



Publication Year	2024
Acceptance in OA	2024-12-11T17:53:55Z
Title	IGM damping wing constraints on the tail end of reionization from the enlarged XQR-30 sample
Authors	Greig, B., Mesinger, A., Bañados, E., Becker, G. D., Bosman, S. E. I., Chen, H., Davies, F. B., D'ODORICO, Valentina, Eilers, A. -C., Gallerani, S., Haehnelt, M. G., Keating, L., Lai, S., Qin, Y., Ryan-Weber, E., Satyavolu, S., Wang, F., Yang, J., Zhu, Y.
Publisher's version (DOI)	10.1093/mnras/stae1080
Handle	http://hdl.handle.net/20.500.12386/35460
Journal	MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY
Volume	530

IGM damping wing constraints on the tail end of reionization from the enlarged XQR-30 sample

B. Greig^{1,2,3★}, A. Mesinger^{1,4}, E. Bañados⁵, G. D. Becker⁶, S. E. I. Bosman^{1,5,7}, H. Chen^{1,8}, F. B. Davies^{1,5}, V. D’Odorico^{1,4,9,10}, A.-C. Eilers^{1,11}, S. Gallerani⁴, M. G. Haehnelt^{12,13}, L. Keating^{1,14}, S. Lai^{1,15}, Y. Qin^{1,3}, E. Ryan-Weber^{3,16}, S. Satyavolu¹⁷, F. Wang¹⁸, J. Yang¹⁸ and Y. Zhu^{1,7,18}

¹*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*

²*Research School of Astronomy & Astrophysics, Australian National University, Canberra, ACT 2611, Australia*

³*ARC Centre of Excellence for All-Sky Astrophysics in 3 Dimensions (ASTRO 3D)*

⁴*Scuola Normale Superiore, Piazza dei Cavalieri, I-56126 Pisa, Italy*

⁵*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

⁶*Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA*

⁷*Institute for Theoretical Physics, Heidelberg University, Philosophenweg 12, D-69120 Heidelberg, Germany*

⁸*Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St George St, Toronto, ON M5R 2M8, Canada*

⁹*INAF – Osservatorio Astronomico di Trieste, Via G.B. Tiepolo, 11, I-34143, Trieste, Italy*

¹⁰*IFPU – Institute for Fundamental Physics of the Universe, via Beirut 2, I-34151 Trieste, Italy*

¹¹*MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

¹²*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

¹³*Kavli Institute of Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

¹⁴*Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK*

¹⁵*Commonwealth Scientific and Industrial Research Organisation (CSIRO), Space & Astronomy, PO Box 1130, Bentley, WA 6102, Australia*

¹⁶*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, Victoria 3122, Australia*

¹⁷*Tata Institute of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India*

¹⁸*Steward Observatory, University of Arizona, 933 N Cherry Ave, Tucson, AZ 85721, USA*

Accepted 2024 April 18. Received 2024 April 17; in original form 2024 April 1

ABSTRACT

The attenuation of Ly α photons by neutral hydrogen in the intergalactic medium (IGM) at $z \gtrsim 5$ continues to be a powerful probe for studying the epoch of reionization. Given a framework to estimate the intrinsic (true) Ly α emission of high- z sources, one can infer the ionization state of the IGM during reionization. In this work, we use the enlarged XQR-30 sample of 42 high-resolution and high signal-to-noise quasar spectra between $5.8 \lesssim z \lesssim 6.6$ obtained with VLT/X-shooter to place constraints on the IGM neutral fraction. This is achieved using our existing Bayesian QSO reconstruction framework which accounts for uncertainties such as the: (i) posterior distribution of predicted intrinsic Ly α emission profiles (obtained via covariance matrix reconstruction of the Ly α and N V emission lines from unattenuated high-ionization emission line profiles; C IV, Si IV + O IV], and C III]) and (ii) distribution of ionized regions within the IGM using synthetic damping wing profiles drawn from a 1.6^3 Gpc^3 reionization simulation. Following careful quality control, we used 23 of the 42 available QSOs to obtain constraints/limits on the IGM neutral fraction during the tail-end of reionization. Our median and 68th percentile constraints on the IGM neutral fraction are: $0.20^{+0.14}_{-0.12}$ and $0.29^{+0.14}_{-0.13}$ at $z = 6.15$ and 6.35 . Further, we also report 68th percentile upper limits of $\bar{x}_{\text{HI}} < 0.21, 0.20, 0.21$, and 0.18 at $z = 5.8, 5.95, 6.05$, and 6.55 . These results imply reionization is still ongoing at $5.8 \lesssim z \lesssim 6.55$, consistent with previous results from XQR-30 (dark fraction and Ly α forest) along with other observational probes considered in the literature.

Key words: intergalactic medium – quasars: emission lines – dark ages, reionization, first stars – early Universe – cosmology: observations – cosmology: theory.

1 INTRODUCTION

The epoch of reionization (EoR) denotes the final major baryonic phase change of the Universe, when the pervasive cold and dense

neutral hydrogen fog in the intergalactic medium (IGM) is evaporated by the intense glow of ultraviolet (UV) ionizing radiation from the first sources. The morphological evolution of the IGM proceeds in a patchy manner whereby individual ionized (H II) regions form first around the earliest structures before growing and merging with their nearest neighbours, percolating and eventually ionizing almost all of the neutral hydrogen in the IGM. Observing and subsequently

★ E-mail: greigb@unimelb.edu.au

characterizing the EoR is of fundamental importance as it will advance our knowledge of the nature of the first structures to form in the Universe.

Our most promising observational tool for probing the EoR is the detection of the 21-cm hyperfine spin-flip transition of neutral hydrogen via radio interferometry. However, this continues to remain elusive with the best-yet upper limits on the 21-cm signal (Mertens et al. 2020; Trott et al. 2020; Abdurashidova et al. 2022a) still a few orders of magnitude away from our theoretical expectations (Ghara et al. 2020, 2021; Greig, Mesinger & Koopmans 2020; Mondal et al. 2020; Greig et al. 2021; Abdurashidova et al. 2022b; HERA Collaboration 2023).

In the absence of directly measuring the neutral hydrogen in the IGM, we can instead infer its presence via its impact on the Ly α photons emitted by bright astrophysical sources such as galaxies and QSOs. At lower redshifts, one can directly infer the ionization state of the IGM by measuring the absorption of Ly α photons within the Ly α forest along the line of sight from the source. However, at $z \gtrsim 5$ the scattering cross-section of Ly α photons is sufficiently large that even trace amounts of neutral hydrogen ($x_{\text{H I}} \gtrsim 10^{-4}$) are enough to produce completely saturated transmission (Fan et al. 2006). As a result, fluctuations in the neutral fraction owing to the patchy nature of reionization are almost indistinguishable from fluctuations post-reionization of the UV background, density or temperature (e.g. D’Aloisio, McQuinn & Trac 2015; Keating, Puchwein & Haehnelt 2018). However, physical models predict different large-scale fluctuations and biases of these fields, which can be used to constrain the EoR history and galaxy properties (see e.g. fig. 2 in Qin et al. 2021).

Alternatively, a more robust probe of the IGM neutral fraction is to measure the imprint of the Ly α damping wing (e.g. Rybicki & Lightman 1979; Miralda-Escudé 1998). This approach takes advantage of the Lorentzian wings away from the resonant central core of the Ly α absorption cross-section. In these wings, the amplitude of the scattering cross-section is several orders of magnitude lower and imprints a smooth absorption profile away from the central saturated absorption. Thus, even a completely neutral IGM will impart a smooth, measurable imprint in the transmission spectrum of the background source sufficiently redward from the line centre.

This IGM damping wing signature has been successfully measured to obtain constraints on the IGM neutral fraction using both galaxies and QSOs. For QSOs, their intrinsic brightness allows the damping wing imprint to be inferred from individual spectra (e.g. Mesinger & Haiman 2007; Bolton et al. 2011; Mortlock et al. 2011; Greig et al. 2017b; Bañados et al. 2018; Davies et al. 2018a; Wang et al. 2020; Yang et al. 2020). On the other hand, for galaxies, the extraction of the damping wing signal requires averaging over a sufficiently large statistical sample owing to the fainter nature of galaxies (e.g. Mesinger et al. 2015; Mason et al. 2018, 2019; Hoag et al. 2019; Umeda et al. 2023), however, it can be more difficult to distinguish the IGM component from the host galaxy environment than is the case for QSOs (e.g. Heintz et al. 2023; Keating et al. 2023, 2024).

Importantly, the key requirement for extracting the IGM damping wing imprint is a robust method to predict the intrinsic Ly α emission, given we inherently observe an attenuated spectrum. For QSOs, numerous methods exist in the literature which differ considerably in their methodology (e.g. Davies et al. 2018b; Fathivavari 2020; Reiman et al. 2020; Āurovčíková et al. 2020; Bosman et al. 2021; Liu & Bordoloi 2021; Chen et al. 2022; Sun, Ting & Cai 2023) but fundamentally boil down to predicting the Ly α profile from unattenuated spectral information redward of Ly α . In Greig et al. (2024), a detailed comparison of the various reconstruction pipelines

in the literature, focusing on the reconstruction of the Ly α line profile, are performed as part of a blind reconstruction challenge. This extends on the work of Bosman et al. (2021), comparing methods to predict the placement of the Ly α forest continuum.

In this work, we use the Bayesian reconstruction pipeline initially introduced in Greig et al. (2017a) which predicts the intrinsic Ly α profile by drawing from an N -dimensional normal distribution with a covariance matrix containing the correlations between the Ly α and N V line and the unattenuated high-ionization emission lines (C IV, Si IV + O IV], and C III]). This approach assumes that the emission lines can be modelled as either a single or double component Gaussian profile characterized entirely by the width, height, and velocity offset. We then couple this reconstruction framework to distributions of synthetic IGM damping wings drawn from a large-volume EoR simulation (Mesinger, Greig & Sobacchi 2016). In doing so, we account for the statistical uncertainties that arise since we are using single lines of sight to each QSO spectra to place constraints on the volume averaged IGM neutral fraction during the EoR.

Previously, IGM damping wing analyses have focused solely on the highest redshift QSOs to be discovered at $z \gtrsim 7$ (Mortlock et al. 2011; Bañados et al. 2018; Wang et al. 2020; Yang et al. 2020).¹ However, in this work we take advantage of the Ultimate X-shooter legacy survey of QSOs at $5.8 \lesssim z \lesssim 6.6$, XQR-30 (D’Odorico et al. 2023). This programme obtained 30 high signal-to-noise (S/N) QSO spectra at high resolution ($\sim 30 \text{ km s}^{-1}$). These were supplemented by a further 12 high-quality $z \sim 6$ QSOs from the X-shooter archive to produce a final sample of 42 QSOs. The sheer size of this QSO sample and the corresponding redshift coverage will enable unique constraints on the tail end of the EoR to be recovered.

This work is structured as follows. In Section 2, we provide a brief description of the observational data, and in Section 3, we describe our full analysis pipeline. In Section 4, we provide the main results along with detailed discussions. In Section 5, we conclude with our closing remarks. Unless stated otherwise, we quote all quantities in co-moving units and adopt the cosmological parameters: $(\Omega_{\Lambda}, \Omega_{\text{M}}, \Omega_b, n, \sigma_8, H_0) = (0.69, 0.31, 0.048, 0.97, 0.81, 68 \text{ km s}^{-1} \text{ Mpc}^{-1})$, consistent with recent results from the *Planck* mission (Planck Collaboration XIII 2016).

2 DATA

2.1 Sample

The data used in this work come from the XQR-30 sample (D’Odorico et al. 2023), which is an ESO (European Southern Observatory) Large Programme ‘XQR-30: the Ultimate X-shooter legacy survey of quasars at $z \simeq 5.8\text{--}6.6$ arcsec (PI: V. D’Odorico; 1103.A-0817). Targeting 30 of the brightest known QSOs within this redshift range with X-shooter (Vernet et al. 2011), we obtained a median resolution of $R \sim 9500\text{--}13\,700$ in the visible arm and $R \sim 7600\text{--}11\,000$ in the near-infrared arms. For this work, we rebinned the observed QSO spectra to 50 km s^{-1} . In addition to these 30 QSOs, a further 12 QSOs with similar redshift range, magnitude, resolution, and S/N to the original XQR-30 sample were added from the X-shooter archive to yield the enlarged XQR-30 sample. This entire sample was reduced using a custom pipeline (López et al. 2016; Becker et al. 2019), with minor modifications made for XQR-30 (see D’Odorico et al. 2023, for further details).

¹Although, see Āurovčíková et al. (2024) who very recently performed an analysis of a much larger sample of 18 QSOs from $6.0 < z < 7.1$.

Table 1. A summary of the enlarged XQR-30 data set. Columns correspond to the QSO name, redshift, emission line used for obtaining the redshift, recovered IGM neutral fraction (68 percentiles, either represented as an upper limit or constraint) and brief comments highlighting the reason for each QSOs exclusion or retention (see the main text for further details).

QSO name	Redshift (z)	Redshift method	$\bar{x}_{\text{H I}}$	Comments
PSO J029 – 29	5.976	Mg II	<0.22	
VST-ATLAS J029 – 36	6.013	Mg II	<0.40	
VDES J0224 – 4711	6.525	Mg II	<0.33	
PSO J060 + 24	6.170	Mg II	$0.29^{+0.24}_{-0.19}$	
VDES J0408 – 5632	6.033	Mg II	$0.24^{+0.19}_{-0.15}$	Mini BAL or weak BAL ^a
PSO J108 + 08	5.955	Mg II	<0.31	
SDSS J0842 + 1218	6.0754	[C II]	<0.30	Mini BAL or weak BAL ^a
PSO J158 – 14	6.0685	[C II]	<0.32	
PSO J217 – 16	6.1498	[C II]	$0.29^{+0.20}_{-0.18}$	
PSO J231 – 20	6.5869	[C II]	<0.29	Mini BAL or weak BAL ^a
PSO J242 – 12	5.840	Mg II	<0.25	
PSO J308 – 27	5.799	Mg II	$0.33^{+0.22}_{-0.20}$	
PSO J323 + 12	6.5872	[C II]	$0.23^{+0.19}_{-0.15}$	
PSO J359 – 06	6.1722	[C II]	<0.34	
SDSS J0100 + 2802	6.3269	[C II]	$0.41^{+0.19}_{-0.18}$	
VST-ATLAS J025 – 33	6.3373	[C II]	<0.33	
ULAS J0148 + 0600	5.977	Mg II	$0.32^{+0.23}_{-0.20}$	Mini BAL or weak BAL ^b
PSO J036 + 03	6.5405	[C II]	<0.31	
SDSS J0836 + 0054	5.773	Mg II	<0.33	
SDSS J0927 + 2001	5.7722	[C II]	<0.33	
SDSS J1030 + 0524	6.304	Mg II	<0.40	
SDSS J1306 + 0356	6.0330	[C II]	<0.36	
CFHQS J1509 – 1749	6.1225	[C II]	<0.38	
PSO J007 + 04	6.0015	[C II]	–	pDLA
PSO J009 – 10	6.0040	[C II]	–	BAL
PSO J023 – 02	5.817	Mg II	–	BAL, pDLA
PSO J025 – 11	5.816	Mg II	–	pDLA
PSO J065 – 26	6.1871	[C II]	–	pDLA
PSO J065 + 01	5.804	Mg II	–	Mini BAL or weak BAL
PSO J089 – 15	5.972	Mg II	–	BAL
J0923 + 0402	6.6330	[C II]	–	BAL
PSO J183 + 05	6.4386	[C II]	–	pDLA
PSO J183 – 12	5.893	Mg II	–	Mini BAL or weak BAL
PSO J217 – 07	6.166	Mg II	–	BAL
DELS J1535 + 1943	6.358	Mg II	–	Heavily reddened
PSO J239 – 07	6.1102	[C II]	–	BAL, pDLA
VDES J2211 – 3206	6.3394	[C II]	–	BAL
VDES J2250 – 5015	5.985	Mg II	–	BAL
SDSS J2310 + 1855	6.0031	[C II]	–	Mini BAL or weak BAL, pDLA
WISEA J0439 + 1634	6.5188	[C II]	–	BAL
SDSS J0818 + 1722	5.960	Mg II	–	pDLA
ULAS J1319 + 0950	6.1347	[C II]	–	pDLA

Note.^a Although identified as a possible mini/weak BAL by Bischetti et al. (2023), the associated BAL complex does not impact the Ly α + N V emission line complex. ^b A potential weak BAL complex however the large blueshift should not contaminate the Ly α + N V emission line complex (Becker et al. 2015).

2.2 Final damping wing sample

Unfortunately not all of the 42 QSOs from the enlarged XQR-30 sample can be used for IGM damping wing studies. For example, we must avoid QSOs which exhibit absorption by dense, neutral absorption systems in close proximity to the host QSO (e.g. damped Ly α systems; DLAs) which produce their own damping wing imprint which can imitate the damping wing signal from a partially neutral IGM (e.g. Prochaska, Hennawi & Herbert-Fort 2008; Bañados et al.

2019; Davies et al. 2023). Further, QSOs which exhibit broad absorption lines (BALs) must also be excluded as these lines either contaminate the observed flux around Ly α and N V or impact our ability to accurately measure the continuum and emission line properties of the QSOs redward of Ly α which are pre-requisite for our intrinsic Ly α reconstruction pipelines. In Table 1, we summarize the full XQR-30 sample, highlighting which QSOs are retained for our damping wing studies, and discuss the exclusions of certain objects in greater detail below.

Bischetti et al. (2022) performed an extensive C IV BAL analysis of the XQR-30 sample, with a follow up analysis exploring N V, Si IV + O IV], and Mg II BAL systems (Bischetti et al. 2023). Following this, 14 BAL QSOs were identified. However, of these, three (VDES J0408 – 5632, SDSS J0842+1218, and PSO J231 – 20) were deemed suitable for our damping wing analysis as the velocities of the BAL systems do not contaminate the Ly α + N V line complex. In D’Odorico et al. (in preparation) proximate DLA systems (pDLAs) have been identified within 7 of the XQR-30 QSOs, however only four of these are unique (i.e. are not also identified as BAL QSOs). We further exclude the heavily reddened QSO J1535 + 1943, which is thought to be obscured (Yang et al. 2021).

Of the additional 12 archival X-shooter spectra in the XQR-30 sample, two exhibit pDLA systems (J0818 + 1722 and J1319 + 0950). Further, J0439 + 1634 is known to exhibit a BAL system. J0148 + 0600 also likely exhibits a BAL system (Becker et al. 2015) however it is highly blueshifted and should not contaminate the Ly α + N V line complex.

In total, this results in 23 QSOs from the XQR-30 sample that we deem suitable for extracting the IGM damping wing imprint.

3 METHOD

Our damping wing analysis relies on the covariance matrix reconstruction pipeline introduced by Greig et al. (2017a) and subsequently extended in Greig et al. (2022) in order to simultaneously reconstruct the Ly α and N V emission lines. This is combined with a Bayesian sampling of synthetic IGM damping wing profiles drawn from large-volume EoR simulations to be able to infer the IGM neutral fraction. Below, we summarize the key components of our full analysis pipeline and refer the reader to the aforementioned works for more in-depth discussions.

3.1 Intrinsic Ly α profile reconstruction

In Greig et al. (2017a), we introduced a method to reconstruct the intrinsic Ly α line profile of QSOs using a covariance matrix of correlations between Ly α and other measurable high-ionization lines (namely, C IV, Si IV + O IV], and C III]). This covariance matrix was determined from a carefully selected (visually inspected) training set of 1673 moderate- z ($2.08 < z < 2.5$), high S/N (>15) QSOs from the Sloan Digital Sky Survey (SDSS-III) Baryon Oscillation Spectroscopic Survey (BOSS) DR12 (Dawson et al. 2013; Alam et al. 2015) to avoid the presence of BAL and DLA systems. In Greig et al. (2022), we expanded the emission line covariance matrix to allow the simultaneous reconstruction of Ly α and N V.

The underlying assumption of this pipeline is that each emission line can be modelled by a sum of Gaussian profiles, each fully described by its amplitude, width, and velocity offset away from systemic. In Greig et al. (2017a), we performed basic model selection to determine that the Ly α and C IV line profiles are best characterized by a two-component Gaussian, consisting of both a broad and narrow component. In addition to the Gaussian emission lines, we model the QSO continuum as a single two-parameter power law ($\propto \lambda^{\alpha_1}$) over our entire fitting region (1180–2300 Å) and normalize all QSO spectra to unity at 1450 Å. For all remaining emission lines, a single-component Gaussian profile was preferred.² This results in

a 21×21 covariance matrix of emission line parameters that is used to construct a 21D Normal distribution which describes the full properties of our five emission lines. The two parameters describing the power-law continuum do not enter the covariance matrix as they do not correlate with the parameters describing the emission lines.

3.1.1 Enlarged QSO training set

For this work we significantly expand our original training set of QSOs, updating our training set to use the BOSS DR16Q database (Lyke et al. 2020). We both expand the redshift range from our original $2.08 < z < 2.5$ to a broader $2.08 < z < 4.0$ and reduce our S/N cut down from $S/N > 15$ to > 6.5 . Owing to this expansion of the training set, we now forego visual inspection of the QSOs in our training set. The primary reason for expanding this training set is twofold: (i) to avoid biasing our training set based on the quality of QSO spectra (large amplitude Ly α peaks etc.) and (ii) to increase the sampled distribution of emission line parameters to ensure a larger region of validity for our covariance matrix.

After broadening our training set, we are left with $\sim 55\,000$ QSOs. We then reduce this sample further by removing QSOs with too many coincident absorption features³ around our emission lines which can make determining the true profile problematic. Thus, we restrict our sample to QSOs with fewer than 20 absorption features across our fitting range of 1180–2300 Å and fewer than 8 in the 1180–1260 Å region around Ly α and N V. While these choices are arbitrary, they strike a balance between providing a sufficient number of QSOs for evaluating our covariance matrix and limiting the number of outlier spectra. Note, in fitting these absorption features we do not distinguish by their origin (i.e. Ly α forest, intervening metal absorption and/or QSO associated absorption). Simply, from experience we find that for an increasing number of individual absorption lines within the 1180–1260 Å region, determining the intrinsic Ly α profile from the training QSOs becomes problematic as the Ly α components become strongly degenerate with these numerous absorption features. Beyond just the Ly α region, we additionally find that too many absorption lines can impact our ability to fit the two component power-law continuum, impacting all subsequent emission line fits. After these cuts, our final training set consists of 30 166 QSOs, corresponding over an order of magnitude increase from the 1673 QSOs used in Greig et al. (2017a). In Appendix A, we provide the correlation coefficient matrix demonstrating the various emission line correlation strengths. With such a large number of QSOs in the training set, it is unlikely that our resultant correlations will be strongly biased. Interestingly, despite substantially increasing the size of our QSO training set for this work, we find very similar correlations strengths between our emission lines as we have found previously (Greig et al. 2022). Although the correlation strengths remain similar, the notably larger training set demonstrates that these line correlations exist over broader ranges for the individual emission line parameter properties.

C II, He II, O III/Al II/Fe II, and Al III (see Greig et al. 2017a for further details). Further, we do not consider an emission component from an Fe continuum, however, for the purposes of our studies we find such a component unlikely to impact on our QSO fits.

³Absorption features are identified through an automated procedure based on a rolling 2 Å window, flagging all flux bins below 3σ (5σ around Ly α) and then roughly determining the absorber line centre which is assigned a Gaussian profile (see Greig et al. 2017a for further details). All absorption features are simultaneously fit to the QSO spectrum along with the emission lines, varying their width, amplitude and velocity offset.

²Note, we fit all observed emission lines that fall within this specific wavelength region, but only use the high-ionization lines for the covariance matrix. That is, we also simultaneously fit for the following lines Si II, O I/Si II,

3.1.2 Reconstruction pipeline for XQR-30

With our updated covariance matrix of emission line parameters, we now outline our method for reconstructing the intrinsic Ly α and N V emission lines. The key steps are as follows:

- (i) The observed rest-frame QSO spectrum is fit over $\lambda = 1275\text{--}2300\text{ \AA}$ using the systemic QSO redshifts⁴ outlined in Table 1.
- (ii) We jointly fit the two-parameter power-law continuum and the high-ionization lines Si IV + O IV], C IV (two-component) and C III] emission line profiles using our Gaussian profiles. Additionally, to improve the overall fit to the observed spectrum, we simultaneously allow a variable number of absorption lines, each modelled as a single Gaussian profile, to be fit.
- (iii) Using the corresponding fits to the C IV, Si IV + O IV], and C III] emission lines (amplitude, width, and velocity offset) of our observed high- z QSO spectrum, we use the maximum a -posteriori values for these lines and evaluate our 21D Normally distributed covariance matrix model. This allows us to reduce the full 21-parameter model down to a 9D estimate of just the joint Ly α and N V intrinsic emission line profile (two Gaussian components for Ly α and one for N V).
- (iv) We then draw joint intrinsic Ly α and N V profiles from this 9 dimensional covariance matrix. We further utilize the amplitude of the observed spectrum redward of $1250\text{ \AA}$ as an additional prior, to ensure the reconstructed profiles match the observed flux within this range.

3.2 Modelling the IGM damping wing during reionization

In order to extract the imprint of a partially neutral IGM, we compare the observed signal against our intrinsic profile reconstructions convolved with synthetic IGM damping wing profiles. Our synthetic IGM damping wing profiles are obtained from the Evolution of 21-cm Structure (Mesinger et al. 2016)⁵ 2016 simulations. These semi-numerical reionization simulations have side lengths of 1.6 Gpc over 1024 voxels and include state-of-the-art subgrid prescriptions for inhomogeneous recombinations and the photoheating suppression of star formation. Further, through judicious choices of star formation efficiency and the typical masses of haloes hosting star-forming galaxies, three distinct EoR morphologies are explored.

Following Greig et al. (2022), we only consider the INTERMEDIATE H II model, which corresponds to reionization driven by galaxies residing in $M_h \gtrsim 10^9 M_\odot$ haloes.⁶ This model is consistent with recent results from forward modelling the Ly α forest opacity distributions and UV galaxy luminosity functions (LFs; Qin et al. 2021).

We statistically account for the fact we only obtain a single observed QSO sightline by drawing a large sample of synthetic IGM

damping wing profiles varying both: (i) the sightline orientation (selecting 10 different lines of sight through the IGM centred on the host halo) and (ii) the diversity of host QSO environments in the IGM. In total, we draw 10^5 synthetic IGM damping wing profiles obtained from the 10^4 largest identified haloes (corresponding to roughly $6 \times 10^{11} < M_h < 3 \times 10^{12} M_\odot$)⁷ and 10 randomly oriented sightlines originating from the central host halo. To have these synthetic damping wing profiles varied by IGM neutral fraction we take the ionization fields obtained from different redshift snapshots from the simulation, which assumes the halo locations do not significantly change due to bulk motions across the different redshift snapshots. Following this, we obtain 24 unique, unevenly spaced values of the IGM neutral fraction, corresponding to the original redshift sampling of our seminumerical simulations (i.e. 24 different snapshots during the EoR), spanning from $\bar{x}_{\text{H I}} = 0.02$ to 0.95 for our analysis.

These synthetic damping wing profiles are constructed by assuming a minimum local H II bubble size, R_{min} , which for the $z > 7$ QSOs (Greig et al. 2017b; Greig et al. 2019; Greig et al. 2022) was adopted to be their corresponding proximity zone size. This is necessary since within our EoR simulations we do not model the impact that the QSO has on growing its own local H II bubble. The actual size of the local H II region is the combination of the cumulative ionizing output of the resident local galaxies (which is incorporated within our simulations) and the unknown QSO contribution, which is crudely implemented by assuming that at a minimum this QSO plus galaxy contribution has to have at least ionized the size of the measured QSO proximity zone. The contribution of the QSO to the local H II region is strongly dependent on its accretion history, obscuration and emission geometry all of which are highly uncertain and thus difficult to impose physically motivated priors. Therefore, in this work we only search for the IGM damping wing imprint between $\lambda = 1218\text{--}1230\text{ \AA}$ to specifically avoid having to simultaneously model the host QSO environment blueward of Ly α , thus minimizing additional modelling uncertainties. This choice makes our approach more agnostic to the QSO emission model and its associated priors. The lower bound of $1218\text{ \AA}$ ensures we are sufficiently far from the gravitational potential of the halo, avoiding potential signatures of inflowing gas. This approach differs to that of Davies et al. (2018a) who utilize numerical simulations to develop a hybrid scheme allowing for the search of the damping wing imprint over an extended region of $\lambda \sim 1190\text{--}1230\text{ \AA}$. This enables these QSO emission model uncertainties such as the QSO lifetime and environment to be folded into their analysis. A detailed discussion of the consequences of these assumptions are provided in Greig et al. (2022) but in summary our redward approach implicitly imposes a complex prior on the QSO lifetime, but should not strongly impact our results on the inferred neutral fraction. A quantitative estimate of the impact of this prior is beyond the scope of this work, as it requires a detailed comparison of the two different damping wing extraction methods.

Due to the significantly increased number of QSOs in the XQR-30 sample relative to our previous $z > 7$ QSO studies and the additional computational burden of having to regenerate these synthetic sightlines for each possible R_{min} , we instead choose to only use the existing synthetic damping wing profiles from the four previously analysed $z > 7$ QSOs. Therefore, for each XQR-30 QSO we simply average over our existing data, drawing 10^5 profiles for $R_{\text{min}}; \sim 11\text{--}16$

⁴Note, our BOSS QSO training set is converted to rest frame using the provided pipeline redshift, rather than the [C II] or Mg II redshifts of the XQR-30 sample. Although this may lead to biases in the reconstructions due to differences in the obtained red/blueshifts from systemic, this is likely a subdominant effect relative to the large scatter in the reconstructed Ly α + N V line and of the synthetic IGM damping wings (see Greig, Mesinger & Bañados 2019, for detailed discussions)

⁵<http://homepage.sns.it/mesinger/EOS.html>

⁶In Greig et al. (2019), we explored the impact that different EoR morphologies had on the recovered IGM neutral fraction constraints, finding only weak evidence for morphology dependence. The median IGM neutral fraction due to the different EoR morphologies was found to be only ± 0.05 relative to the typical 68th percentile uncertainty of $\pm 0.15\text{--}0.20$. This should not impact our results in any discernible way.

⁷Although the XQR-30 QSOs are known to be bright and powered by massive blackholes (Mazzucchelli et al. 2023), they are consistent with the range of halo masses considered within our seminumerical simulations and thus should not strongly impact on our inferred constraints on the IGM neutral fraction.




Figure 1. A visual demonstration of the initial covariance reconstruction for ATLAS J025 – 33 at $z = 6.34$ highlighting problems with our original reconstruction pipeline. Left panel: the observed QSO spectrum (black curve), corresponding ML (red curve) reconstruction and 300 random profile draws (grey curves) from our full reconstructed covariance matrix. The solid purple curve corresponds to the median reconstruction profile obtained from the full distribution of reconstructed profiles while the dashed (dotted) curves correspond to the 68th (95th) distribution of profiles. The vertical blue dashed line corresponds to rest-frame $\text{Ly}\alpha$, while the purple box denotes the region over which we fit for the IGM damping wing imprint (1218–1230 Å). Right panel: the observed $\text{Ly}\alpha$ transmission profile (black curve) obtained by dividing the observed spectrum by the median reconstructed profile. The coloured lines correspond to the median synthetic IGM damping wing profiles for each IGM neutral fraction value sampled by our EoR simulations. Grey-shaded regions correspond to absorption features that are masked by our damping wing fitting pipeline.

comoving Mpc.⁸ By averaging over these $R_{\min}$ we effectively broaden our resultant posteriors, leading to a more conservative estimate of the IGM neutral fraction. In Appendix B, we explore the impact of assuming different $R_{\min}$ for one of our QSOs in our sample. Between our lower and upper limits on this minimum radii we find the inferred IGM neutral fraction shifts by ~ 0.02 , which is considerably smaller than the width of the resultant posterior. Therefore, this choice of convolving over $R_{\min}$ should have little impact on our inferred results on the IGM neutral fraction.

3.3 Initial QSO reconstruction performance

Our initial attempts at performing intrinsic reconstructions following the procedure outlined in Section 3.1.2 for the XQR-30 sample proved to be problematic. As a visual demonstration, in the left panel of Fig. 1 we provide our initial attempt at reconstructing the intrinsic profile of ATLAS J025 – 33. The black curve corresponds to the observed spectrum, while the red line corresponds to the maximum-likelihood (ML) profile reconstruction. The grey lines correspond to 300 random posterior draws of our reconstructed profiles, while the purple solid line corresponds to the median profile and the dashed (dotted) purple lines demarcate the 68th and 95th percentile ranges determined over the full distribution of reconstructed profiles. Finally, the purple-shaded box corresponds to the interval over which we extract the smooth IGM damping wing component (1218–1230 Å).

In this example, the ML reconstruction (also the median, but to a lesser extent as it is obtained over the distribution instead of an individual profile) drops below the observed spectrum between

rest-frame 1210–1220 Å.⁹ Owing to the strong blueshift of this particular QSO, this corresponds to the flux between the $\text{Ly}\alpha$ and N v emission lines. Importantly, this ‘feature’ is not just restricted to one QSO, instead it is persistent across the full QSO sample, systematically occurring between the $\text{Ly}\alpha$ and N v emission lines. In the right panel of Fig. 1, we demonstrate the impact that this behaviour can have on our recovery of the smooth IGM damping wing imprint. Here, we show the $\text{Ly}\alpha$ transmission profile (black curve), obtained by dividing the observed spectrum (black curve, left panel) by the median reconstruction profile (purple curve, left panel). The coloured curves correspond to the median synthetic IGM damping wing profiles for each value of the IGM neutral fraction sampled by our dataset (increasing neutral fraction from top to bottom, see Section 3.2). The smooth, monotonically increasing synthetic damping wing profiles are incompatible with the decreasing amplitude of the observed $\text{Ly}\alpha$ transmission within the 1218–1230 Å region owing to this underestimate of the intrinsic flux. Attempting to jointly fit the observed $\text{Ly}\alpha$ transmission and synthetic damping wing profiles within this region was problematic for our pipeline.

3.4 Calibrating our QSO reconstruction pipeline

We initially assumed that the origin of this feature was the higher quality of the X-shooter sample relative to the training set of SDSS (BOSS) QSOs used in the construction of our covariance matrix. The higher resolution would then lead to more prominent emission line features (less blended), in particular around $\text{Ly}\alpha$ and N v driving this difference. Subsequently, we degraded the XQR-30 QSOs to better mimic the spectral properties of the QSOs within our training set but

⁸Note that this is slightly smaller than some of the proximity zone sizes measured for the XQR-30 sample by Satyavolu et al. (2023); 2–7 proper Mpc at $z = 5.8$ –6.6. However, unfortunately we no longer have the base simulation data to extend our upper limit of $R_{\min}$. Nevertheless, this should not significantly impact our results due to the broad range of modelling uncertainties.

⁹The reconstructed profile also drops below the observed spectrum at > 1235 Å, however, this is less prevalent across our sample and occurs outside of the region over which we fit for the damping wing (i.e. 1218–1230 Å). Nevertheless, this likely is related to the issues with the width and placement of the N v line from our reconstructed profile.




Figure 2. The performance of our covariance matrix reconstruction method from two distinct samples of 30 QSOs from X-shooter (left panel) and SDSS (BOSS; right panel) at $3.5 < z < 4.5$. The red solid lines correspond to the flux ratio (reconstructed ML profile over true intrinsic profile) for each individual QSO in the samples. The black solid lines correspond to the median determined over each sample.

found that this behaviour of a reduction of flux between Ly α and N v persisted.

Therefore, in order to establish the origin of this feature in our profile predictions we performed a reconstruction of QSOs drawn from two different samples both at $3.5 < z < 4.5$ which should be unaffected by IGM attenuation (i.e. where we can easily establish the true intrinsic profile shape). For this, we constructed two samples of 30 QSOs each, one observed with X-shooter to obtain QSOs of similar quality and properties of the XQR-30 sample and one with SDSS (BOSS DR16Q)¹⁰ to be representative of the quality of our training set spectra. For the former, we utilize the XQ100 sample (López et al. 2016). Importantly, by considering two samples of QSOs from distinctly different instruments, we can determine whether: (i) the problems in our reconstruction are intrinsic to the quality of the spectrum and (ii) whether it is acceptable for a reconstruction method trained on lower quality spectra to be applied to any instrument.

Originally, the scope of this investigation was to be limited to just our covariance matrix reconstruction pipeline. However, it quickly morphed into a robust and detailed comparison of all available reconstruction pipelines in the literature. For detailed discussions resulting from this comparison, we refer the reader to Greig et al. (2024). In short, the issue was identified to be intrinsic to our covariance reconstruction method (see below) and that the quality of the spectra in the training set (e.g. resolution or S/N) relative to the observed spectra did not impact the reconstruction performance. In the remainder of this section, we focus only on the results intrinsic to our covariance matrix reconstruction pipeline, relevant for this work.

In Fig. 2, we summarize the performance of our covariance matrix reconstruction pipeline across the two distinct samples of 30 QSOs. The left panel corresponds to the X-shooter sample, while the right panel corresponds to the SDSS (BOSS) sample. Here, the red curves correspond to the flux ratios (reconstructed ML profile divided by true, unattenuated flux) for each individual QSO in each sample. The

black solid line corresponds to the median flux ratio obtained over the 30 QSOs in each sample. For both, we follow the reconstruction procedure outlined in Section 3.1.2.

Immediately from Fig. 2, it is evident that the flux ratios across both samples follow very similar behaviour, exemplified by the remarkably similar shape in the median flux ratios, albeit with the SDSS sample exhibiting a slightly larger offset. This highlights that the higher quality X-shooter spectra in comparison to the BOSS training set does not produce any additional source of systematic biases. Thus explaining why the degradation of the XQR-30 spectra did not alleviate the issue. Further, for both samples a clear dip is evident between 1215 and 1230 Å.¹¹ That is, our reconstruction pipeline consistently underestimates the shape of the reconstructed Ly α + N v line profile within this region by ~ 10 percent irrespective of the quality of the observed spectra being reconstructed. Thus, the problematic feature identified earlier is systematic to our reconstruction pipeline.

The origin of this behaviour is due to the modelling of the QSO flux within this specific wavelength region. Here, the QSO flux is the sum of three Gaussian components,¹² the sum of a broad and narrow component for Ly α and a single component for N v. Deviations in the predicted locations and profile shapes of these Gaussian components can produce this decrement in the QSO flux between Ly α and N v. For example, a larger predicted separation between the broad Ly α and N v lines caused by either (or both) a bluer than expected Ly α line centre or redward N v line centre or a narrower than expected width in these line profiles would produce a flux decrement within this spectral region. Owing to the notably lower strength of the N v emission line correlations (compared to the Ly α components) with the other emission lines in our covariance matrix (see e.g. Fig. A1) the N v is going to be less strongly constrained (i.e. larger uncertainty)

¹⁰Although some QSOs from our expanded training set likely exist in this sample of 30 BOSS QSOs, since our reconstruction method relies on drawing from a Normal distribution based on a covariance matrix of line properties our reconstructions will not be biased by their appearance.

¹¹Note that this feature occurs over a different rest-frame wavelength range purely due to the blueshift of ATLAS J025 – 33 used in our previous example.

¹²Although in theory we could consider additional Gaussian components to improve the ability to fit the Ly α + N v line complex, the overall correlation strengths between additional (beyond two) Ly α components and the other high-ionization lines would diminish as additional components are not always required to fit the observed training data.