macquart J. P., et al., 2020, [Nature](#), [581](#), 391
- Mandelker N., Nagai D., Aung H., Dekel A., Birnboim Y., van den Bosch F. C., 2020, [MNRAS](#), [494](#), 2641
- Martizzi D., Hahn O., Wu H.-Y., Evrard A. E., Teyssier R., Wechsler R. H., 2016, [MNRAS](#), [459](#), 4408

- Martizzi D., et al., 2019, [MNRAS](#), [486](#), 3766
- Mathur S., Weinberg D. H., Chen X., 2003, [ApJ](#), [582](#), 82
- McAlpine S., et al., 2016, [Astronomy and Computing](#), [15](#), 72
- McCammon D., et al., 2002, [ApJ](#), [576](#), 188
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, [MNRAS](#), [465](#), 2936
- McQuinn M., 2014, [ApJ](#), [780](#), L33
- Meiring J. D., Tripp T. M., Werk J. K., Howk J. C., Jenkins E. B., Prochaska J. X., Lehner N., Sembach K. R., 2013, [ApJ](#), [767](#), 49
- Mernier F., et al., 2018, [Space Sci. Rev.](#), [214](#), 129
- Mernier F., et al., 2020, [A&A](#), [642](#), A90
- Mewe R., Gronenschild E. H. B. M., 1981, [A&AS](#), [45](#), 11
- Miller M. J., Bregman J. N., 2015, [ApJ](#), [800](#), 14
- Mitchell P. D., Schaye J., 2021, arXiv e-prints, p. [arXiv:2103.10966](#)
- Mitchell P. D., Schaye J., Bower R. G., Crain R. A., 2020, [MNRAS](#), [494](#), 3971
- Monaghan J. J., Lattanzio J. C., 1985, [A&A](#), [149](#), 135
- Monari G., et al., 2018, [A&A](#), [616](#), L9
- Morrison R., McCammon D., 1983, [ApJ](#), [270](#), 119
- Morton D. C., 2003, [The Astrophysical Journal Supplement Series](#), [149](#), 205
- Mroczkowski T., et al., 2019, [Space Sci. Rev.](#), [215](#), 17
- Nelson D., et al., 2018, [MNRAS](#), [477](#), 450
- Nelson D., et al., 2019, [MNRAS](#), [490](#), 3234
- Nicastro F., 2018, preprint, ([arXiv:1811.03498](#))
- Nicastro F., et al., 2002, [ApJ](#), [573](#), 157
- Nicastro F., et al., 2005, [ApJ](#), [629](#), 700
- Nicastro F., Krongold Y., Mathur S., Elvis M., 2017, [Astronomische Nachrichten](#), [338](#), 281
- Nicastro F., et al., 2018, [Nature](#), [558](#), 406
- Oliphant T. E., 2006, Guide to NumPy. Trelgol Publishing, USA
- Oppenheimer B. D., 2018, [MNRAS](#), [480](#), 2963
- Oppenheimer B. D., Schaye J., 2013, [MNRAS](#), [434](#), 1063
- Oppenheimer B. D., et al., 2016, [MNRAS](#), [460](#), 2157

- Oppenheimer B. D., Segers M., Schaye J., Richings A. J., Crain R. A., 2018, [MNRAS](#), [474](#), 4740
- Oppenheimer B. D., et al., 2020a, [MNRAS](#), [491](#), 2939
- Oppenheimer B. D., et al., 2020b, [ApJ](#), [893](#), L24
- Pérez F., Granger B. E., 2007, [Computing in Science & Engineering](#), [9](#), 21
- Perna R., Loeb A., 1998, [ApJ](#), [503](#), L135
- Pierre M., Bryan G., Gastaud R., 2000, [A&A](#), [356](#), 403
- Pillepich A., et al., 2018, [MNRAS](#), [473](#), 4077
- Planck Collaboration et al., 2014, [A&A](#), [571](#), A1
- Planck Collaboration et al., 2020, [A&A](#), [641](#), A6
- Planelles S., Schleicher D. R. G., Bykov A. M., 2015, [Space Sci. Rev.](#), [188](#), 93
- Ploeckinger S., Schaye J., 2020, [MNRAS](#), [497](#), 4857
- Porquet D., Dubau J., 2000, [A&AS](#), [143](#), 495
- Prochaska J. X., Zheng Y., 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 648
- Prochaska J. X., Weiner B., Chen H. W., Mulchaey J., Cooksey K., 2011, [ApJ](#), [740](#), 91
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, [MNRAS](#), [430](#), 2427
- Rahmati A., Schaye J., Crain R. A., Oppenheimer B. D., Schaller M., Theuns T., 2016, [MNRAS](#), [459](#), 310
- Rasmussen A., Kahn S. M., Paerels F., 2003, *X-ray IGM in the Local Group*. Kluwer Academic Publishers, Dordrecht, p. 109, [doi:10.1007/978-94-010-0115-1_20](https://doi.org/10.1007/978-94-010-0115-1_20)
- Ravera L., et al., 2014, in Takahashi T., den Herder J.-W. A., Bautz M., eds, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9144, Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*. p. 91442L, [doi:10.1117/12.2055884](https://doi.org/10.1117/12.2055884)
- Ravi V., 2019, [ApJ](#), [872](#), 88
- Roca-Fàbrega S., et al., 2019, [MNRAS](#), [484](#), 3625
- Rohatgi A., 2018, *WebPlotDigitizer*, <https://automeris.io/WebPlotDigitizer>
- Rosas-Guevara Y. M., et al., 2015, [MNRAS](#), [454](#), 1038
- Rupke D., 2018, [Galaxies](#), [6](#), 138
- Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, [MNRAS](#), [454](#), 2277
- Schaye J., 2004, [ApJ](#), [609](#), 667
- Schaye J., Dalla Vecchia C., 2008, [MNRAS](#), [383](#), 1210
- Schaye J., et al., 2010a, [MNRAS](#), [402](#), 1536

- Schaye J., et al., 2010b, [Monthly Notices of the Royal Astronomical Society](#), 402, 1536
- Schaye J., et al., 2015, [MNRAS](#), 446, 521
- Segers M. C., Oppenheimer B. D., Schaye J., Richings A. J., 2017, [MNRAS](#), 471, 1026
- Shorer P., 1979, [Phys. Rev. A](#), 20, 642
- Shull J. M., Smith B. D., Danforth C. W., 2012, [ApJ](#), 759, 23
- Simionescu A., et al., 2013, [ApJ](#), 775, 4
- Simionescu A., et al., 2021, [Experimental Astronomy](#),
- Smith R. K., et al., 2016, in *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. Society of Photo-Optical Instrumentation Engineers (SPIE), Bellingham, Washington USA, p. 99054M, [doi:10.1117/12.2231778](#)
- Springel V., 2005, [MNRAS](#), 364, 1105
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, [MNRAS](#), 328, 726
- Springel V., et al., 2005, [Nature](#), 435, 629
- Stern J., Fielding D., Faucher-Giguère C.-A., Quataert E., 2019, [MNRAS](#), 488, 2549
- Takei Y., et al., 2011, [ApJ](#), 734, 91
- Tanimura H., et al., 2019, [MNRAS](#), 483, 223
- Tepper-García T., Richter P., Schaye J., Booth C. M., Dalla Vecchia C., Theuns T., Wiersma R. P. C., 2011, [MNRAS](#), 413, 190
- Tepper-García T., Richter P., Schaye J., Booth C. M., Dalla Vecchia C., Theuns T., 2012, [MNRAS](#), 425, 1640
- Tepper-García T., Richter P., Schaye J., 2013, [MNRAS](#), 436, 2063
- The CHIME/FRB Collaboration et al., 2021, arXiv e-prints, [p. arXiv:2106.04352](#)
- The EAGLE team 2017, arXiv e-prints, [p. arXiv:1706.09899](#)
- The Lynx Team 2018, arXiv e-prints,
- Tilton E. M., Danforth C. W., Shull J. M., Ross T. L., 2012, [ApJ](#), 759, 112
- Tripp T. M., Sembach K. R., Bowen D. V., Savage B. D., Jenkins E. B., Lehner N., Richter P., 2008, [ApJS](#), 177, 39
- Truong N., et al., 2018, [MNRAS](#), 474, 4089
- Truong N., et al., 2020, [MNRAS](#), 494, 549
- Tumlinson J., et al., 2011, [Science](#), 334, 948
- Tumlinson J., Peebles M. S., Werk J. K., 2017, [ARA&A](#), 55, 389
- Tuominen T., et al., 2021, [A&A](#), 646, A156

- Van de Voort F., Schaye J., 2013, [MNRAS](#), [430](#), 2688
- Van de Voort F., Schaye J., Booth C., Haas M. R., Dalla Vecchia C., 2011, *Monthly Notices of the Royal Astronomical Society*, 414, 2458
- Verner D. A., Verner E. M., Ferland G. J., 1996, [Atomic Data and Nuclear Data Tables](#), [64](#), 1
- Vogelsberger M., et al., 2014, [MNRAS](#), [444](#), 1518
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, [Nature Reviews Physics](#), [2](#), 42
- Voit G. M., 2019, [ApJ](#), [880](#), 139
- Walsh S., McBreen S., Martin-Carrillo A., Dauser T., Wijers N., Wilms J., Schaye J., Barret D., 2020, [A&A](#), [642](#), A24
- Walters A., Ma Y.-Z., Sievers J., Weltman A., 2019, [Phys. Rev. D](#), [100](#), 103519
- Wendland H., 1995, [Advances in Computational Mathematics](#), 4, 389
- Werk J. K., et al., 2014, [ApJ](#), [792](#), 8
- Werk J. K., et al., 2016, [ApJ](#), [833](#), 54
- Werner N., Mernier F., 2020, *Hot Atmospheres of Galaxies, Groups, and Clusters of Galaxies*. Springer International Publishing, pp 279–310, [doi:10.1007/978-3-030-38509-5_10](#)
- Wiersma R. P. C., Schaye J., Smith B. D., 2009a, [MNRAS](#), [393](#), 99
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, [MNRAS](#), [399](#), 574
- Wijers N. A., Schaye J., Oppenheimer B. D., Crain R. A., Nicastro F., 2019, [MNRAS](#), [488](#), 2947
- Wijers N. A., Schaye J., Oppenheimer B. D., 2020, [MNRAS](#), [498](#), 574
- Wright R. J., Lagos C. d. P., Power C., Mitchell P. D., 2020, [MNRAS](#), [498](#), 1668
- Wu C., Gao X., 2019, [Scientific Reports](#), [9](#), 7463
- XRISM Science Team 2020, arXiv e-prints, [p. arXiv:2003.04962](#)
- Yoshikawa K., Sasaki S., 2006, [Publications of the Astronomical Society of Japan](#), [58](#), 641
- Yoshikawa K., Yamasaki N. Y., Suto Y., Ohashi T., Mitsuda K., Tawara Y., Furuzawa A., 2003, [PASJ](#), [55](#), 879
- Zaritsky D., 1999, in Gibson B. K., Axelrod R. S., Putman M. E., eds, *Astronomical Society of the Pacific Conference Series Vol. 165, The Third Stromlo Symposium: The Galactic Halo*. p. 34 ([arXiv:astro-ph/9810069](#))
- Zhang Y. N., et al., 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. p. 114449F, [doi:10.1117/12.2560964](#)
- Zhuravleva I., et al., 2014, [Nature](#), [515](#), 85
- Zinger E., et al., 2020, [MNRAS](#), [499](#), 768

de Graaff A., Cai Y.-C., Heymans C., Peacock J. A., 2019, [A&A](#), [624](#), A48

de Plaa J., Zhuravleva I., Werner N., Kaastra J. S., Churazov E., Smith R. K., Raassen A. J. J., Grange Y. G., 2012, [A&A](#), [539](#), A34

den Herder J. W., et al., 2001, [A&A](#), [365](#), L7


Nederlandse samenvatting

Van gas naar sterrenstelsels

In dit proefschrift heb ik onderzocht hoe we het hete gas om en tussen sterrenstelsels kunnen vinden. Dat heb ik gedaan met behulp van computersimulaties: de EAGLE-simulaties. EAGLE staat voor 'Evolution and Assembly of GaLaxies and their Environments': evolutie en opbouw van sterrenstelsels en hun omgevingen. In die simulaties beginnen we met 3 miljard 'stukjes' donkere materie en evenveel 'stukjes' gas, in een groot stuk (nagebootst) heelal. Dit stuk is groot genoeg dat er honderden sterrenstelsels in vormen die ongeveer even zwaar zijn als onze Melkweg. Donkere materie is een stof die overal gewoon doorheen beweegt, inclusief andere donkere materie. Het oefent wel zwaartekracht uit, op andere donkere materie en op 'gewone' materie zoals gas, sterren, planeten en proefschriften. We hebben donkere materie gevonden door te ontdekken dat er iets in het heelal is dat zwaartekracht uitoefent op gas en sterren. Dat 'iets' kunnen we niet zien en we weten ook niet precies wat het is. Daarom wordt het donkere materie genoemd.

In de EAGLE-simulatie delen we die donkere materie en het gas op in stukjes, omdat je in computer alleen een beperkte lijst temperaturen, dichtheden, posities, etc. kunt berekenen en opslaan. Deze stukjes zijn een beetje zoals de pixels in een digitale foto: met genoeg kleine stukjes kun je goed benaderen wat je in werkelijkheid ziet. De grootte van de pixels bepaalt hoeveel details je kunt zien. De 'stukjes' gas in de simulatie stellen bollen gas voor, die bewegen, uitzetten of worden samengedrukt door de zwaartekracht en de druk van het gas eromheen. Elk stukje gas heeft een positie, snelheid, temperatuur, dichtheid etc. In werkelijkheid kan er binnen het gebied dat één zo'n stukje gas beslaat, gas bestaan van verschillende temperaturen en dichtheden. In de simulatie hebben we echter alleen één temperatuur en dichtheid van het stukje gas, zoals één enkele groene pixel in een foto soms alle verschillende bladeren van een boom voorstelt. Hoewel gas ook in werkelijkheid uit deeltjes bestaat, zijn deze heel veel kleiner dan de stukjes in de simulatie. Elk stuk gas heeft een massa 2 miljoen keer zo groot als die van de zon. De massa aan sterren in de Melkweg is ongeveer 20.000 keer zo groot als zo'n stuk gas. Als sterren vormen uit dat gas, volgt de simulatie ook niet individuele sterren, maar grote groepen sterren die tegelijk vormen en samen bewegen.

In dit soort simulaties zijn er dingen die we redelijk (denken te) begrijpen: hoe het gas en de donkere materie er uitzagen in het hele jonge heelal, hoe de zwaartekracht werkt, en hoe gas beweegt, koelt en opwarmt door druk van ander gas. We denken ook redelijk te begrijpen hoe gas kan koelen door straling uit te zenden. In het heel vroege heelal waren het gas en de donkere materie ongeveer uniform verdeeld, wat wil zeggen dat de dichtheid en temperatuur van het gas overal ongeveer gelijk was. We weten dat het gas uniform verdeeld was door waarnemingen van licht uit het vroege heelal. Uit de verdeling van de zichtbare materie kunnen we ook afleiden wat de verdeling van de donkere materie geweest moet zijn. De zwaartekracht van de donkere materie heeft namelijk de verdeling van zichtbare materie



Figuur 1: Het kosmische web uit de EAGLE-simulatie. Het grootste plaatje laat een stuk van ongeveer 300 miljoen bij 300 miljoen lichtjaar zien, met een diepte van ongeveer 60 miljoen lichtjaar. De kleur geeft de temperatuur van het gas aan. Het blauwe gas is relatief koel (minder dan $30.000\text{ }^\circ\text{C}$), het groene gas is warmer (tot ongeveer $300.000\text{ }^\circ\text{C}$), en het rode gas is het heetst. Ik heb vooral het rode gas onderzocht. De kleinere plaatjes laten een sterrenstelsel zien en de omgeving van dit sterrenstelsel. Dit is fig. 1 uit [Schaye et al. \(2015\)](#).

mede gevormd.

In de eerste stap van sterrenstelselvorming zijn kleine dichtheidsverschillen belangrijk. Waar net iets meer materie is, trekt de zwaartekracht nog meer materie aan. Zo groeien er ‘pannenkoeken’: vlakke structuren met een hogere dichtheid dan de gebieden aan beide kanten ervan. Binnen de pannenkoeken en waar verschillende pannenkoeken elkaar snijden, ontstaan er slierten met nog hogere dichtheden: de filamenten. Binnen de filamenten, en waar filamenten elkaar snijden, ontstaan uiteindelijk min of meer bolvormige concentraties donkere materie en gas: de halo’s, waar de dichtheden het hoogst zijn. We noemen dit netwerk van pannenkoeken, filamenten, en halo’s het kosmische web, afgebeeld in Fig. [1](#)

In de halo’s is de gemiddelde dichtheid ongeveer honderd keer groter dan gemiddeld in het heelal. De dichtheid is het hoogst in het midden van de halo. De gasdeeltjes zitten daar dicht bij elkaar, dus ze botsen ook vaker. Bij zo’n botsing kan straling vrijkomen. Deze straling bestaat uit zichtbaar licht, maar ook uit lichtdeeltjes (fotonen) met hogere energie (UV en röntgenstraling) en lagere energie (infrarood- en radiostraling). De temperatuur is een maat voor de hoeveelheid energie in het gas. Als het gas fotonen uitzendt, verliest het gas energie en koelt het gas dus. Dit koelen gebeurt daarom vooral in het dichte gas midden in halo’s.

Als gas afkoelt, neemt de druk in dat gas af, waardoor het door het omliggende gas verder wordt samengedrukt. In dit koele, dichte gas ontstaan uiteindelijk klompen gas die dicht genoeg zijn dat ze door hun eigen zwaartekracht verder samentrekken. Hieruit vormen zich

sterren. Zo ontstaan verzamelingen sterren midden in de halo's van het kosmische web: dit zijn sterrenstelsels. Sterrenstelsels vormen nieuwe sterren uit gas midden in halo's, maar kunnen ook groeien door samen te smelten met andere sterrenstelsels. Als halo's samensmelten, maar de sterrenstelsels erin (nog) niet, ontstaan satellietsterrenstelsels: sterrenstelsels binnen een halo die om een 'centraal' sterrenstelsel midden in de halo draaien. We zien dat sterrenstelsels nu nog steeds sterren vormen. Er worden nog zoveel sterren geboren dat die stervorming niet verklaard kan worden met alleen het gas dat zich tussen de sterren bevindt, binnen de sterrenstelsels. Er moet dus nieuw gas aangevoerd worden uit de halo rondom het sterrenstelsel. Dat gas wordt het circumgalactisch medium (CGM) genoemd, en het gas tussen de halo's heet het intergalactisch medium (IGM).

In het algemeen ontstaan sterrenstelsels met meer massa in halo's met meer massa. Binnen deze halo's is er meestal (ongeveer) een evenwicht tussen de zwaartekracht, die gas dichter naar het midden van de halo trekt, en gasdruk, die het gas juist van het midden van de halo wegdukt. In zwaardere halo's moet het hete gas in het CGM daarom meer gasdruk hebben, dus warmer zijn, om te voorkomen dat het door de zwaartekracht in elkaar getrokken wordt. In die zin kun je een halo een beetje zien als een put. Een halo met meer massa is als een diepere put: als je op de bodem van de put staat en een balletje tot de bovenrand wilt gooien, zul je in een diepere put harder moeten gooien. Op dezelfde manier moet gas in een zwaardere halo sneller bewegen (dat wil zeggen, heter zijn) om de halo 'tot de rand' te blijven vullen. Op deze manier kun je een typische temperatuur uitrekenen voor halo's van een gegeven massa, waarbij de gasdruk en de zwaartekracht in evenwicht zijn.

Er zijn echter ook processen die we nog niet zo goed begrijpen, maar die wel belangrijk zijn voor de vorming van (realistische) sterrenstelsels. Als het gas alleen maar uit het IGM het CGM in zou stromen, zou koelen en sterren zou vormen, zouden sterrenstelsels meer sterren vormen dan we in waarnemingen vinden. Er moeten dus ook processen zijn die stervorming remmen. Die processen worden feedback genoemd. We weten dat supernova's en heel zware zwarte gaten genoeg energie kunnen produceren om stervorming voldoende te remmen door de toevoer van gas te beïnvloeden. Een supernova is een enorme ontploffing die plaatsvindt aan het einde van het leven van een zware ster. Veel van die ontploffingen samen kunnen gas een sterrenstelsel uit blazen. Als gas eenmaal een zwart gat in is gevallen, komt het er niet meer uit. Er kan echter wel veel energie vrijkomen uit gas dat naar het zwarte gat toe gezogen wordt, maar er nog niet in zit. Er komt namelijk veel straling vrij uit het zeer hete gas om het zwarte gat, en een deel van het gas wordt heel snel weggeblazen van het zwarte gat in plaats van dat het er in valt. Hoe supernova's en zwarte gaten de stervorming precies remmen en hoeveel energie ze daarvoor precies gebruiken, is echter niet bekend. Dat komt onder andere omdat supernova's individuele ontploffende zware sterren zijn. In EAGLE kunnen we individuele sterren en het gas eromheen niet fijnmazig genoeg bekijken; in plaats daarvan moeten we werken met grote klompen sterren en gas die vele malen groter zijn dan het gebied waar een enkele supernova in afgaat. In de EAGLE-simulatie is de feedback dus het resultaat van wat heel veel supernova's samen doen met een grote klomp gas. Op dezelfde manier kan er veel energie van het gas rondom superzware zwarte gaten komen, maar dit gebeurt ook op een schaal die te klein is voor de EAGLE-simulaties.

In de EAGLE-simulatie worden deze onzekerheden opgevangen door een model te maken van hoe de energie van supernova's en zwarte gaten het gas eromheen beïnvloedt. In de EAGLE-simulatie wordt dit gas heel sterk verhit; in andere simulaties wordt het gas bijvoorbeeld in beweging gezet. Die modellen hebben een aantal parameters (vrij instelbare getallen), zoals hoeveel energie er in een keer vrijkomt als er feedback plaatsvindt. In simulaties werden er verschillende waarden voor die parameters uitgetoet. De waarden

werden gekozen waarbij de hoeveelheid sterrenstelsels met verschillende massa's, de groottes van die sterrenstelsels en de massa's van zwarte gaten in sterrenstelsels van een gegeven massa het beste klopten met wat er wordt waargenomen.


Als we EAGLE met andere simulaties vergelijken, zien we dat er verschillende manieren zijn om vrij realistische sterrenstelsels te maken. De grootste verschillen daartussen zitten in de feedback. Een belangrijke vraag daarbij is hoe het precies de vorming van sterren remt. In EAGLE zorgt de feedback van supernova's ervoor dat veel gas waar sterren uit hadden kunnen vormen niet alleen uit het sterrenstelsel geblazen wordt, maar dat het ook veel gas de hele halo mee uitblaast die om het sterrenstelsel zit. Daarmee wordt stervorming dus geremd door de hoeveelheid stervormend gas in het sterrenstelsel direct te verminderen, maar ook door de bron van dit gas te legen. In een andere simulatie, namelijk IllustrisTNG, valt het gas dat sterrenstelsels uit geblazen wordt eerder terug het sterrenstelsel in. Dat zorgt er voor dat er in de EAGLE-simulatie veel minder CGM gas rondom sterrenstelsels zoals onze Melkweg zit dan in de IllustrisTNG-simulatie. Als we weten hoeveel gas er in het CGM is, hebben we dus ook een beter idee van hoe de feedback werkt.

Straling: hoe we het gas kunnen waarnemen

Op dit moment weten we niet hoeveel gas zich bevindt om sterrenstelsels als de Melkweg. We weten ongeveer hoeveel koel en warm gas (blauw en groen in Fig. 1) er rond sterrenstelsels is, maar we weten minder over het hete gas (rood in Fig. 1), omdat we dat moeilijker kunnen waarnemen. Dat komt omdat het koelere gas te vinden is met waarnemingen van UV-straling die we nu goed kunnen doen. Het hetere gas zorgt voor vergelijkbare effecten, maar dan in röntgenstraling, waar gelijksoortige signalen moeilijker te vinden zijn omdat de energierijkere fotonen moeilijker te meten zijn. Er zijn verschillende manieren waarop we het hete gas zouden kunnen vinden. Een aantal daarvan is al toegepast rondom sterrenstelsels die veel meer massa hebben dan de Melkweg. In dit proefschrift richt ik me op twee opties: absorptie- en emissielijnen. Die lijnen worden gevormd doordat straling van een heel specifieke energie wordt opgenomen (absorptielijn) of uitgezonden (emissielijn) door een elektron dat vast blijft zitten aan een atoomkern. Het hoofddoel van mijn promotieonderzoek is om uit te zoeken of en hoe we het hete gas in het CGM en IGM kunnen vinden door middel van waarnemingen van röntgen- en UV-straling.

In Fig. 2 laat ik voorbeelden van zulke emissie- en absorptielijnen zien. Als stoffen worden verhit gaan ze straling uitzenden. In het voorbeeld van Fig. 2 is dat zichtbaar licht. Bovenaan staat een voorbeeld van een continu spectrum: er wordt licht van alle kleuren uitgezonden. Naarmate iets heter wordt, wordt er meer straling uitgezonden en wordt er relatief meer straling met hogere energie uitgezonden. Straling bestaat uit lichtdeeltjes, die fotonen heten. Bij zichtbaar licht vormen de minst energetische fotonen rood licht, en de meest energetische vormen paars licht. Fotonen kunnen echter nog meer energie hebben dan die van zichtbaar licht: dan vormen ze ultraviolette (UV) straling, of zelfs röntgenstraling.

Absorptielijnen ontstaan doordat bepaalde gasdeeltjes licht absorberen op hele specifieke golflengtes. Als je het licht bijvoorbeeld met een prisma in verschillende golflengtes opsplijt, zodat je een regenboog ziet, dan mist er een heel specifieke kleur licht. Het onderste plaatje in Fig. 2 laat zien hoe het eruit ziet als zo'n stof voor een bron staat die het licht uit het bovenste plaatje produceert. Op een soortgelijke manier kun je in het licht van sommige bronnen juist extra felle lijnen in een regenboog zien: dat zijn emissielijnen. Een voorbeeld van emissielijnen staat midden in Fig. 2. Een bepaalde stof produceert emissie- en absorptielijnen op dezelfde energie (met dezelfde kleur). De golflengtes van die lijnen zijn



Figuur 2: Voorbeelden van verschillende soorten spectra. De voorbeeldbron bovenaan zendt licht uit in alle kleuren. De bron in het midden zendt alleen licht uit van heel specifieke kleuren: emissielijnen. De onderste bron zendt hetzelfde licht uit als het bovenste, maar er zit wat materiaal tussen de bron en de waarnemer. Dat materiaal absorbeert licht met heel specifieke kleuren: absorptielijnen. Het materiaal in de middelste bron dat licht uitzendt is hetzelfde als het materiaal dat onderaan licht absorbeert. Stoffen absorberen namelijk licht van dezelfde kleuren die ze uitzenden als ze heet zijn. Deze figuur komt van http://teacherlink.ed.usu.edu/tlnasa/reference/imaginedvd/files/imagine/docs/science/how_l1/spectra.html

zo specifiek dat we aan de lijn kunnen zien wat voor stof de lijn geproduceerd heeft. Aan de hand van de lijnen die we vinden, kunnen we er dus achter komen wat voor stoffen er in een wolk gas zitten. Voor de röntgen- en UV-lijnen die ik onderzocht heb, zijn die stoffen verschillende soorten atomen en ionen.

Absorptie- en emissielijnen ontstaan in röntgen- en UV-straling op dezelfde manier als in zichtbaar licht. Het hete gas om en tussen sterrenstelsels dat ik bestudeer is zo heet dat het vooral röntgenstraling uitzendt en absorbeert. Als we willen zoeken naar spectraallijnen (absorptie- of emissielijnen), kunnen we voor dit hete gas dus het beste naar röntgenlijnen zoeken. Met de telescopen die nu beschikbaar zijn is dit erg moeilijk. We kunnen met die telescopen namelijk niet genoeg röntgenstraling opvangen en die niet scherp genoeg in verschillende kleuren splitsen, waardoor we maar heel zelden duidelijke lijnen kunnen identificeren. Dat geldt zelfs voor de helderste achtergrondbronnen die beschikbaar zijn. Voor emissielijnen is het nog lastiger omdat het hete CGM gas relatief ijl is, en daarom überhaupt vrij weinig straling uitzendt. Met toekomstige instrumenten zullen we hopelijk meer kunnen leren over deze lijnen. Dat komt omdat ze meer straling zullen kunnen opvangen, en ze de energie ('kleur') van de fotonen beter zullen kunnen meten.

De stoffen waarvan we de emissie- en absorptielijnen zoeken zijn vooral bepaalde ionen. Ionen zijn deeltjes met een positieve of negatieve lading. Gas zoals we dat op Aarde kennen bestaat uit moleculen, die weer uit atomen bestaan. Die atomen bestaan weer uit een kern, met een wolk elektronen eromheen. Die elektronen kleven aan de kern omdat ze negatief geladen zijn, terwijl de kern positief geladen is. Een atoom is een kern met precies zoveel elektronen eromheen dat het geheel neutraal of ongeladen is. Als we beginnen met gas zoals we hier op Aarde om ons heen vinden, en we dat gaan verhitten, botsen de moleculen op een gegeven moment zo vaak zo hard op elkaar, dat de moleculen uit elkaar geslagen worden, en losse atomen overblijven. Dat is het geval voor bijvoorbeeld het blauwe gas in Fig. 1. Als het gas nog heter wordt, worden zelfs elektronen van de atomen afgeslagen. Hoe heter het gas wordt, hoe minder elektronen er overblijven. Een atoom dat een of meerdere elektronen kwijt is heet een ion. Elektronen kunnen ook door fotonen (straling) van de kernen worden losgeslagen. Dat is vooral belangrijk in relatief ijl gas, waar atomen en ionen snelle, losgeslagen elektronen minder vaak tegenkomen.

Een ion kan ook juist een elektron 'vangen' als het positief geladen ion en het negatief geladen elektron elkaar tegenkomen. Of ion een elektron vangt of kwijtraakt bij een bot-

sing hangt vooral van de snelheid van het elektron af. Een langzaam elektron wordt eerder gevangen, een snel elektron slaat juist eerder een ander elektron los. Die snelheid hangt af van de temperatuur: een hogere temperatuur betekent dat er meer snelle elektronen zijn, en minder langzame. Bij elke temperatuur kan er zo een evenwicht ontstaan. Hierbij is van elk element het aandeel ionen/atomen van elk elektronenaantal stabiel, doordat de ionen per seconde evenveel elektronen vangen als dat ze ze verliezen. Elk ion heeft daarbij een 'optimale temperatuur', waarbij het aandeel van dat ion in alle atomen en ionen van een element het grootst is. De dichtheid bepaalt ook mede het aantal elektronen dat elk ion per seconde vangt en verliest, omdat de elektronen elkaar bij een hogere dichtheid vaker tegenkomen. Deze toename van het aantal botsingen met de dichtheid maakt echter niet uit voor het evenwicht. Het evenwicht wordt namelijk alleen bepaald door de verhouding tussen het aantal botsingen waarbij een elektron wordt gevangen en het aantal waarbij een elektron wordt losgeslagen. Het evenwicht hangt daarom vaak alleen van de temperatuur af. Alleen als het gas zo ijl is dat botsingen zeldzaam worden dan kan ook het wegslaan van elektronen door straling belangrijk worden. Dat wegslaan door straling hangt niet van de dichtheid af. Omdat het vangen van elektronen wel van de dichtheid afhangt, wordt het evenwicht in dit ijle gas ook afhankelijk van de dichtheid.

In mijn berekeningen ben ik uitgegaan van zo'n evenwicht in de ionverdeling, inclusief het effect van ionisatie (het losslaan van een elektron) door straling. In het CGM rond sterrenstelsels die minstens zoveel massa hebben als onze Melkweg, is het gas meestal dicht genoeg dat de ionverdeling niet afhangt van de dichtheid, in elk geval in de binnendelen. De straling is echter wel van belang in de filamenten van het kosmische web, waaruit absorptielijnen ook te meten zullen zijn.

Wat we kunnen zien, en wat het betekent

Het kan dus vrij ingewikkeld zijn om te bepalen welke spectraallijn we in welke omgeving verwachten te vinden, en als we een spectraallijn van een bepaald ion vinden, is het niet eenvoudig om te achterhalen hoeveel en wat voor gas die lijn gemaakt heeft. Een belangrijk deel van het werk in mijn proefschrift was om daarop beter zicht te krijgen. Een voorbeeld van die onzekerheid ligt in het aandeel zuurstof in het gas. Lijnen van twee zuurstofionen zijn de sterkste die we verwachten te vinden in het hete gas. Zuurstof maakt echter maar een klein deel uit van alle atomen/ionen in dat gas; bijna alle massa van het gas zit in waterstof en helium. In het hete gas dat ik voor mijn proefschrift onderzocht heb, zijn alle elektronen echter van deze kernen afgeslagen, waardoor ze geen spectraallijnen kunnen vormen. Verder wordt zuurstof, net als alle andere elementen zwaarder dan waterstof en helium, in sterren gevormd, en was het niet aanwezig in het vroege heelal. Dat zuurstof komt in het gas rond en tussen sterrenstelsels terecht door bijvoorbeeld de feedback van supernova's en zwarte gaten die gas het sterrenstelsel uit blaast. Het aandeel zuurstof in het gas neemt dus over het algemeen toe naarmate je dichterbij een sterrenstelsel komt, en zegt iets over de feedbackprocessen. Dat het aandeel zuurstof in het gas niet constant is, betekent echter ook dat het moeilijk is de hoeveelheid gas om een sterrenstelsel te bepalen door middel van alleen de spectraallijnen van dat element. Ik heb voorspellingen van bijvoorbeeld de zuurstofverdeling rond sterrenstelsels uit de EAGLE-simulaties gebruikt om te bepalen wat voor spectraallijnen we met verschillende instrumenten zouden kunnen zien. Ik heb ook voorspellingen gemaakt over wat we wel en niet kunnen concluderen over het gas dat we waarnemen aan de hand van de spectraallijnen die we vinden.

We zagen eerder dat bij elke massa van een halo een bepaalde typische temperatuur

hoort. Dat betekent dat we in halo's van verschillende massa ook verschillende spectraallijnen verwachten, omdat bij elke temperatuur andere ionen optimaal aanwezig zijn. Het blijkt dat het aandeel ionen van de soorten die ik onderzocht heb, en dus ook de sterkte van de absorptielijnen van dat ion, het grootst is wanneer deze typische temperatuur goed overeenkomt met de temperatuur waar dat ion vaak voorkomt. Dat is niet geheel verrassend. Toch is het niet vanzelfsprekend: een specifieke temperatuur waarbij het ion vaak voorkomt bestaat alleen in relatief dicht gas en aan de randen van halo's is het gas niet per se dicht genoeg hiervoor. Hier kan ionisatie door straling namelijk belangrijk zijn. Daarnaast bestaat het CGM niet alleen uit gas van een temperatuur nabij de typische waarde, maar uit gas van allerlei verschillende temperaturen. Toch is de typische temperatuur een goede maat voor de temperatuur van het meeste hete gas in een halo.

De sterkte van een absorptielijn hangt er vanaf hoeveel van het ion dat die lijn veroorzaakt tussen de lichtbron en de waarnemer staat. Dat hangt weer af van de dichtheid van het gas, het aandeel van het element (de atoomkern) dat bij dat ion hoort in dat gas, en de ionverdeling van dat element.

De sterkte van röntgenemissielijnen hangt sterker van de dichtheid van het gas af dan de sterkte van absorptielijnen. Dat komt omdat botsingen tussen een ion en elektron (of lichtdeeltje) nodig zijn om het ion de energie te geven die nodig is om licht uit te zenden. De sterktes van absorptie- en een emissielijnen hangen dus allebei af van hoeveel ionen er in het gas zijn. Voor de emissielijnen neemt de sterkte van de lijn echter ook nog toe naarmate er meer elektronen in de buurt zijn om op de ionen botsen. Als het gas dicht genoeg is dat de botsingen met lichtdeeltjes onbelangrijk zijn, hangt de temperatuur waar de emissie het sterkst is echter niet van de dichtheid af. Dat komt omdat de ionverdeling, zoals eerder uitgelegd, alleen van de temperatuur afhangt in dicht gas. Het aantal elektronen in de buurt van een ion wordt twee keer zo groot als de dichtheid twee keer zo groot wordt, voor elke gastemperatuur. Als de dichtheid hoog genoeg wordt, is de temperatuur waar een emissielijn het sterkst is dus niet afhankelijk van de dichtheid. De sterkte van die emissielijn neemt echter nog wel toe naarmate de dichtheid hoger wordt. Op deze manier is er ook een karakteristieke temperatuur voor emissielijnen.

Dit zorgt ervoor dat de sterkte van een bepaalde emissielijn uit een halo op een vergelijkbare manier als de absorptielijnen afhangt van de typische temperatuur van een halo van een gegeven massa. Voor sommige emissielijnen blijft de hoeveelheid lijnemissie echter toenemen met halomassa wanneer de typische halotemperatuur groter is dan de temperatuur waar de emissielijn het sterkst is. Dit komt voor bij de sterkste emissielijn van ionen die nog maar een elektron vasthouden. Die toenemende emissie heeft er mee te maken dat deze emissielijnen nog relatief sterk zijn in gas dat heter is dan bij het emissiemaximum, en dat een halo met meer massa, meer CGM heeft om die straling uit te zenden. De emissie mag dan per kilogram typisch zwakker zijn, maar dat wordt met meer kilo's gecompenseerd.

In het hete gas dat ik bestudeer zijn de lichtste elementen, waterstof en helium, hun elektronen allemaal kwijtgeraakt. Omdat absorptie- en emissielijnen gevormd worden door elektronen die energie opnemen (absorptie) of kwijtraken (emissie) terwijl ze vast blijven zitten aan een atoomkern, kunnen we in dit hete gas geen lijnen van waterstof of helium vinden. Dat is belangrijk omdat dit de elementen zijn die het meeste voorkomen, en waarvan we anders dus sterke lijnen hadden kunnen verwachten. Verder zijn dit de enige elementen die aanwezig waren in het hele vroege heelal. Zwaardere elementen zijn later ontstaan uit waterstof en helium. Deze elementen zijn ontstaan in sterren, dus in sterrenstelsels. Waar we (lijnen van) deze zwaardere elementen vinden heeft dus ook te maken met hoe (ver) gas uit sterrenstelsels geblazen is.

De sterkste absorptielijn die we uit dit hete gas verwachten komt van het ion O VII; dat is een zuurstofatoom dat zes elektronen is kwijtgeraakt, en er dus maar twee overheeft. Het ion O VIII, een zuurstofatoom dat nog maar één elektron over heeft, zorgt ook voor een sterke absorptielijn, en de sterkste emissielijn in het hete gas.

Met de toekomstige ruimtetelescoop Athena kunnen we relatief sterke röntgenabsorptielijnen van O VII en O VIII vinden. Deze lijnen komen vooral van relatief dicht gas, met relatief veel elementen die in sterren gevormd zijn, en vooral van het gas rondom sterrenstelsels. Je vindt daarmee vooral gas met een temperatuur die optimaal is voor dat ion, ook als het merendeel van het aanwezige gas een andere temperatuur heeft. Met de voorgestelde ruimtetelescopen Arcus en Lynx zou je ook zwakkere lijnen moeten kunnen vinden. Deze zwakkere lijnen komen meestal van ijler gas. In dit ijlere gas is het losslaan van elektronen door straling belangrijk, naast het losslaan van elektronen door botsingen met andere elektronen. Deze bijdrage van straling zorgt ervoor dat atomen bij een gegeven temperatuur meer elektronen kwijtraken, waardoor een zwakkere absorptielijn van een gegeven ion van gas op lagere temperaturen kan komen dan sterkere absorptie van datzelfde ion.

De sterkste absorptielijnen die we met deze telescopen kunnen vinden zullen van het relatief dichte gas met relatief veel zuurstof komen. Dat wil zeggen dat deze lijnen vooral van het CGM, rondom sterrenstelsels, komen. Zwakkere lijnen kunnen ook van gas in het CGM komen, maar ook van ijler gas in het IGM, zoals in de filamenten van het kosmische web. Athena is onder andere bedoeld om röntgenabsorptielijnen te zoeken, door een gepland aantal röntgenbronnen waar te nemen. De telescoop zal gevoelig genoeg zijn om een O VII absorptielijn te vinden in de meeste röntgenbronnen die zo onderzocht worden, als de halo van een sterrenstelsel zoals onze Melkweg tussen ons en die bron staat. Hoewel het aandeel van elementen als zuurstof, die in sterren gemaakt worden, in het CGM afneemt naarmate de afstand tot het centrale sterrenstelsel toeneemt, neemt het aandeel van deze elementen in het gas dat voor röntgenabsorptielijnen zorgt, nauwelijks af. Dat komt omdat de zuurstoflijn alleen sterk genoeg is om waar te kunnen nemen op plekken waar toevallig relatief veel zuurstof zit. Zulke plekken zijn zeldzamer naarmate we verder van een sterrenstelsel kijken, waar de zuurstof vandaan komt. Bij een waarneming kunnen we echter alleen die zeldzame sterke lijnen vinden, waardoor we een verkeerde indruk kunnen krijgen van wat typisch is voor het gas ver van een sterrenstelsel.

Emissielijnen komen vooral van gas met een relatief hoge dichtheid en een hoge temperatuur. Dat betekent dat emissielijnen vooral komen van gas in de binnendelen van het CGM, waar het gas veel dichter en een beetje heter is dan in de buitendelen. Met Athena zullen we in een paar emissielijnen (bijna) het hele CGM in sommige halo's kunnen zien, met name die (veel) zwaarder zijn dan de halo rondom de Melkweg. Rondom sterrenstelsels als de onze zullen we een paar emissielijnen uit de binnendelen van het CGM kunnen vinden.

Ik heb ook een aantal gemeten absorptielijnen vergeleken met de voorspellingen uit de EAGLE-simulaties. Die metingen zijn wat onzeker, omdat röntgenabsorptielijnen met de huidige instrumenten moeilijk waar te nemen zijn. Daarnaast heb ik de simulaties met lijnen van maar vier gaswolken vergeleken. Dat maakt die vergelijking moeilijk: in een simulatie vind je namelijk nooit precies dezelfde gaswolk als in de waarnemingen. Wat je wel kunt doen is kijken of die wolken ongeveer even vaak voorkomen als een simulatie voorspelt en of de eigenschappen van die wolken redelijk overeenkomen met die uit de simulatie. Dit soort vergelijkingen van populaties zijn lastig met weinig metingen, omdat je dan niet weet of je een typische gaswolk of net een atypische gevonden hebt.

Gegeven de sterktes van de absorptielijnen lijken de temperaturen van twee gemeten gaswolken redelijk te kloppen met de voorspellingen uit EAGLE. In eerste instantie leken

wolken met röntgenabsorptielijnen die zo sterk zijn dat we ze kunnen meten wat vaker voor te komen dan voorspeld in EAGLE, maar dat bleek toch mee te vallen: er is op een aantal plekken gezocht naar die lijnen en op veel plekken is niks gevonden.

Mogelijk ligt het gas dat sterke röntgenabsorptielijnen produceert in het echt wat verder van sterrenstelsels dan voorspeld in EAGLE, maar dat is gebaseerd op twee gemeten wolken. Van één van de twee is niet helemaal zeker of die er ook echt is. Op dit moment is het dus moeilijk te zeggen of dit een probleem is, of dat één of twee wolken toevallig wat verder van sterrenstelsels lagen dan veel andere. Als we in de toekomst meer gevoelige instrumenten hebben, zullen we meer van dit soort metingen kunnen doen. Het zal interessant zijn om deze op dezelfde manier met simulaties te vergelijken. Met deze vergelijkingen kunnen we namelijk leren hoe ver stoffen die in sterren gemaakt worden van sterrenstelsels worden weggeblazen door feedback.

Al met al lijken de vooruitzichten goed voor waarnemingen van het hete gas rond sterrenstelsels met toekomstige ruimtetelescopen. Met röntgenemissielijnen zullen we veel kunnen leren over gas rond sterrenstelsels zwaarder dan onze Melkweg. Het zal moeilijk worden om emissielijnen te vinden van heet gas rond sterrenstelsels die ongeveer even zwaar zijn als onze Melkweg, en van gas in de filamenten van het kosmische web. We kunnen met absorptielijnen echter ook dit gas waarnemen. Met deze waarnemingen zullen we meer te weten komen over de gasstromen rondom sterrenstelsels, en hoe deze de vorming van sterren beïnvloeden.

English summary

From gas to galaxies

In this thesis, I have investigated ways to find the hot gas around and between galaxies. I have done this using computer simulations: the EAGLE simulations. EAGLE stands for 'Evolution and Assembly of GaLaxies and their Environments'. In those simulations, we start with 3 billion 'pieces' of dark matter and as many 'pieces' of gas, in a large (simulated) region of the universe. This region is large enough to contain hundreds of galaxies with roughly the mass of our Milky Way. Dark matter is a substance that moves right through everything, including other dark matter. It does exert gravity, on other dark matter and on 'normal' matter, like gas, stars, planets, and theses. We found dark matter by discovering that something in the Universe is exerting gravity on gas and stars. We cannot see that something, and we do not know exactly what it is, which is why it is called dark matter.

In the EAGLE simulation, we divide that dark matter and the gas into chunks, because in a computer, you can only store and calculate a finite list of temperatures, densities, positions, etc. These chunks are a bit like the pixels in a digital photo: with enough small pieces, you can make a good approximation of what you would actually see. The size of the pixels determines how much detail you can capture. The chunks of gas in the simulation represent balls of gas that move, expand, and contract due to gravity, and pressure from the surrounding gas. Each chunk of gas has a position, velocity, temperature, density, etc. In reality, there may be gas at a range of temperatures and densities within the volume encompassed by one of these chunks. However, in the simulation, we only have the single temperature and density of the chunk, like one green pixel in a photo sometimes represents all the leaves in a tree. Though gas is made up of particles, these particles are much, much smaller than the chunks of gas in the simulation. Each chunk contains as much mass as 2 million suns. The stars in the Milky Way contain about 20,000 times this mass. When stars form from these chunks of gas, the simulation does not track individual stars. Instead, it tracks big groups of stars that form at the same time and move together.

In these simulations, we think we understand some of the ingredients reasonably well. These include the properties of the dark matter and gas in the early Universe, how gravity works, and how gas moves, heats up, and cools due to pressure from other gas. We also think we have a reasonable understanding of how gas cools by emitting light. In the early Universe, the distribution of gas and dark matter was close to uniform. That means that its temperature and density was almost exactly the same everywhere. We know the gas distribution was uniform from observations of light from the early Universe. We can use the distribution of this gas in the early Universe to infer what the dark matter distribution must have been. This is possible because the dark matter's gravity has helped shape the distribution of the gas.

The small differences in gas density throughout the early Universe are important for the first step of galaxy formation. Where there is slightly more mass, gravity attracts even more




Figure 3: The cosmic web in the EAGLE simulation. The largest picture shows a roughly 300 million by 300 million light year region, with a depth of about 60 million light years. The colour indicates the temperature of the gas. The blue gas is relatively cool (less than $30,000\text{ }^\circ\text{C}$), the green gas is warmer (up to about $300,000\text{ }^\circ\text{C}$), and the red gas is the hottest. I have primarily investigated the red gas. The smaller pictures show a galaxy and its environment. This is fig. 1 of [Schaye et al. \(2015\)](#).

matter. This causes ‘pancakes’ of gas and dark matter to grow. These are flat structures with a higher density than the areas to either side of them. Within these pancakes, and where pancakes intersect, filaments form with even higher densities. Within those filaments, and where the filaments intersect, roughly spherical concentrations of dark matter and gas eventually form. These are haloes, and they have the highest densities of these structures. This network of pancakes, filaments, and haloes is called the cosmic web, shown in Fig. 3

In the haloes, the average density is about a hundred times that of the average throughout the Universe. The density is highest at the centre of the halo. The gas particles are closer together there, so they collide more often. In these collisions, light can be emitted. That can be visible light, but it can also consist of light particles (photons) with more energy (ultraviolet and X-ray photons), or with lower energy (infrared and radio photons). Temperature is a measure of the energy content of the gas. When gas emits photons, it loses energy and therefore, it cools. This cooling therefore mostly happens in the dense gas at the centres of haloes.

When gas cools, the pressure in the gas decreases, and it is compressed by the pressure of the gas around it. In this cool, dense gas, clumps eventually form that are dense enough to collapse further under their own gravity. Stars form from these clumps. This is how collections of stars form in the middle of the haloes of the cosmic web. These are galaxies. Galaxies form new stars from gas in halo centres, but they can also grow by merging with other galaxies. When haloes merge, but the galaxies within them do not merge (yet), satellite galaxies form. These are galaxies within a halo that orbit the central galaxy at its centre. We

still see galaxies form stars. So many new stars are born that their formation cannot be explained with only the gas that is present among the stars of a galaxy. This means this gas must be supplemented by new gas from the halo surrounding the galaxy. That gas is called the circumgalactic medium (CGM), and the gas between haloes is called the intergalactic medium (IGM).

In general, more massive galaxies form within more massive haloes. Within these haloes, there is usually a rough equilibrium between the gravity pulling the gas towards the halo centre, and the gas pressure, pushing the gas apart. In more massive haloes, the hot gas in the CGM must therefore have a higher pressure, and must therefore be hotter, to prevent it from collapsing under the halo's gravity. In this sense, the halo is a bit like a well. A more massive halo is like a deeper well: if you're standing at the bottom and throw a ball up to the edge, you need to throw it harder in a deeper well. In the same way, gas in a more massive halo needs to move faster (i.e., be hotter) to keep filling the halo up to its 'edge'. Using this assumed balance, we can calculate a typical temperature for haloes of a given mass, at which the gas pressure and gravity are in equilibrium.

There are other processes which are important to galaxy formation that we do not understand very well. If gas only ever flowed from the IGM sheets and filaments into the CGM in the haloes of the cosmic web, cooled there, and formed stars, the Universe would have formed many more stars than we observe. There must therefore also be processes at play that decrease or stop star formation. Those processes are known as feedback. We know that supernova explosions and supermassive black holes can produce enough energy to decrease star formation to the degree required by observations by affecting the gas supply of galaxies. A supernova is an enormous explosion at the end of the life of a massive star. The combined effects of many of those explosions can blow gas out of a galaxy. Once gas falls into a black hole, it cannot come back out. However, energy can be released from gas that is being pulled towards a black hole, but is not yet inside it. This is because the very hot gas swirling around a black hole emits a lot of light, and because some of the gas is blown away from the black hole at very high speed instead of falling in. However, we do not know exactly how supernovae and black holes limit star formation, or exactly how much energy they use to do that. This is in part because supernovae are individual exploding stars. However, we cannot resolve individual stars and the gas around them in EAGLE; instead we have to work with big clumps of gas and stars much larger than the region in which a single supernova explosion takes place. In the EAGLE simulations, the feedback is therefore the result of what many supernovae together would do to a big clump of gas. In the same way, a lot of energy can come from the gas around black holes, but this happens on a scale which is much too small for the EAGLE simulations.

In the EAGLE simulation, these uncertainties were dealt with by making a model of how the energy of supernova explosions and supermassive black holes affects the gas around them. In EAGLE, this gas is heated up very much; in other simulations, the gas may, for example, be set in motion. These models have a few free parameters: numbers that can be set, such as how much energy is released at once when feedback occurs. In simulations, different values for these parameters were tried. The values that were chosen were the ones that produced the right numbers of galaxies of different masses, reasonable sizes for those galaxies, and black holes of the right masses in galaxies of a given mass. The correct values are known from observations.

Comparing EAGLE to other simulations, we see that there are different ways to make realistic galaxies. The greatest differences between these methods lie in the feedback. An important question there is exactly how it reduces star formation. In EAGLE, supernova

feedback expels much of the gas that could have formed stars from galaxies, but it also takes a lot of gas from the halo with it, expelling it from the halo surrounding galaxy. This means that star formation is decreased by reducing the gas supply in the galaxy, but also by depleting the source of this galactic gas. In a different simulation, IllustrisTNG, gas expelled from the galaxy tends to fall back onto the galaxy. This means that in the EAGLE simulation, there is much less CGM gas around galaxies like our Milky Way than there is in the IllustrisTNG simulation. If we learn how much gas there is in the CGM, we will therefore also have a better idea of how feedback operates.

Radiation: how we can observe the gas

We do not currently know how much gas there is around galaxies like the Milky Way. We know roughly how much cool and warm gas (blue and green in Fig. 3) there is around galaxies, but we know less about the hot gas (red in Fig. 3), because it is harder to observe. This is because we can find the cooler gas with observations of ultraviolet (UV) radiation which we can currently do well. The hotter gas has similar effects, but in X rays, where similar signals are more difficult to find because the more energetic photons are more difficult to measure. There are different ways we could find the hot gas. A number of them have already been applied around galaxies much more massive than the Milky Way. In this thesis, I focus on two options: absorption and emission lines. Those lines are formed because radiation at very specific energies is absorbed (absorption line) or emitted (emission line) by an electron that remains bound to an atomic nucleus. The main goal of my thesis is to figure out if and how we can observe the hot gas in the CGM and IGM using X-ray and UV observations.

In Fig. 4 I show examples of such emission and absorption lines. When substances are heated, they emit radiation. In the example in Fig. 4, that radiation is visible light. The top image shows a continuous spectrum: light of all colors is emitted. As something is heated up, it emits more radiation, and relatively more radiation at high energies. This radiation is made up of light particles, called photons. For visible light, the lowest-energy photons form red light, and the highest-energy photons are purple. However, photons can have more energy than the ones we can see: they then form ultraviolet (UV) radiation or even X rays.

Absorption lines form because certain particles in a gas absorb light at very specific energies. For example, if you split light into colors using a prism, creating a rainbow, a very specific color of light will be missing. The bottom image in Fig. 4 demonstrates what that looks like, if those particles are between the observer and the source producing the light in the top picture. Similarly, there are extra bright lines in a rainbow produced from the light in some sources: those are emission lines. The middle image in Fig. 4 shows an example of those. A given substance produces emission and absorption lines at the same energy (with the same color). The energies of these lines are so specific that we can tell by the line which substance produced it. Using the lines we find, we can therefore figure out what substances make up a cloud of gas. For the X-ray and UV lines I have investigated, those substances are different kinds of atoms and ions.

Absorption and emission lines in UV radiation and X rays form in the same way as in visible light. The hot gas I study around and between galaxies is so hot it mainly emits and absorbs X rays. Therefore, a search for spectral lines (absorption and emission lines) from this gas is best done in X rays. This is very difficult with the telescopes that are currently available. This is because we cannot catch enough X-ray photons or split the light into different colors precisely enough, making it difficult to clearly identify lines. This is true even for the brightest background sources available. It is even more difficult to find emission




Figure 4: Examples of different kinds of spectra. The example source at the top emits light of all colors. The source in the middle only emits very specific colors of light: emission lines. The bottom source emits the same light as the one at the top, but there is some material between the source and the observer. That material absorbs light of very specific colors: absorption lines. The material emitting light in the middle panel is the same as the material absorbing light at the bottom. We can tell because materials absorb light of the same colors they emit when they are hot. This figure is from http://teacherlink.ed.usu.edu/tlnasa/reference/imaginedvd/files/imagine/docs/science/how_l1/spectra.html

lines, because the CGM gas is relatively diffuse, and therefore emits fairly little radiation at all. With future instruments, we will hopefully learn more about these lines. That will be possible because they will be able to ‘catch’ more of the radiation, and they will measure the energy (‘color’) of the photons more precisely.

The substances of which we want to find the absorption and emission lines are mostly specific ions. Ions are particles with a positive or negative charge. Gas we are familiar with on Earth consists of molecules, which are made up of atoms. Those atoms consist of a nucleus, surrounded by a cloud of electrons. The electrons are bound to the nucleus because they carry a negative charge, and the nucleus carries a positive charge. An atom is a nucleus with exactly enough electrons around it that the combination carries no net charge. If we take gas like we have around us here on Earth and heat it up, the molecules will start to collide harder and more often, until they are smashed apart and only loose atoms remain. This is the state of, for example, the blue gas in Fig. 1. If the gas gets even hotter, the electrons are knocked loose from their atoms. As the gas gets hotter, fewer and fewer electrons remain bound to their nuclei. An atom that has lost one or more electrons is called an ion. Electrons can also be knocked loose from their nuclei by photons. This is mainly important in relatively diffuse gas, where atoms and ions encounter fast, unbound electrons less often.

An ion can also ‘catch’ an electron when the positively charged ion and the negatively charged electron encounter each other. Whether an ion catches an electron or loses one in an encounter mainly depends on the electron’s speed. A slow electron is more likely to be caught, while a fast electron is more likely to knock off another one. That speed depends on the temperature of the gas: a higher temperature means there are more fast electrons, and fewer slow ones. At each temperature, an equilibrium can be reached. In this situation, the fraction of atoms/ions with each number of electrons is stable, because the ions catch and lose the same number of electrons each second. Each ion has an ‘optimal temperature’ where its fraction of the total number of atoms and ions reaches a maximum. The density also helps determine the number of electrons an ion catches and loses each second, because the atoms/ions and electrons encounter each other more often when the density is higher. However, this increase in the number of collisions with density does not change the equilibrium. This is because the equilibrium is only determined by the *ratio* of the rates at which electrons are captured and lost. Therefore, the equilibrium only depends on the temperature. The effect of photons knocking electrons loose only becomes important when the gas is diffuse enough that electron-ion collisions become rare. Those photon collisions are in-

dependent of the gas density. Because the rate at which electrons are caught does depend on the gas density, the equilibrium depends on density as well as temperature in this diffuse gas.

In my calculations, I have assumed the gas is in such an equilibrium state. This includes the effect of ionisation (knocking off an electron) by photons. In the CGM around galaxies at least as massive as our Milky Way, the gas is mostly dense enough that the ion distribution does not depend on density, at least in the inner regions. However, the effects of radiation on the ion distribution are important in the filaments of the cosmic web, where absorption lines should also become detectable.

What we can see and what it means

We have seen that it can be fairly complicated to figure out which spectral line we can find in which environment, and if we find a spectral line from a given ion, it is not easy to figure out how much and what sort of gas produced it. An important part of the work in my thesis is to get a better idea of that. An example of the uncertainty is the fraction of the gas that consists of oxygen. Lines from two oxygen ions are the strongest ones we expect to find in the hot gas. However, oxygen only makes up a small fraction of all atoms/ions in that gas; almost all the gas mass is in hydrogen and helium. In the hot gas I have studied in this thesis, hydrogen and helium have lost all their electrons. This means they cannot form spectral lines. Like the other elements heavier than hydrogen and helium, oxygen is formed in stars, and was not present in the early Universe. The oxygen is therefore formed in galaxies, and ended up in the CGM and IGM because, for example, the feedback from supernovae and black holes blew gas out of the galaxies. The fraction of oxygen in the gas therefore tends to increase as you get closer to galaxies, and tells us something about feedback processes. However, the fact that the fraction of oxygen in the gas varies means that it is difficult to determine the amount of gas around a galaxy using only spectral lines from that element. I have used predictions of, for example, the distribution of oxygen around galaxies in the EAGLE simulations to determine which spectral lines should be observable with which instruments. I have also made predictions for what we can and cannot conclude about the gas we observe based on the spectral lines we find.

We have seen previously that there is a typical temperature for a halo of a given mass. That means we also expect different spectral lines in haloes of different masses, because different ions are optimally present at different temperatures. It turns out that the fractions of the ion species I have studied, and therefore the strengths of their absorption lines, are greatest when this typical halo temperature matches the temperature where that ion is most prevalent. That is not entirely surprising. However, it is not a trivial result: that specific temperature for ions only applies to relatively dense gas, and at the edges of haloes, gas may be too diffuse for that temperature to apply. That is because there, ionisation by radiation can be important. Another issue is that the CGM does not just consist of gas close to the typical temperature, but is instead made of of gas at different temperatures. Nonetheless, the typical temperature is a good measure of the temperature of most of the hot gas in a halo.

The strength of an absorption line depends on how much of the ion causing the absorption is between the light source and the observer. This depends on the gas density, on the fraction of the element (nucleus) that makes up the ion in the gas, and on what fraction of those nuclei are found in the different ions.

The strength of X-ray emission lines depends more strongly on the density of the hot gas than the strength of the absorption lines. This is because a collision between an electron (or a

photon) and an ion is necessary to give the ion the energy it needs to emit light. The strengths of emission and absorption lines therefore both depend on how many ions producing the lines there are in the gas. However, the strengths of emission lines also increase as the number of nearby electrons an ion can collide with increases. However, if the gas is dense enough that photon collisions are unimportant, the temperature where an emission line is strongest does not depend on the density. This is because the ion distribution, as explained previously, only depends on temperature in dense gas. The number of electrons around an ion doubles when the density doubles at each gas temperature. In sufficiently dense gas, the temperature where an emission line is strongest therefore does not depend on the density. The strength of that line does, however, increase as the density increases. Therefore, there is also a characteristic temperature for emission lines.

This means that the strength of a given emission line from a halo depends on the typical temperature of a halo of a given mass in the the same way as for absorption lines. However, the strengths of some emission lines continue to increase with halo mass even when typical halo temperatures are above the optimal value for the emission line. This happens with the strongest emission lines of ions that only have a single electron left. This emission continues to increase because these emission lines are still relatively strong in gas hotter than the temperature where the emission is greatest, and because these hotter, more massive haloes simply contain more gas to emit the line. The emission per pound of gas does decrease, but that is compensated by more pounds.

In the hot gas I have studied, the lightest elements, hydrogen and helium, have lost all their electrons. Because absorption and emission lines are formed by electrons that absorb energy (absorption lines) or lose energy (emission lines) while remaining bound to the nucleus, we cannot find hydrogen or helium lines in this hot gas. That is important because those are the most prevalent elements, from which we would otherwise expect strong lines. These are also the only elements that were present in the early Universe. Heavier elements formed later from these two. They formed in stars, and therefore in galaxies. Where we find (lines from) these heavier elements is therefore related to how (far) gas has been blown out of galaxies.

The strongest absorption line we expect from this gas is from the ion O VII; this is an oxygen atom that has lost six electrons, and therefore only has two left. The ion O VIII, an oxygen atom that has lost all but one electron, also causes a strong absorption line, and the strongest emission line from the hot gas.

With the future space telescope Athena, we will be able to find relatively strong absorption lines from O VII and O VIII. These lines mostly come from relatively dense gas, with many elements formed in stars, and mostly from the gas around galaxies. These measurements are mostly sensitive to gas at temperatures optimal for those ions, even if most of the gas present has other temperatures. With the proposed Arcus and Lynx telescopes we should also be able to find weaker lines. These lines mostly come from more diffuse gas. In this diffuse gas, the effect of photons knocking electrons off ions is important, next to ionisation by electrons. This radiation contribution means that nuclei lose more electrons at a given temperature, meaning that a weak absorption line from a given ion can come from lower-temperature gas than stronger absorption lines from the same ion.

The strongest absorption lines we expect to find with these telescopes come from relatively dense gas with a relatively high oxygen content. That means that these lines mostly come from the CGM, around galaxies. Weaker lines can also come from CGM gas, but also from more diffuse gas in the IGM, like that of the filaments of the cosmic web. Athena is designed, in part, to look for X-ray absorption lines by observing a planned set of X-ray sources.

The telescope will be sensitive enough to find an O VII absorption line in most X-ray sources, if the halo of a galaxy like our Milky Way is between us and the source. Although the fraction of elements like oxygen that are created in stars decreases in the CGM as we get further from the central galaxy, the fraction of these elements in the gas causing X-ray absorption lines barely decreases. That is because the oxygen lines are only strong enough to observe where there happens to be a lot of oxygen. Those places are rarer as we look further from a galaxy, where the oxygen originates. However, we can only observe those rare, strong lines, meaning we can get a wrong impression of what is typical for gas far from galaxies.

Emission lines mainly come from gas at relatively high temperatures and densities. This means emission lines mainly come from the inner parts of the CGM, where gas is much denser and somewhat hotter than at the edges. With Athena, we will be able to see almost the entire CGM in a few emission lines in a few haloes, especially those that are (much) more massive than that of our Milky Way. Around galaxies like ours, we should be able to find a few emission lines from the inner CGM.

I have also compared a few measured absorption lines to the predictions from the EAGLE simulations. Those measurements are somewhat uncertain, because X-ray absorption lines are difficult to measure with current instruments. Another limitation is that I have compared the simulation to only four gas clouds. This makes comparisons difficult: in a simulation, you will never find the exact same gas cloud as in an observation. What we can do is see if the number of clouds we observe is similar to the prediction from a simulation, and whether the properties of those clouds match those in the simulations. These population comparisons are difficult with few measurements, because you do not know whether an observed cloud is typical or just happened to be odd.

Given the strengths of the absorption lines, the temperatures of two measured gas clouds seem to correspond reasonably with the EAGLE predictions. At first glance, the number of observed clouds seemed to exceed the EAGLE predictions, but a closer examination showed that this discrepancy was not so severe: the observers searched for those lines in a number of different places, and did not find lines in many of them.

The gas producing strong X-ray absorption lines may lie further from galaxies than predicted in EAGLE, but that is based on only two measured clouds. The detection of one of those clouds is uncertain. Therefore, it is currently difficult to say whether this is a problem for the EAGLE model, or whether one or two clouds happened to be a bit further from galaxies than typical. In the future, when we have more sensitive instruments, we will be able to make more measurements like these. It will be interesting to compare them to simulations in the same way. These comparisons can teach us how far elements created in stars are blown away by feedback.

All in all, the prospects for observations of hot gas around galaxies with future telescopes are good. Using X-ray emission lines, we will learn much about the gas around galaxies more massive than the Milky Way. It will be difficult to find emission lines from galaxies roughly as massive as the Milky Way, and from gas in the filaments of the cosmic web. With absorption lines, we will also be able to observe this gas. Using these observations, we will learn more about gas flows around galaxies, and how they affect the formation of stars.

Publications

First-author papers

Nastasha A. Wijers, Joop Schaye, & Benjamin D. Oppenheimer. *The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray and UV line absorption.*

2020, MNRAS, 498, 574–598

Nastasha A. Wijers, Joop Schaye, Benjamin D. Oppenheimer, Robert A. Crain, & Fabrizio Nicastro. *The abundance and physical properties of O VII and O VIII X-ray absorption systems in the EAGLE simulations.*

2019, MNRAS, 488, 2947–2969

Co-authored papers

J. Dorigo Jones, S. D. Johnson, Sowgat Muzahid, J. Charlton, H.-W. Chen, A. Narayanan, Sameer, J. Schaye, & **N. A. Wijers**. *Improving blazar redshift constraints with the edge of the Ly α forest: 1ES 1553+113 and implications for observations of the WHIM.*

2022, MNRAS, 509, 4330–4343

Jussi Ahoranta, Alexis Finoguenov, Massimiliano Bonamente, Evan Tilton, **Nastasha Wijers**, Sowgat Muzahid, & Joop Schaye. *Discovery of a multiphase O VI and O VII absorber in the circumgalactic/intergalactic transition region.*

2021, A&A, 656, A107

Adam J. Batten, Alan R. Duffy, **Nastasha A. Wijers**, Vivek Gupta, Chris Flynn, Joop Schaye, & Emma Ryan-Weber. *The cosmic dispersion measure in the EAGLE simulations.*

2021, MNRAS, 505, 5356–5369

Aurora Simionescu, Stefano Ettori, Norbert Werner, Daisuke Nagai, Franco Vazza, Hiroki Akamatsu, Ciro Pinto, Jelle de Plaa, **Nastasha Wijers**, Dylan Nelson, Etienne Pointecouteau, Gabriel W. Pratt, Daniele Spiga, Giuseppe Vacanti, Erwin Lau, Mariachiara Rossetti, Fabio Gastaldello, Veronica Biffi, Esra Bulbul, Maximilien J. Collon, and Jan-Willem den Herder, Dominique Eckert, Filippo Fraternali, Beatriz Mingo, Giovanni Pareschi, Gabriele Pezzulli, Thomas H. Reiprich, Joop Schaye, Stephen A. Walker, & Jessica Werk. *Voyage through the hidden physics of the cosmic web.*

2021, Experimental Astronomy

F. Nicastro, J. Kaastra, C. Argiroffi, E. Behar, S. Bianchi, F. Bocchino, S. Borgani, G. Branduardi-Raymont, J. Bregman, E. Churazov, M. Diaz-Trigo, C. Done, J. Drake, T. Fang, N. Grosso, A. Luminari, M. Mehdipour, F. Paerels, E. Piconcelli, C. Pinto, D. Porquet, J. Reeves, J. Schaye, S. Sciortino, R. Smith, D. Spiga, R. Tomaru, F. Tombesi, **N. Wijers**, & L. Zappacosta. *The Voyage of Metals in the Universe from Cosmological to Planetary Scales: the need for a Very High-Resolution, High Throughput Soft X-ray Spectrometer*. 2021, *Experimental Astronomy*

T. Tuominen, J. Nevalainen, E. Tempel, T. Kuutma, **N. Wijers**, J. Schaye, P. Heinämäki, M. Bonamente, & P. Ganeshaiah Veena. *An EAGLE view of the missing baryons*. 2021, *A&A*, 646, A156

Sarah Walsh, Sheila McBreen, Antonio Martin-Carrillo, Thomas Dauser, **Nastasha Wijers**, Jörn Wilms, Joop Schaye, & Didier Barret. *Detection capabilities of the Athena X-IFU for the warm-hot intergalactic medium using gamma-ray burst X-ray afterglows*. 2020, *A&A*, 642, A24

Benjamin D. Oppenheimer, Ákos Bogdán, Robert A. Crain, John A. ZuHone, William R. Forman, Joop Schaye, **Nastasha A. Wijers**, Jonathan J. Davies, Christine Jones, Ralph P. Kraft, & Vittorio Ghirardini. *EAGLE and Illustris-TNG Predictions for Resolved eROSITA X-Ray Observations of the Circumgalactic Medium around Normal Galaxies*. 2020, *ApJ*, 893, L24

Jussi Ahoranta, Jukka Nevalainen, **Nastasha Wijers**, Alexis Finoguenov, Massimiliano Bonamente, Elmo Tempel, Evan Tilton, Joop Schaye, Jelle Kaastra, & Ghassem Gozaliasl. *Hot WHIM counterparts of FUV O VI absorbers: Evidence in the line-of-sight towards quasar 3C 273*. 2020, *A&A*, 634, A106

Benjamin D. Oppenheimer, Jonathan J. Davies, Robert A. Crain, **Nastasha A. Wijers**, Joop Schaye, Jessica K. Werk, Joseph N. Burchett, James W. Trayford, & Ryan Horton. *Feedback from supermassive black holes transforms centrals into passive galaxies by ejecting circumgalactic gas*. 2020, *MNRAS*, 491, 2939–2952

Sean D. Johnson, John S. Mulchaey, Hsiao-Wen Chen, **Nastasha A. Wijers**, Thomas Connor, Sowgat Muzahid, Joop Schaye, Renyue Cen, Scott G. Carlsten, Jane Charlton, Maria R. Drout, Andy D. Goulding, Terese T. Hansen, & Gregory L. Walth. *The Physical Origins of the Identified and Still Missing Components of the Warm-Hot Intergalactic Medium: Insights from Deep Surveys in the Field of Blazar 1ES1553+113*. 2019, *ApJ*, 884, L31

Deborah Lokhorst, Roberto Abraham, Pieter van Dokkum, **Nastasha Wijers**, Joop Schaye, John A. ZuHone, Akos Bogdan, Ralph Kraft, & Alexey Vikhlinin. *On the Detectability of Visible-wavelength Line Emission from the Local Circumgalactic and Intergalactic Medium*. 2019, *ApJ*, 877, 4

Benjamin Oppenheimer, Juna Kollmeier, Andrey Kravtsov, Joel Bregman, Daniel Angles-

Alcazar, Robert Crain, Romeel Dave, Lars Hernquist, Cameron Hummels, Joop Schaye, Grant Tremblay, G. Mark Voit, Rainer Weinberger, Jessica Werk, **Nastasha Wijers**, John A. ZuHone, Akos Bogdan, Ralph Kraft, & Alexey Vikhlinin. *Imprint of Drivers of Galaxy Formation in the Circumgalactic Medium*.

2019, BAAS, 51, 280

Joel Bregman, Edmund Hodges-Kluck, Benjamin D. Oppenheimer, **Nastasha Wijers**, Laura Brenneman, Juna Kollmeier, Jelle Kaastra, Jiangtao Li, Eric D. Miller, Andrew Ptak, Randall Smith, Pasquale Temi, Feryal Ozel, & Alexey Vikhlinin. *A Survey of Hot Gas in the Universe*.

2019, BAAS, 51, 450

F. Nicastro, J. Kaastra, Y. Krongold, S. Borgani, E. Branchini, R. Cen, M. Dadina, C. W. Danforth, M. Elvis, F. Fiore, A. Gupta, S. Mathur, D. Mayya, F. Paerels, L. Piro, D. Rosa-Gonzalez, J. Schaye, J. M. Shull, J. Torres-Zafra, **N. Wijers**, & L. Zappacosta. *Observations of the missing baryons in the warm-hot intergalactic medium*.

2018, Nature, 558, 406–409

Preprints

Adam J. Batten, Alan R. Duffy, Chris Flynn, Vivek Gupta, Emma Ryan-Weber, & **Nastasha Wijers**. *Fast radio bursts as probes of feedback from active galactic nuclei*.

2021, submitted to MNRAS Letters, arXiv:2109.13472

Nastasha A. Wijers, & Joop Schaye. *The warm-hot circumgalactic medium around EAGLE-simulation galaxies and its detection prospects with X-ray line emission*.

2021, submitted to MNRAS, arXiv:2108.04847

Curriculum Vitae

My path to astronomy was relatively straightforward: my father is an astronomer, and I grew up stargazing with him on clear nights on vacations far from city lights. My father taught me all sorts of interesting things about astronomy, and my mother taught me important principles of research, such as the basics of internet research, and how to write in my own words. Both my parents encouraged me to be curious and to keep asking questions. This has helped me on my path towards astronomy, and has helped me keep learning during my PhD.

In school I was interested in most subjects, but I gravitated towards the physical sciences. I chose the ‘natuur en techniek’ (nature and technology) set of subjects in high school, and for my final project, a friend and I did actual observations at the University of Amsterdam! Unfortunately, these data were not usable, but we did analyse some data others took, to try to find the age of a globular cluster. In the last two years of high school, I also participated in a number of academic olympiads. I participated in the finals for astronomy (third place) and geography in 2010, and for astronomy (first place) and physics (second place, and a honourable mention in the international olympiad) in 2011. These olympiads were a first taste of these subjects as I would learn them at university, and I met many great people there. In 2011, I graduated high school summa cum laude.

Despite this early interest in astronomy, I was not sure whether I wanted to study (astro)physics or mathematics, so I did both. The University of Amsterdam offered a double BSc program for this. I did an end-of-year project in astronomy in my first year, mathematics in my second, and astronomy again in my third. I also participated in two International Astronomical Youth Camps during my BSc, which I enjoyed very much. I met a lot of great people there, and during my first IAYC in 2011, I got my first taste of programming, with a three-body simulation in C. In 2014, I obtained my BSc degrees in (astro)physics and mathematics, cum laude and with honours.

By the end of my BSc, I had decided I wanted to continue in astrophysics or particle physics rather than in mathematics. My interests at this point were mainly in theoretical and high-energy (astro)physics. I received a scholarship from Clare College to attend Cambridge University for the one-year Part III Maths program. Here, I followed courses in particle physics, black holes, and cosmology. I obtained an MSc degree here, with a distinction, but I wanted to gain some research experience and get a sense of the research being done before starting a PhD.

With exemptions for most of the coursework based on the courses I had taken in Cambridge, and another scholarship, I started an MSc research project with Shin’ichiro Ando at the University of Amsterdam. I worked on an analytical model for the distribution of dark matter in the universe, with the goal of understanding the effect of structure within dark matter haloes (subhaloes) on the strength of possibly detectable signals from dark matter. I enjoyed the project and learned a lot from it. However, it is highly uncertain when, if ever, signals like the one I was working on would be found. I decided to apply for PhD projects

focussed on the gas in the cosmic web and galaxy formation.

I chose a PhD project with Joop Schaye in Leiden, working on numerical simulations of galaxy formation. I focussed on hot, diffuse gas that had been difficult to detect. My research has shown that this gas may be detected within a decade! I have started a postdoctoral fellowship at CIERA (Northwestern University). Here, I will continue to work on diffuse gas, focussing on the physical processes that shape this gas, and how they connect to different observables.

Acknowledgements

There are many people who have helped me throughout my education and PhD, and who have helped me become the person I am today. First, I would like to thank my parents, Astrid and Ralph, for all the love and support they have given me throughout my life, and in particular during this strange and isolating pandemic year. Thank you for letting me stay with you during these difficult times. You have encouraged my curiosity throughout my life, and keep reminding me to stop and enjoy the view every once in while. I aim to take that advice one of these days.

Paula, Lydia, and Babs, you are great friends and have helped me face many of the challenges I have met during my PhD. I will miss our evenings that always last longer than planned. Paula, I will miss your hugs, and Lydia, I will miss our treks through Leiden. I will visit you when I'm in the Netherlands, and I hope you can visit me in Evanston sometime. I am very grateful for the technology that will let us keep in contact when I'm half a world away. You have helped me in so many ways. I am a better person since I met you.

Joop, you have taught me so much these past five years. Besides the research skills needed for a PhD, you have helped me find many collaborators who have expanded my knowledge and connected me to the astronomical community, and you have taught me a lot about critically evaluating ongoing research. I have enjoyed the opportunity to learn from you, and from the others within the community you connected me with.

Thank you to the computational galaxy formation group here in Leiden for welcoming me, and for many fun evenings of bouldering. You helped me get started on my adventure in Leiden. Your code and instructions were the starting point for my work. Thank you for letting me barge into your offices with all my questions.

To Mariska and Xander, thank you for helping me keep sane through all the pressures and uncertainty that come with research and my own insecurities. You have helped me grow as a person, and I know I am much better prepared for the future because of you.

Thanks to the Sterrewacht IT team for keeping everything running smoothly, and for helping me install, and in one case, debug, obscure astronomical software with strange error messages. Thank you to the the Secretariat team for helping me with all sorts of practical and administrative issues that I couldn't seem to figure out.

To the Leiden Observatory community: thank you for being open and welcoming, for organizing all sorts of movie nights, barbecues, and borrels to bring us together, and for caring about the wider issues of the world.

