

Lawrence Berkeley National Laboratory

LBL Publications

Title

New Measurements of the Ly α Forest Continuum and Effective Optical Depth with LyCAN and DESI Y1 Data

Permalink

<https://escholarship.org/uc/item/06s3333t>

Journal

The Astrophysical Journal, 976(1)

ISSN

0004-637X

Authors

Turner, Wynne
Martini, Paul
Karaçaylı, Naim Göksel
[et al.](#)

Publication Date

2024-11-01

DOI

10.3847/1538-4357/ad8239

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



New Measurements of the Ly α Forest Continuum and Effective Optical Depth with LyCAN and DESI Y1 Data

Wynne Turner^{1,2}, Paul Martini^{1,2,3}, Naim Göksel Karaçaylı^{1,2,3}, J. Aguilar⁴, S. Ahlen⁵, D. Brooks⁶, T. Claybaugh⁴, A. de la Macorra⁷, A. Dey⁸, P. Doel⁶, K. Fanning^{9,10}, J. E. Forero-Romero^{11,12}, S. Gontcho A Gontcho⁴, A. X. Gonzalez-Morales^{13,14}, G. Gutierrez¹⁵, J. Guy⁴, H. K. Herrera-Alcantar¹⁴, K. Honscheid^{2,3}, S. Juneau⁸, T. Kisner⁴, A. Kremin⁴, A. Lambert⁴, M. Landriau⁴, L. Le Guillou¹⁶, A. Meisner⁸, R. Miquel^{17,18}, J. Moustakas¹⁹, E. Mueller²⁰, A. Muñoz-Gutiérrez⁷, A. D. Myers²¹, J. Nie²², G. Niz^{14,23}, C. Poppett^{4,24}, F. Prada²⁵, M. Rezaie²⁶, G. Rossi²⁷, E. Sanchez²⁸, E. F. Schlafly²⁹, D. Schlegel⁴, Michael F. Schubnell³⁰, H. Seo³¹, D. Sprayberry⁸, G. Tarlé³⁰, B. A. Weaver⁸, and H. Zou²²

¹ Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA; turner.1839@osu.edu

² Center for Cosmology & AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

³ Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

⁴ Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

⁵ Physics Department, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA

⁶ Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK

⁷ Instituto de Física, Universidad Nacional Autónoma de México, Ciudad de México C.P. 04510, Mexico

⁸ NSF NOIRLab, 950 North Cherry Ave., Tucson, AZ 85719, USA

⁹ Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Menlo Park, CA 94305, USA

¹⁰ SLAC National Accelerator Laboratory, Menlo Park, CA 94305, USA

¹¹ Departamento de Física, Universidad de los Andes, Carrera 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

¹² Observatorio Astronómico, Universidad de los Andes, Carrera 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia

¹³ Consejo Nacional de Ciencia y Tecnología, Avenida Insurgentes Sur 1582. Colonia Crédito Constructor, Del. Benito Juárez C.P. 03940, México D.F. Mexico

¹⁴ Departamento de Física, Universidad de Guanajuato—DCI, C.P. 37150, Leon, Guanajuato, Mexico

¹⁵ Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA

¹⁶ Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France

¹⁷ Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain

¹⁸ Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra Barcelona, Spain

¹⁹ Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA

²⁰ Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, UK

²¹ Department of Physics & Astronomy, University of Wyoming, 1000 East University, Department 3905, Laramie, WY 82071, USA

²² National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100012, People's Republic of China

²³ Instituto Avanzado de Cosmología A. C., San Marcos 11—Atenas 202. Magdalena Contreras, 10720. Ciudad de México, Mexico

²⁴ Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, USA

²⁵ Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain

²⁶ Department of Physics, Kansas State University, 116 Cardwell Hall, Manhattan, KS 66506, USA

²⁷ Department of Physics and Astronomy, Sejong University, Seoul, 143-747, Republic of Korea

²⁸ CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain

²⁹ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

³⁰ Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

³¹ Department of Physics & Astronomy, Ohio University, Athens, OH 45701, USA

Received 2024 May 10; revised 2024 September 4; accepted 2024 September 5; published 2024 November 19

Abstract

We present the Ly α Continuum Analysis Network (LyCAN), a convolutional neural network that predicts the unabsorbed quasar continuum within the rest-frame wavelength range of 1040–1600 Å based on the red side of the Ly α emission line (1216–1600 Å). We developed synthetic spectra based on a Gaussian mixture model representation of nonnegative matrix factorization (NMF) coefficients. These coefficients were derived from high-resolution, low-redshift ($z < 0.2$) Hubble Space Telescope/Cosmic Origins Spectrograph (COS) quasar spectra. We supplemented this COS-based synthetic sample with an equal number of DESI Year 5 mock spectra. LyCAN performs extremely well on testing sets, achieving a median error in the forest region of 1.5% on the DESI mock sample, 2.0% on the COS-based synthetic sample, and 4.1% on the original COS spectra. LyCAN outperforms principal component analysis (PCA) and NMF-based prediction methods using the same training set by 40% or more. We predict the intrinsic continua of 83,635 DESI Year 1 spectra in the redshift range of $2.1 \leq z \leq 4.2$ and perform an absolute measurement of the evolution of the effective optical depth. This is the largest sample employed to measure the optical depth evolution to date. We fit a power law of the form $\tau(z) = \tau_0(1+z)^\gamma$ to our measurements and find $\tau_0 = (2.46 \pm 0.14) \times 10^{-3}$ and $\gamma = 3.62 \pm 0.04$. Our results show particular agreement



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

with high-resolution, ground-based observations around $z = 2$, indicating that LyCAN is able to predict the quasar continuum in the forest region with only spectral information outside the forest.

Unified Astronomy Thesaurus concepts: Convolutional neural networks (1938); Cosmology (343); Dark energy (351); Intergalactic medium (813); Large-scale structure of the universe (902); Lyman alpha forest (980)

1. Introduction

Observations indicate that the Universe is expanding at an accelerated rate. Measurements of Type Ia supernovae provided the first evidence for this accelerated expansion (A. G. Riess et al. 1998; S. Perlmutter et al. 1999), the cause of which is unknown and called dark energy. Since then, multiple cosmological surveys have contributed to the study of the accelerating expansion, including observations of the cosmic microwave background (e.g., Planck Collaboration et al. 2016) and measurements of the baryon acoustic oscillation (BAO) standard ruler with galaxy and quasar spectra (e.g., S. Cole et al. 2005; D. J. Eisenstein et al. 2005). D. H. Weinberg et al. (2013) provided a comprehensive review of observational probes of cosmic acceleration.

Quasars are among the most luminous objects in the Universe, acting as cosmic flashlights that illuminate the matter-density distribution. The Ly α forest, a collection of absorption lines detected blueward of the Ly α emission line in the spectra of distant quasars, captures this matter-density distribution along the line of sight. The first to measure the BAO scale with the Ly α (Ly α) forest was the Baryon Oscillation Spectroscopic Survey (BOSS; K. S. Dawson et al. 2013) program within the third stage of the Sloan Digital Sky Survey (SDSS-III; D. J. Eisenstein et al. 2011) (e.g., N. G. Busca et al. 2013; A. Slosar et al. 2013; A. Font-Ribera et al. 2014). The successor to BOSS, extended-BOSS (eBOSS; K. S. Dawson et al. 2016) as a part of SDSS-IV (M. R. Blanton et al. 2017), performed the same measurements on a sample of ~ 4 million galaxies and quasars (e.g., R. Ahumada et al. 2020; H. du Mas desBourboux et al. 2020).

While BOSS and eBOSS made great strides in constraining cosmological parameters, the nature of dark energy remains unknown. The Dark Energy Spectroscopic Instrument (DESI), the largest spectroscopic survey to date, aims to achieve the tightest constraints on the dark energy equation of state by measuring spectra of ~ 40 million galaxies and quasars over a five-year period (DESI Collaboration et al. 2016a, 2016b). Roughly 1 million DESI targets are devoted to the high-redshift ($z \geq 2.1$) quasars used to constrain cosmological parameters with the Ly α forest. Each forest absorption feature provides a measurement of the matter-density distribution of the intergalactic medium (IGM) along the line of sight to the quasar. At the DESI resolution, this can correspond to well over a hundred independent absorption measurements per quasar if the entire forest is visible, whereas a quasar used as a discrete tracer only provides one measurement of the matter-density field.

The Ly α forest can also be used to probe properties of the IGM, such as its temperature and ionization state (e.g., J. Schaye et al. 2000; P. McDonald et al. 2001; M. Zaldarriaga et al. 2001; J. S. Bolton et al. 2005; A. A. Meiksin 2009; G. D. Becker et al. 2011, 2015; V. Khaire et al. 2019; M. Walther et al. 2019; P. Gaikwad et al. 2021). Measurements of the effective optical depth of the forest τ_{eff} are one way to study these properties, and the scatter of the optical depth also reveals information about the intensity of the ionizing background radiation produced by active galactic nuclei and star-forming galaxies (e.g.,

G. D. Becker et al. 2015; G. Kulkarni et al. 2019). Other works have also studied the escape fraction of ionizing photons (e.g., V. Khaire et al. 2016) and the evolution of the UV background, including with the development of empirically constrained models (e.g., V. Khaire & R. Srianand 2019; E. Puchwein et al. 2019; C.-A. Faucher-Giguere 2020). M. McQuinn (2016) provided a comprehensive review of the evolution of the IGM. While the Ly α forest is readily accessible at $z > 2$ with ground-based observations, the optical depth is not straightforward to measure because it is difficult to determine the unabsorbed quasar continuum due to the forest absorption.

In order to exploit the Ly α forest for cosmological analysis and studies of the IGM, accurate knowledge of the quasar continuum is required. This becomes increasingly difficult at higher redshifts because the increase in the neutral fraction and density of the high-redshift Universe prohibits direct measurements of the true continuum in the forest region. In the BOSS/eBOSS and current DESI Ly α analysis pipeline, `picca`,³² the flux transmission field $\delta_q(\lambda)$ is measured in the Ly α forest region as a function of observed wavelength λ according to

$$\delta_q(\lambda) = \frac{f_q(\lambda)}{\bar{F}(\lambda)C_q(\lambda)} - 1 \quad (1)$$

where $f_q(\lambda)$ is the observed flux, and the mean expected flux $\bar{F}(\lambda)C_q(\lambda)$ is the product of the mean transmission $\bar{F}(\lambda)$ and the unabsorbed continuum $C_q(\lambda)$ (H. du Mas desBourboux et al. 2020). These deltas $\delta_q(\lambda)$ are measured for each Ly α forest pixel and then auto-correlated with each other or cross-correlated with quasars to achieve a three-dimensional correlation function. Lacking knowledge of the true continuum, a two-parameter fit of the mean continuum $\bar{F}(\lambda)C_q(\lambda)$ in the forest region is performed on the observed flux. However, this introduces spurious correlations that complicate the analysis since the fitting process biases the mean of each delta toward zero for each line of sight. This biasing also results in the loss of information on large scales and projects out large-scale modes, which impacts the power spectrum in one dimension (e.g., N. G. Karaçaylı et al. 2020, 2022; C. Ravoux et al. 2023; N. G. Karaçaylı et al. 2024) and in three dimensions (e.g., M. L. A. Karim et al. 2023; R. de Belsunce et al. 2024). The standard approach is now to use a distortion matrix to accurately model and account for these spurious correlations. J. E. Bautista et al. (2017) provide a detailed description of the distortion matrix.

There have been a number of attempts to measure the unabsorbed Ly α forest continuum. Some previous works have traced the transmission peaks in the Ly α forest region in high-resolution spectra to estimate the true continuum (e.g., M. Rauch et al. 1997; J. Schaye et al. 2003), but this is likely inaccurate for lower resolution spectra and at higher redshifts where the transmission peaks do not represent the unabsorbed continuum. Others have used composites of

³² <https://github.com/igmhub/picca/>

medium-resolution spectra to measure the continuum (e.g., G. D. Becker et al. 2013), but this is also not a direct measurement of the continuum at high redshifts.

Because a correlation exists between the continuum shape including broad emission lines on the red and blue sides of the Ly α emission line (e.g., I. Pâris et al. 2011; B. Greig et al. 2017), other works have attempted to predict the true continuum using only longer wavelength information (redward of the Ly α emission line) with principal component analysis (PCA; e.g., N. Suzuki et al. 2005; I. Pâris et al. 2011; K.-G. Lee et al. 2012; F. B. Davies et al. 2018). K.-G. Lee et al. (2012) also corrected their continuum predictions using external constraints on the Ly α forest mean flux, making their method dependent on the Ly α forest region. In general, dimensionality reduction techniques such as PCA or nonnegative matrix factorization (NMF) are limited by their ability to learn linear representations of the data, while neural networks are able to learn more complex, nonlinear representations. We will show that PCA- and NMF-based prediction models do not work as well as a deep learning approach that has the ability to learn nonlinear relationships.

Several recent efforts have illustrated the potential of deep learning approaches. Some have used such approaches to infer properties of the IGM (e.g., L. Huang et al. 2021; P. Nayak et al. 2023; F. Nasir et al. 2024) or fit the absorption lines in the forest with Voigt profiles (P. Jalan et al. 2024). Additionally, deep learning can be used to predict the unabsorbed continuum in the forest region. B. Liu & R. Bordoloi (2021) developed a feedforward neural network called iQNet that was designed to predict the quasar continuum in the Ly α forest region using only information on the red side. However, their training set was entirely based on Cosmic Origins Spectrograph (COS) quasars and the resulting model is likely not optimized for DESI quasars, e.g., the DESI quasars are generally much more luminous. Another effort from Z. Sun et al. (2023) modeled the quasar continuum using an unsupervised probabilistic model. This model achieves promising results, although is not independent of the observed flux in the Ly α forest region. The dependence on flux in the forest region also does not enable an independent measurement of the mean absorption.

In this work, we present a convolutional neural network (CNN) to predict the Ly α forest continuum that we call the Ly α Continuum Analysis Network (LyCAN). We used a Gaussian mixture model (GMM) representation of the NMF coefficients of low- z ($z < 0.2$) Hubble Space Telescope HST/COS spectra to develop a sample of 40,000 synthetic spectra. We altered the emission line strengths of half of these spectra so that they are more representative of high-luminosity quasars. We trained our neural network on a random subset of this sample of 40,000 COS-based synthetic spectra combined with an equal number of DESI Year 5 (Y5) mock spectra (e.g., J. Farr et al. 2020; H. K. Herrera-Alcantar et al. 2023) and used the resulting network to predict the Ly α forest continua for the training and testing samples, a larger sample of DESI Y5 mock spectra, and the real quasar spectra from DESI Year 1 (Y1).

This paper is organized as follows. In Section 2, we introduce the data used in this study, including the low- z HST/COS spectra, DESI mocks, and DESI Y1 spectra. In Section 3, we detail the generation of our synthetic spectra. These synthetic spectra were generated from archival COS observations of low-redshift quasars, which we augmented with mock spectra generated for the DESI Ly α forest analysis. We discuss

the creation and performance of our neural network in Section 4, and compare its performance with PCA- and NMF-based prediction models. These PCA and NMF models use the same training set as LyCAN and are distinct from the NMF representation that is used to help generate synthetic spectra. In Section 5, we use the LyCAN continuum predictions to calculate the evolution of the effective optical depth on DESI Y1 quasars and compare these results to previous works in the literature. Finally, we discuss our conclusions and planned future work in Section 6.

2. Data

We used several data sets for several purposes. The first purpose is the generation and assembly of a synthetic data set for training, testing, and validation of LyCAN. This data set includes archival observations from HST/COS. We describe this sample in Section 2.1, along with our processing steps to prepare these data to create the training set. DESI observations are important to both ensure our synthetic data are representative of DESI data and for our measurement of the mean optical depth. We describe DESI observations in Section 2.2. Finally, we supplemented the COS sample with DESI mocks to construct our synthetic data set, as well as used synthetic mock observations to characterize uncertainties in our fit to the mean optical depth evolution. We describe these DESI mocks in Section 2.3.

2.1. HST/COS

We used low-redshift ($z < 0.2$) quasars observed by the COS instrument on the HST to help construct our training, testing, and validation samples. Low-redshift quasars are very minimally affected by Ly α forest absorption, and they therefore provide the best measurement of the true continuum in that region.

We retrieved spectra from the Hubble Spectroscopic Legacy Archive (HSLA; M. Peeples et al. 2017), and then selected the subset with rest-frame wavelength coverage from 1070 to 1600 Å and a median signal-to-noise ratio (SNR) of at least five per resolution element. We chose this minimum rest wavelength requirement to maximize our sample size because very few quasars in the HSLA fully span the entire Ly α forest region down to 1040 Å. Furthermore, some spectra suffer from a significant number of missing data pixels in this more narrow forest region. We removed QSOs that were missing more than 30% of the native pixels in the minimum forest region (1070–1200 Å), and then manually removed suspected broad absorption line (BAL) quasars. This resulted in a sample of 48 QSOs. Finally, we transformed each quasar spectrum to units of relative flux by normalizing by the mean flux in the 2 Å interval centered around 1280 Å.

Following the work of B. Liu & R. Bordoloi (2021), we used the LineTools API³³ to interactively fit a smooth continuum model to each of these quasar spectra. This included interpolating over geocoronal Ly α emission and Galactic absorption features. Some spectra were missing data that extend to 1040 Å. In these cases, we extrapolated the LineTools continua to 1040 Å using the eBOSS DR14 (M. Blomqvist et al. 2019; V. de Sainte Agathe et al. 2019) mean quasar continuum. We also corrected each quasar for extinction due to

³³ <https://linetools.readthedocs.io>




Figure 1. HST/COS quasar spectrum and the smooth representation from LineTools. The original COS spectrum is in gray and the continuum fit is in blue. The continuum includes the extrapolation to 1040 Å (see Section 3.1).

interstellar dust. We retrieved $E(B - V)$ values for each of our COS objects from E. F. Schlafly & D. P. Finkbeiner (2011) and dereddened each spectrum with the K. D. Gordon et al. (2023) extinction model assuming $R_V = 3.1$. We then eliminated 10 COS objects that appeared least similar to DESI-like quasars, which yielded a final sample of 38 COS quasars. We detail the continuum extrapolation and outlier removal procedures in Section 3.1. Figure 1 shows an example of a LineTools continuum fit and its continuum extrapolation. The redshift distribution of the initial 48 COS objects and the final 38 objects after outlier removal is shown in Figure 2. In Section 3.1, we describe how we used this sample of COS objects to generate 40,000 synthetic spectra.

2.2. DESI

DESI was installed at the Mayall 4 m telescope at Kitt Peak National Observatory (DESI Collaboration et al. 2022) and is in the midst of the largest spectroscopic survey of the Universe. The aim of DESI is to measure spectra for ~ 40 million galaxies and quasars from the local Universe to $z > 3.5$ (DESI Collaboration et al. 2016a) and study the nature of dark energy. The main cosmological tools are the BAO standard ruler to measure distances and constrain cosmological parameters, and redshift-space distortions to measure the growth of structure and test possible modifications to general relativity. DESI has 10 spectrographs and 5000 fibers that may be positioned across a 3° diameter field of view, which allows for efficient observations of 5000 targets simultaneously. The main DESI survey began in 2021 May upon the completion of commissioning and survey validation (DESI Collaboration et al. 2024a). An early data release included data from the survey validation period and the first two months of the main survey (DESI Collaboration et al. 2024b; C. Ramírez-Pérez et al. 2024). The DESI collaboration recently completed a series of papers that describe cosmological results from the first year of data collection (DESI Collaboration et al. 2024c, 2024d, 2024e). Additional key papers (DESI Collaboration et al. 2024, in preparation) based on Data Release 1 (DR1; DESI Collaboration et al. 2024, in preparation) are forthcoming.

The purpose of LyCAN is to predict the Ly α forest continuum for DESI quasars. The DESI Y1 quasar catalog contains $\sim 450,000$ Ly α forest ($z \geq 2.1$) quasars. In this paper, we demonstrate the performance of LyCAN with a subset of

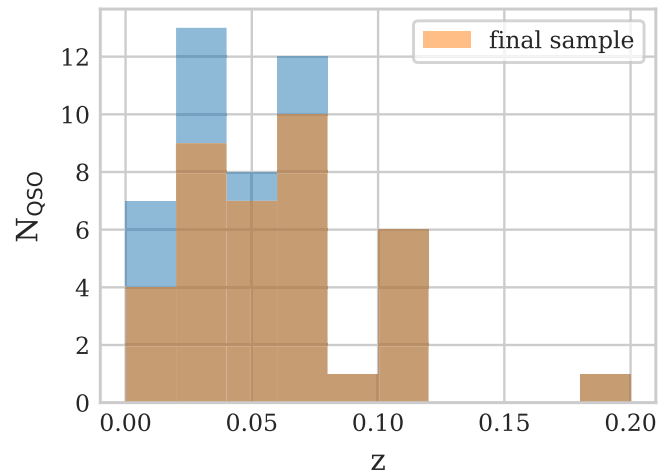


Figure 2. Redshift distribution of the HST/COS objects. The blue distribution is the original sample of 48 objects, while the orange distribution shows the final sample of 38 objects after outlier removal (see Section 3.1).

these quasars where the rest wavelength coverage spans at least 1160–1600 Å. We also eliminated spectra with redshift warnings and those with BALs and/or damped Ly α absorbers (DLAs). A total of 250,955 quasars satisfy these conditions. We resampled these spectra onto a rest wavelength grid from 1040 to 1600 Å with a dispersion of 0.8 Å per pixel and normalized by the 40 pixels surrounding 1280 Å. We corrected each quasar for extinction due to interstellar dust, as described in Section 2.1.

We used a subset of the DESI Y1 quasars with median r -band SNR ≥ 5 to measure the evolution of the Ly α optical depth. For this subset, we did not eliminate quasars with known DLA systems because we later corrected our measurements for optically thick absorbers (see Section 5.2). The other selection criteria were unchanged. This sample contains 83,635 quasars, which is the largest sample to date used for measurements of the effective optical depth.

2.3. DESI Mocks

We supplemented our synthetic data set for training, testing, and validation with DESI mock spectra that were constructed for the main cosmological analysis (e.g., J. Farr et al. 2020; H. K. Herrera-Alcántar et al. 2023). The DESI mock spectra were generated with a software suite called `quickquasars`.³⁴ The quasar continua used by `quickquasars` are generated with the `simqso`³⁵ library, which models the continuum as a broken power law and is tuned to better reproduce the mean continuum observed in eBOSS DR16 (H. du Mas desBourboux et al. 2020) data via a PCA. These unabsorbed continua are multiplied by transmission files generated with the `LyCoLoRe` program (J. Farr et al. 2020). Instrumental effects and astrophysical contaminants are also included in `quickquasars` to produce “fully contaminated” mocks that include BALs, DLAs, high column density systems, and metal lines.

The final sample of mock spectra are representative of the DESI magnitude and redshift distribution, including with realistic noise and instrumental effects. For more details about

³⁴ <https://github.com/desihub/desim/blob/main/py/desim/scripts/quickquasars.py>

³⁵ <https://github.com/imcgreer/simqso>

the synthesis of the DESI mocks, see H. K. Herrera-Alcantar et al. (2023). For this work, we used the “fully contaminated” Year 5 (Y5) mock spectra because these best represent the eventual quality of DESI data and already span the noise properties of Y1. The main distinction is that the Ly α quasars will receive multiple observations over the course of the survey, and so the Y1 quasars will eventually have somewhat higher mean SNR. We used a realization of the mocks that contains the true continua in the forest region, in addition to the simulated (observed) forest. We conducted our analysis with a random sample of 40,000 mock spectra that exclude BALs and DLAs. These spectra fully cover the rest wavelength range of 1040–1600 Å. This sample of mock spectra spans from $2.43 < z < 2.72$.

We used a different realization of Y5 mocks with median r -band SNR ≥ 5 to compute the evolution of the effective optical depth with both LyCAN continuum predictions and the truth continua. This sample consists of 311,833 mock quasars with $2.1 < z < 3.8$, and includes DLAs but not BALs. In Section 5, we show these results and discuss how we used them to calibrate our systematic uncertainties.

3. COS-based Synthetic Spectra

We constructed a sample of 40,000 synthetic spectra from our base sample of archival COS observations (see Section 2.1) to produce a sufficient sample to train our neural network. This process of data augmentation is a standard practice for training neural networks to increase the size of the training set, e.g., to help teach the network to ignore certain features such as noise. We combined this COS-based synthetic sample with DESI mock spectra to construct our training, testing, and validation data sets. In Section 3.1 we detail the preprocessing steps required to generate our synthetic spectra and in Section 3.2 we describe the methods we used to create the set of 40,000 synthetic spectra.

3.1. Preprocessing

We require full coverage of the Ly α forest region for training of the neural network. It is therefore important that the basis of our sample of synthetic spectra covers the full forest; however, many of the COS spectra are missing data on the blue end of the forest region. We extended these spectra to cover the full rest-frame wavelength range of 1040–1600 Å with the mean continuum measured by picca in eBOSS DR14 (M. Blomqvist et al. 2019; V. de Sainte Agathe et al. 2019). For each COS spectrum with missing data on the blue end of the Ly α forest region, we performed this extrapolation as follows: we first normalized the eBOSS mean continuum to the COS LineTools continuum over the first 10 Å of available data. Within that 10 Å range, we computed the COS continuum as a linear combination of the LineTools continuum and the eBOSS mean continuum, beginning with 100% and 0% contributions from the COS and eBOSS continua on the red end, respectively, and ending with 0% and 100% contributions from the COS and eBOSS continua on the blue end of the available data. The remaining COS continuum in the region with missing data was then taken to be equal to the eBOSS mean continuum in that region.

After completing this extrapolation, we had 48 COS continua with full coverage in the rest wavelength range of 1040–1600 Å. However, low- z quasars such as those in the COS data set tend to be lower luminosity than typical DESI




Figure 3. Reduced χ^2 for NMF models with a number of components ranging from 2 to 20. These components were derived from our initial sample of 48 COS continua. We used models with six components throughout this paper.

quasars. This is potentially a problem as there is an anti-correlation between the luminosity of the quasar continuum and the equivalent widths of its emission lines (J. A. Baldwin 1977). Because the COS quasars tend to be lower luminosity, they will feature stronger broad emission lines, and therefore may not be representative of the DESI sample. In order to reduce the low-luminosity bias of our synthetic spectra, we removed the COS objects that are least similar to DESI Y1 quasars from our sample.

We identified the COS quasars that are least similar to the DESI quasars through an analysis of their spectra on the red side of the Ly α emission lines. First, we derived NMF components from our 48 COS continua. NMF is a dimensionality reduction technique similar to PCA but the components are restricted to be nonnegative, which is more physically representative of quasar spectra. The components are not necessarily orthogonal as a consequence of this nonnegativity constraint. We used the NMF code developed by G. Zhu (2016), which includes the ability to use weights and mask pixels. We masked pixels where there was a $\geq 5\sigma$ difference between the LineTools continuum and flux, and derived a set of six NMF components from our COS continua. Figure 3 shows the reduced χ^2 curve for NMF models for 2–20 components, which begins to flatten out at six components. We fit a linear combination of these six NMF components to the 1230–1600 Å region for both the 48 COS quasars and 10,000 relatively high-SNR (median r -band SNR ≥ 5) DESI Y1 quasars. The six-dimensional parameter space of coefficients resulting from this fit represents the shape of DESI quasar spectra on the red side of Ly α .

We fit a GMM to the DESI coefficients and chose five classes based on the Bayesian information criterion (BIC). The distribution of coefficients is shown in Figure 4. Ten of the COS quasars fall significantly outside both the observed coefficient distribution of points from DESI and the GMM. The 10 that we eliminated all have probabilities $< 10^{-10}$ of being described by the GMM. The COS quasars that we retained are marked with blue points in Figure 4 and the ones that we eliminated are marked with red points.

3.2. Synthetic Spectra

We generated synthetic spectra based on a GMM of our final sample of COS continua. We first removed two BL-Lac objects

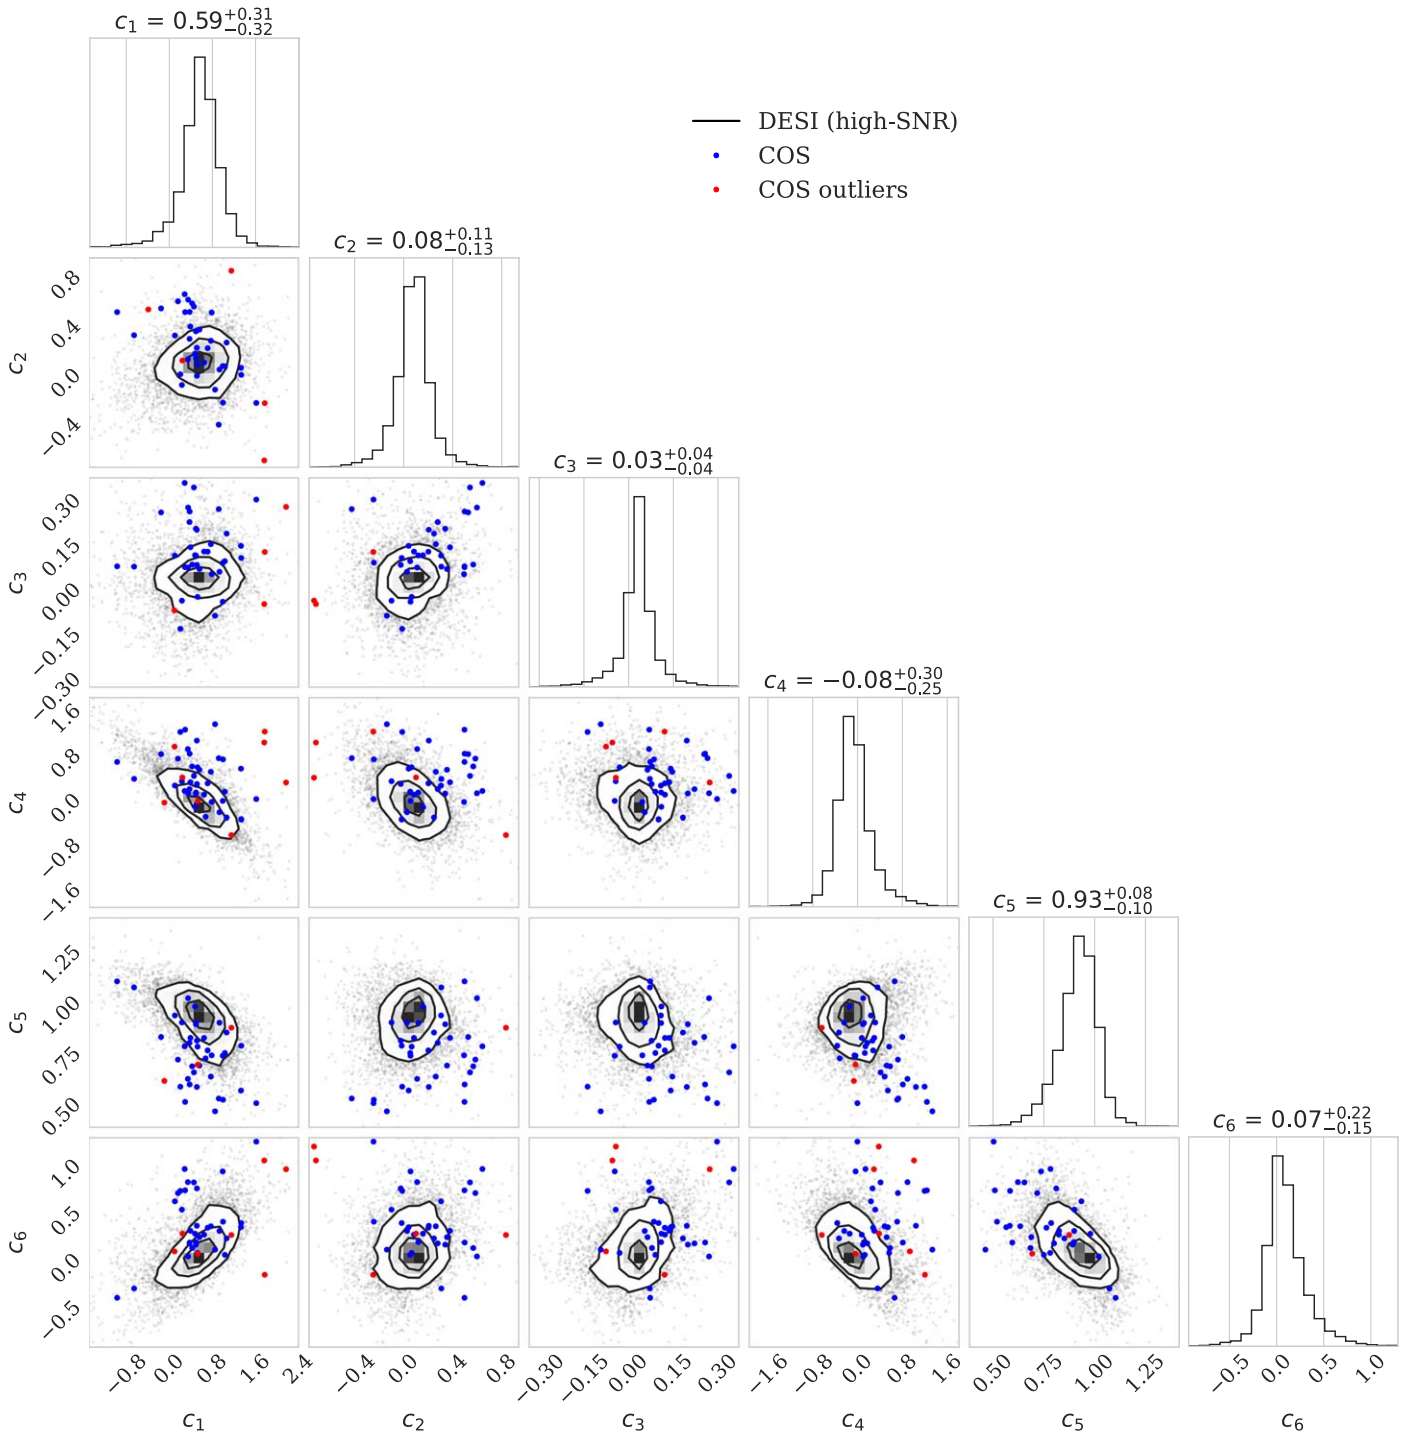


Figure 4. Distribution of coefficients obtained from fitting six NMF components to the 1230–1600 Å region of 10,000 DESI Y1 (gray points and black contours) and 48 COS quasar spectra (large blue and red points). The six components were derived from COS observations, where c_i indicates the coefficient associated with component i . The red points mark the 10 COS objects with the lowest likelihood ($<10^{-10}$) of being described by the GMM representation of the DESI distribution. These 10 quasar spectra (red points) were removed as unrepresentative of DESI quasars. The remaining 38 (blue points) were retained for the analysis.

from our COS sample and we then used the pixel masking criteria described in Section 3.1 to derive six new NMF components from the remaining sample of 36 masked COS continua. To create the synthetic spectra, we generated a GMM with four classes from the six-dimensional parameter space of NMF coefficients. We used the BIC to choose a four-component mixture model. The mean spectrum for each of these four classes is shown in Figure 5 and the main difference between the classes is the strength of the emission lines. For

each of the four classes, we randomly selected an equal number of COS continua that are members of the corresponding class. Afterwards, we added a representative number of BL-Lac spectra based on the fraction of those present in the COS sample (two out of 38). The inclusion of BL-Lac-like objects allows our network to generalize better to DESI-like data with weaker emission line properties. Our sample of synthetic COS continua totals 20,000.




Figure 5. The mean spectra for our four classes of COS continua from the GMM representation of the NMF coefficients.

This sample of continua forms the basis of our synthetic spectra. To create more diversity, we linearly perturbed each of these spectra according to

$$C_{\text{perturb}}(\lambda) = C(\lambda) + M_0 + M_1 \left(\frac{\lambda/\text{\AA} - 1040}{1600 - 1040} \right) \quad (2)$$

as in Z. Sun et al. (2023), where M_1 was drawn from a uniform distribution from -0.2 to 0.2 and affects the slope of the continuum, and M_0 was drawn from a different uniform distribution and affects the strength of emission lines. To create COS-like synthetic spectra, we drew M_0 from a uniform distribution from -0.1 to 0.1 so as to not systematically affect the emission line strengths. For another sample of 20,000 synthetic spectra that are more DESI-like, we drew M_0 from a uniform distribution from 2 to 6 based on visual inspection of DESI Y1 quasars. The higher values of M_0 reduce the strengths of the emission lines relative to the continuum, which is more representative of higher redshift, higher luminosity quasars. These linear perturbations add diversity to our synthetic spectra and increase the generalizability of our neural network.

After adding perturbations, we renormalized each spectrum at 1280 \AA . We then added 2–4 random absorption features to the red side ($1216\text{--}1600 \text{ \AA}$) of our perturbed continua to mimic metal absorption lines. Each absorption line is a Voigt profile with a Lorentzian full width at half maximum (FWHM) of 0.1 and a Gaussian FWHM of either 0.4 for the COS-like spectra or 1.0 for the DESI-like spectra. The Lorentzian amplitude was randomly selected between 2 and 5 for the COS-like spectra or 1.5–5 for the DESI-like spectra. The range of absorption features and the parameter values were chosen based on visual inspection of the COS and DESI Y1 spectra.

Finally, we added noise to each of these synthetic spectra. In order to incorporate a diverse SNR range into our synthetic sample, we added noise according to the SNR range observed in COS and DESI Y1 data. For the COS-like sample, we added Gaussian noise to each of the 20,000 synthetic spectra by randomly selecting error spectra belonging to the associated GMM class. We repeated this for the 20,000 DESI-like synthetic spectra using a sample of 40,000 DESI Y1 quasars that span the noise properties of the full Y1 sample.

Our final sample contains 40,000 COS-based synthetic spectra, of which 20,000 are COS-like in their emission line and noise properties, and the other 20,000 are DESI-like. This diversity in synthetic spectra helps ensure that our neural network learns to be robust against varying levels of noise, emission line strengths, and metal absorption line properties. The dispersion of our DESI-like COS-based synthetic spectra is shown in the second panel of Figure 6.

4. Ly α Continuum Analysis Network

In this section, we present our neural network for predicting the quasar continuum, the LyCAN. We first describe the data set we used to train LyCAN in Section 4.1. This data set is based on the COS-based synthetic spectra and DESI mocks described in the previous sections. We then describe how we optimized the hyperparameters for LyCAN and the network architecture in Section 4.2. Finally, we quantify the performance of our network and compare it to that of PCA- and NMF-based prediction models in Section 4.3.

4.1. Training Set

The goal of our work is to predict continua for DESI quasars. Additionally, we want to avoid biasing our network toward low-luminosity quasars, which are over-represented in the COS data set. We combined our 40,000 COS-based synthetic spectra with a sample of 40,000 DESI Y5 mock spectra with true continua. These 80,000 spectra are the input data for our neural network. We performed a randomized 65%/25%/10% split to select our training, testing, and validation sets, respectively. The testing set was never seen during training and is only used to quantify the performance of our model afterwards, while the validation set was used during the training process to evaluate the mean squared error loss function at each epoch, allowing for the determination of optimal weights.

The diversity of these spectra is shown in Figure 6. This figure shows these samples before the training/testing/validation split, but each class of spectra is equally represented in each of these subsets. The third panel shows the combination of DESI Y5 mock spectra and COS-based synthetic spectra with COS-like and DESI Y1-like noise, which is representative of our training set.

4.2. Network

We created the convolutional neural network (CNN) model LyCAN to predict the forest continuum. A neural network model is a good choice for this application due to its ability to learn nonlinear representations of the data, unlike PCA- and NMF-based methods. A CNN is a type of feedforward neural network with one or more convolutional layers that convolve the input pixels with learned filters. These layers are designed to capture hierarchical features in the input data. Optimizable hyperparameters in these convolutional layers include the number of filters and kernel size, among others. Dense (fully connected) layers are often added to the latter part of the network to combine high-level features learned by the convolutional layers. The number of nodes per dense layer is another optimizable hyperparameter.

We optimized the CNN architecture with a pseudo-random search in which we varied the number of convolutional layers, the kernel size and number of filters for these layers, the number of dense layers, and the number of nodes per dense layer. The number of convolutional layers ranged from one to




Figure 6. Illustration of the diversity of the spectra. From left to right: DESI Y5 mock spectra included in our input data, COS-based synthetic spectra with DESI Y1-like noise, the combination of Y5 mock spectra and COS-based synthetic spectra with COS-like and DESI Y1-like noise, and the real DESI Y1 data. The shaded areas indicate the 1 and 2 σ regions. The third panel is representative of our training set.




Figure 7. AFFE scores for candidate CNN architectures. Filled symbols indicate that a kernel size of three was used for the convolutional layers, while open symbols indicate a kernel size of six. The colors and symbols of the points vary depending on the number of convolutional layers in the network. Our chosen architecture is marked with a large circle in the lower left-hand corner and minimizes both the median AFFE score ($|\delta F|_{m\text{-score}}$; Equation (4)) and the 1 σ range AFFE score ($|\delta F|_{p\text{-score}}$; Equation (5)).

four (including the input layer), and dense layers also ranged from one to four (not including the output layer). We trained and tested 10 different models for each possible combination (e.g., one convolutional layer and one hidden dense layer for the least complex models, up to four convolutional layers and four hidden dense layers for the most complex models). Each of these models was generated by randomly selecting the number of filters per convolutional layer and the number of nodes per hidden dense layer from a range of options. The number of filters was chosen to be a multiple of 16 between 16 and 64 and the number of nodes was chosen to be a multiple of 32 between 32 and 256. The number of nodes in the output

dense layer is always equal to the number of wavelength pixels in the full spectral range.³⁶ We separately trained these 160 configurations with a kernel size of three and a kernel size of six for a total of 320 unique neural network architectures.

We trained each of these unique architectures with the training subset of our data and quantified each model’s performance on the testing subset. Among the testing subset, we evaluated the performance on the COS-based synthetic and DESI Y5 mock spectra with the absolute fractional flux error (AFFE), defined as

$$|\delta F| = \frac{\int_{\lambda_1}^{\lambda_2} \left| \frac{F_{\text{pred}}(\lambda) - F_{\text{true}}(\lambda)}{F_{\text{true}}(\lambda)} \right| d\lambda}{\int_{\lambda_1}^{\lambda_2} d\lambda}, \quad (3)$$

where F_{true} is the true continuum, F_{pred} is the predicted continuum, $\lambda_1 = 1040 \text{ \AA}$, and $\lambda_2 = 1200 \text{ \AA}$. We compared the resulting AFFE values in the forest region for each trained model by calculating a median AFFE score per model of

$$|\delta F|_{m\text{-score}} = \sqrt{|\tilde{\delta F}|_{\text{S}}^2 + |\tilde{\delta F}|_{\text{M}}^2} \quad (4)$$

where $|\tilde{\delta F}|_{\text{S}}$ refers to the median AFFE for the COS-based synthetic testing sample and $|\tilde{\delta F}|_{\text{M}}$ refers to the median AFFE for the DESI mock testing sample. In addition, we evaluated another metric based on the percentile range of AFFEs,

$$|\delta F|_{p\text{-score}} = \sqrt{|\delta F|_{\text{S},1\sigma}^2 + |\delta F|_{\text{M},1\sigma}^2}, \quad (5)$$

where $|\delta F|_{\text{S},1\sigma}$ denotes the 1 σ AFFE range for the COS-based synthetic testing sample and $|\delta F|_{\text{M},1\sigma}$ is the 1 σ AFFE range for the DESI mock testing sample. Figure 7 shows the metrics in Equations (4) and (5) for a selection of the tested models.

We chose the architecture that minimized the median AFFE and 1 σ range scores. This architecture features six hidden layers, three of which are convolutional layers and the

³⁶ Due to edge effects from resampling, the rest-frame wavelength range actually spans from 1040.8–1599.2 \AA (inclusive).

Table 1
LyCAN Architecture

Layer Type	Size	Activation Function
Conv1D (input)	64	ReLU
MaxPooling1D	2	
Conv1D	32	ReLU
MaxPooling1D	2	
Conv1D	48	ReLU
MaxPooling1D	2	
Conv1D	16	ReLU
MaxPooling1D	2	
Flatten		
Dense	96	ReLU
Dense	64	ReLU
Dense	128	ReLU
Dense (output)	699	Linear

Note. The first column specifies whether the layer is a convolutional layer, a max pooling layer (occurring after each convolutional layer), or a dense (fully connected) layer. The size column specifies the number of filters for convolutional layers or neurons for dense layers. In the case of the max pooling layers, pixels are pooled in groups of two. We use the linear activation function for the output layer, and the ReLU activation function for all other layers where applicable. The number of nodes in the output layer is equal to the number of pixels in the predicted continuum. The kernel size is six for the convolutional layers.

remaining three are dense layers. The full architecture is shown in Table 1. LyCAN was compiled and trained using Keras³⁷ with the TensorFlow³⁸ backend. We used the linear activation function for the output layer and the rectified linear unit (ReLU; A. Krizhevsky et al. 2012) activation function for all other layers, the mean squared error loss function, and the Adam optimizer (D. P. Kingma & J. Ba 2017). We trained for a maximum of 50 epochs and implemented early stopping criteria that monitored the validation loss starting at epoch 20 with a patience of five and a minimum delta of 10^{-5} . The patience sets the number of epochs with negligible improvement after which training is stopped and optimal weights are restored, while the delta sets the minimum change in the validation loss to qualify as an improvement. There are a total of 169,851 trainable parameters (i.e., weights and biases) in our network. Finally, the output continua were smoothed using a Gaussian filter with a standard deviation of three in order to reduce the noise in the predictions and minimize the AFFEs.

We tested other versions of CNNs for this work. For example, we trained a CNN with the same architecture on only mock spectra and tested it on each of the same testing sets. We found no improvement in the performance on mock spectra, and a significant reduction in performance on our COS-based synthetic spectra, including the set with DESI-like emission line and noise properties. We also trained a CNN on only a sample of COS-based synthetic spectra with COS-like emission line and noise properties, and found equally good performance on the testing set for the same type of objects, but markedly poorer performance on the COS-based synthetic spectra with DESI-like properties and the DESI mock spectra. Finally, we trained a CNN on only COS-based synthetic spectra containing an equal representation of COS- and DESI-like properties. Median AFFEs on the COS-based synthetic samples remained

low and a moderate median AFFE of ~ 0.08 on the DESI mock spectra was achieved. This shows that the addition of COS-based synthetic spectra with DESI-like emission line and noise properties helps the network more accurately learn how to predict continua for DESI-like objects. Ultimately, a combination of COS-based synthetic spectra with varying emission line and noise properties and DESI mock spectra proved to be most effective.

4.3. Performance and Comparisons

We quantify the performance of our network with the AFFE metric defined in Equation (3). We also used its nonabsolute counterpart, the fractional flux error (FFE), to study systematic biases in the LyCAN continuum predictions.

Figure 8 shows continuum predictions for real DESI Y1 QSOs and Y5 mock QSOs in our testing set at several redshifts. The truth continua for the mock spectra are also shown, but are difficult to distinguish because the LyCAN predictions are nearly identical. AFFE values are shown in the legend for the mock spectra.

We created PCA- and NMF-based prediction models to compare the performance of these two other well-known methods to LyCAN. These models are based on the same training data as LyCAN, and are distinct from the previous NMF models developed to remove COS outliers and help generate COS-based synthetic spectra. We used scikit-learn³⁹ to derive PCA components from the training set continua after z-score standardization. This scaled each spectrum so that its mean continuum is zero and standard deviation is one, which reduces possible biases due to luminosity. Once we computed the PCA continuum predictions, we reversed the standardization. We optimized the PCA and NMF models by testing models with up to 16 components. For each model, we fit the components to the red side of the training data (1216–1600 Å) and then used the best-fit coefficients to extrapolate the continuum from the red side to the full spectral range, including the forest. We used the median AFFE on the DESI mock testing set as a performance metric to find the optimal models. For both PCA and NMF, we found that four components minimized this metric, and we therefore adopted these as the optimal models.

Figure 9 shows a comparison of the PCA and NMF continuum predictions to LyCAN for four types of spectra. Neither technique works as well as LyCAN. In addition to their poorer performance, the components of the PCA and NMF models must be fit to smoothed continua in order to yield reasonable results. We therefore used the red side of the continua output by LyCAN as input to our PCA- and NMF-based models when predicting continua for DESI Y1 quasars.

LyCAN outperforms these PCA- and NMF- based models. Figure 10 shows the distribution of AFFE values for different models and testing samples. The median AFFE values for different data sets and models are shown in Table 2. The first column shows the median AFFE values for these three models on the real COS quasars used to generate the synthetic data in the next two columns. LyCAN is able to predict the forest continua for the COS sample more accurately than the NMF and PCA models, although this sample is quite small. The next two columns of Table 2 list the median AFFEs of each model

³⁷ <https://keras.io/>

³⁸ <https://www.tensorflow.org/>

³⁹ <https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.PCA.html>

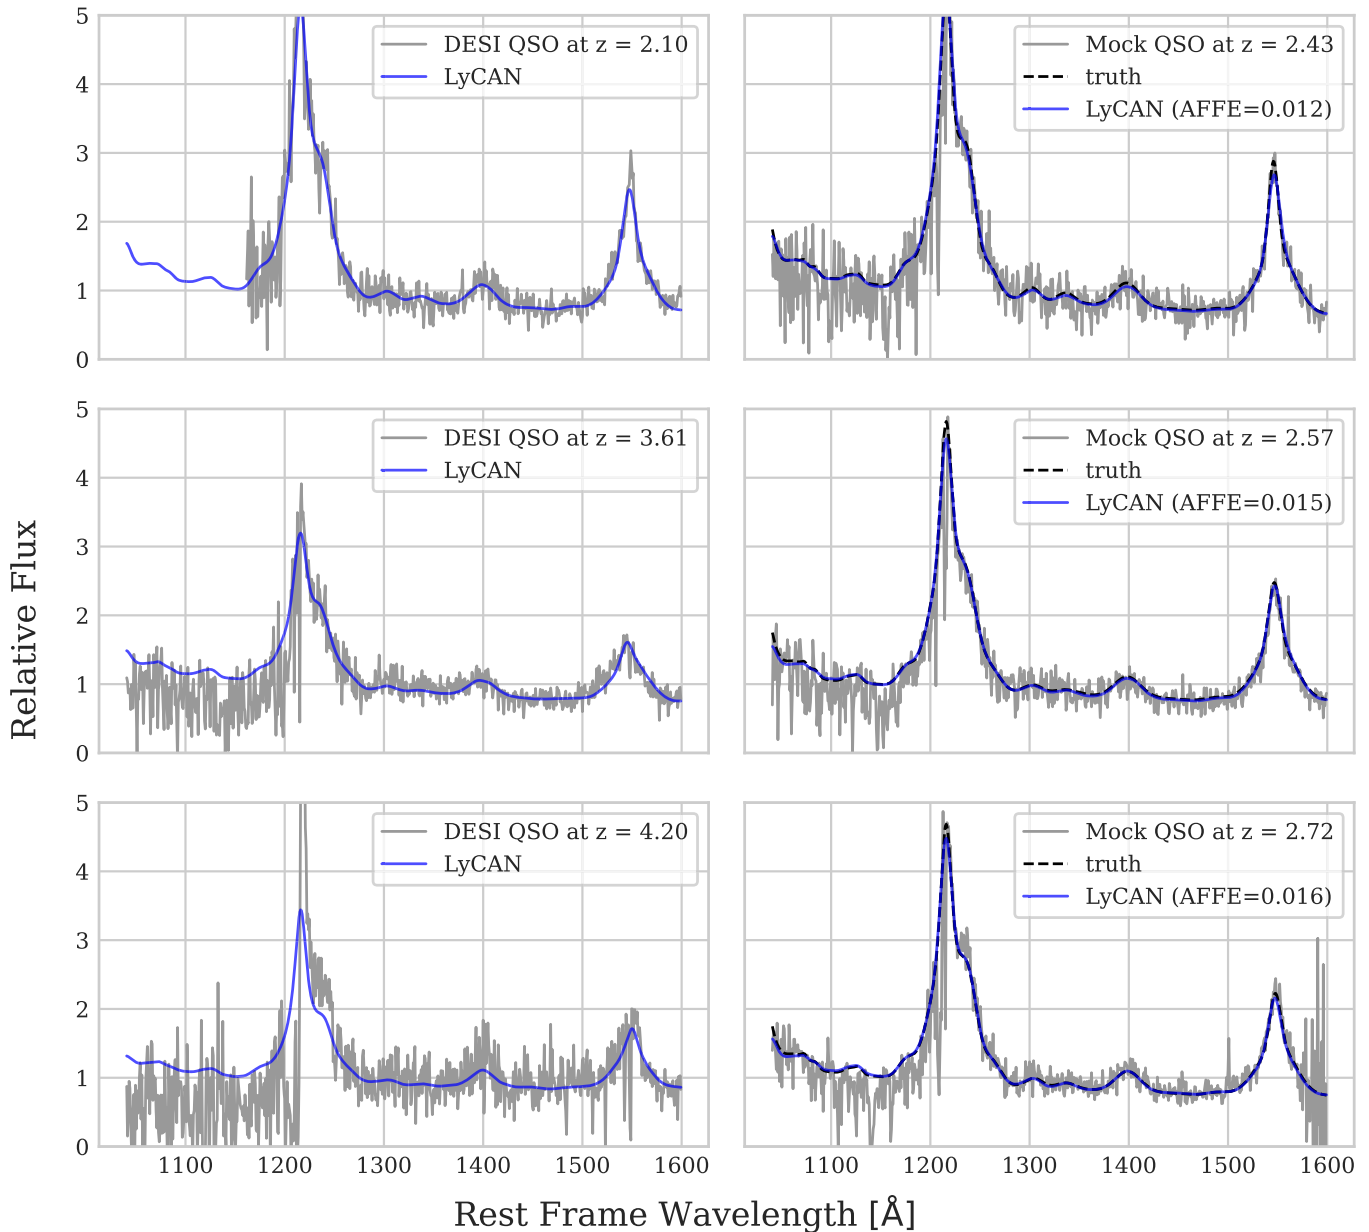


Figure 8. LyCAN continuum predictions for DESI Y1 and DESI Y5 mock quasars. The panels on the left-hand side show predictions for real DESI Y1 quasars at increasing redshifts from top to bottom, and the right-hand panels show predictions on DESI Y5 mock spectra. The panels with mock spectra also include their true continua and the AFFE values in the forest region.

on the synthetic training and testing spectra, respectively, generated from the COS objects. Both of these subsets include synthetic spectra with DESI- and COS-like noise. LyCAN again outperforms the PCA and NMF models, achieving a median AFFE of 2.0%. We see nearly identical performance on the training versus testing subsets across all three models. Finally, the last two columns of Table 2 show the median AFFEs from each prediction model on the DESI Y5 mock training and testing sets. Again, LyCAN performs the best out of all the models, achieving a median AFFE of 1.5%, approximately 40% better than the performance of the PCA model on the same data set (2.6%) and almost five times better than the performance of the NMF model (7.4%). Interestingly, the NMF model performs more poorly than the PCA model in all cases except for the real COS sample. This is a surprising

result and may indicate that the NMF components are not as generalizable to more diverse data sets as the PCA components, whether that is due to the number of components used, the components themselves, or the standardization process. We also note that while the number of components has been optimized for our PCA and NMF models, more sophisticated PCA-based models in the literature that use pixel weighting and/or a projection matrix approach may offer better performance (e.g., N. Suzuki et al. 2005; I. Pâris et al. 2011; K.-G. Lee et al. 2012; F. B. Davies et al. 2018).

We investigated the dependence of our DESI Y1 LyCAN continuum predictions on redshift and SNR. Figure 11 shows the dispersion of continuum predictions for DESI Y1 quasars as a function of SNR and redshift. Overall, there is slightly less dispersion in the high-SNR predictions, while the median

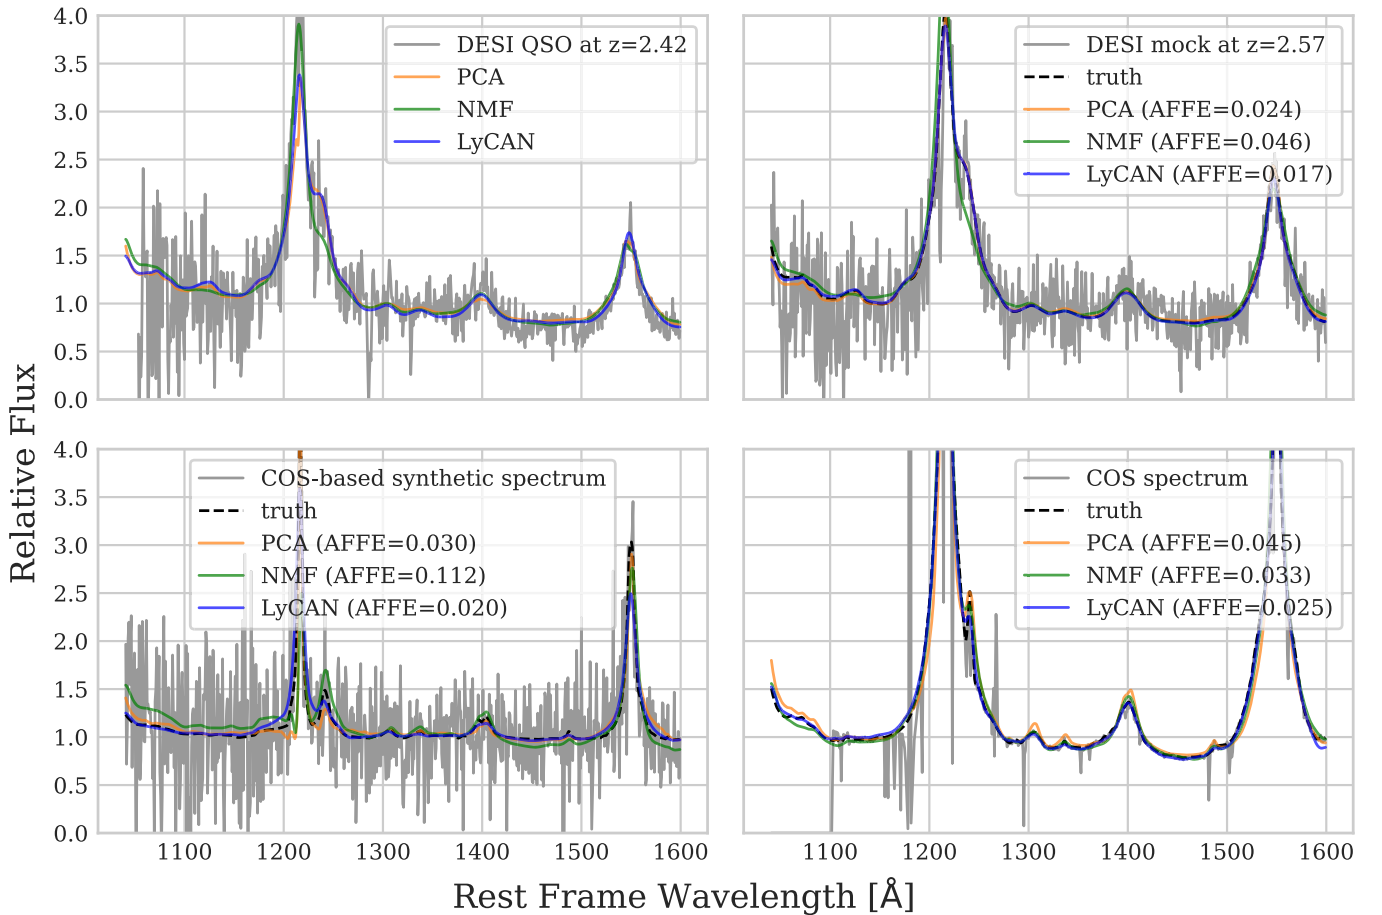


Figure 9. Examples of PCA, NMF, and LyCAN continuum predictions for four different classes of spectra, i.e., a real DESI Y1 spectrum, a DESI Y5 mock spectrum, a COS-based synthetic spectrum, and a real COS spectrum. None of these spectra appear in our training set.

continuum remains consistent. We see a smaller but similar effect in the dispersions of the higher-redshift predictions. The shape of the predictions remains consistent in each SNR and redshift range.

Similarly, we investigated the dependence of both FFE and AFFE for DESI Y5 mock continuum predictions on the redshift and SNR of the quasar. This provides an estimate of any bias in LyCAN that may produce nonphysical trends with redshift and/or SNR. Figure 12 shows these systematic biases. LyCAN performs more accurately on low to intermediate redshift, higher SNR objects. Based on the FFE, LyCAN slightly overpredicts the continuum level at higher redshifts by $\sim 1\%$. We correct for this effect in Section 5.

5. Evolution of the Effective Optical Depth

In this section, we use the DESI Y1 quasar continuum predictions from LyCAN to measure the evolution of the effective optical depth of the Ly α forest. We detail our optical depth calculation in Section 5.1. This calculation is a measurement of the total optical depth in the forest region, which includes contributions from metals and optically thick absorbers, in addition to the Ly α forest. In Section 5.2, we describe how we correct for these contributions, as well as quantify the potential bias in our continuum prediction, to determine the optical depth due to the Ly α forest. Finally, we compare our DESI Y1 results to previous results from the literature in Section 5.3.

5.1. Calculation

LyCAN enables an absolute measurement of the effective optical depth because it does not use any spectral information in the forest region, and we make this measurement with the largest sample to date. We do not use every DESI quasar for our measurement because we require a sufficiently high SNR to accurately measure the absorption present in the forest region. We use DESI Y1 quasars with median r -band SNR ≥ 5 and remove objects with BAL features, but retain DLAs because we later correct our measurements for optically thick absorbers. Our choice of SNR cut and BAL removal is consistent with previous works (e.g., V. Kamble et al. 2020). Among the full sample of $\sim 450,000$ $z \geq 2.1$ DESI Y1 quasars, roughly 90,000 contain BALs and $\sim 14,000$ have redshift warnings. The largest cut is due to the minimum SNR requirement. Our final sample contains 83,635 DESI Y1 quasars.

For each quasar, we measure the optical depth in redshift bins of width $\Delta z_{\text{Ly}\alpha} = 0.1$ between $2.0 \leq z_{\text{Ly}\alpha} \leq 4.2$, where $z_{\text{Ly}\alpha} \equiv \lambda_{\text{obs}}/\lambda_{\text{Ly}\alpha} - 1$ is the redshift of the Ly α pixel in the forest region. We restrict the analysis to the rest-frame wavelength range of 1070–1160 Å in each quasar as in V. Kamble et al. (2020) to mitigate contamination from the Ly α and O VI emission lines. In each redshift bin, we compute the 3σ -clipped mean transmission as

$$\langle F \rangle = \left\langle \frac{f(z_{\text{Ly}\alpha})}{C_q(z_{\text{Ly}\alpha})} \right\rangle \quad (6)$$

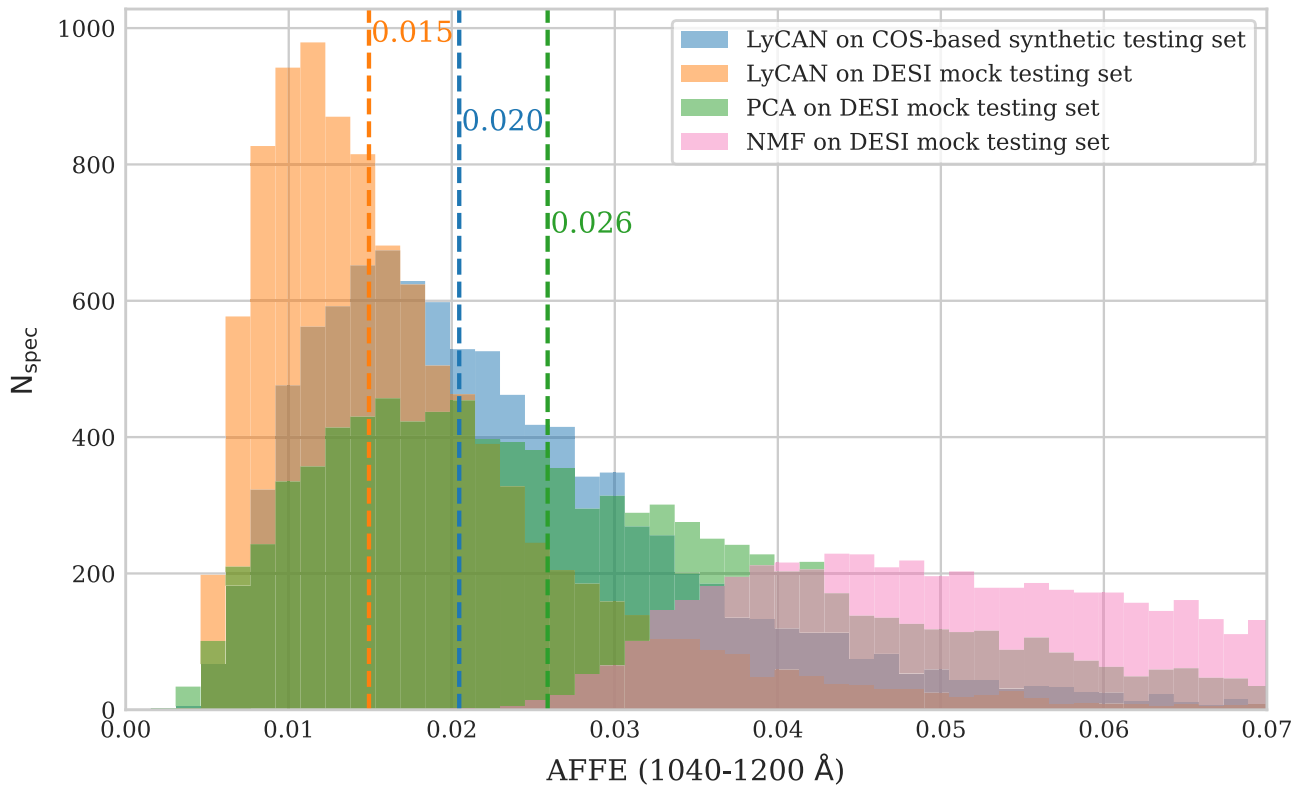


Figure 10. Distribution of AFFEs in the forest region for LyCAN, PCA, and NMF models on the DESI mock testing set. We also show the distribution of AFFEs for LyCAN predictions on our COS-based synthetic testing spectra. The dashed vertical lines indicate the median AFFE for each sample.

Table 2
A Comparison of the Median AFFEs for PCA, NMF, and LyCAN Prediction Models on Different Samples

Model	COS Sample	COS-based Synthetic Sample, Training	COS-based Synthetic Sample, Testing	DESI Mock Sample, Training	DESI Mock Sample, Testing
LyCAN	0.041	0.020	0.020	0.015	0.015
NMF	0.066	0.138	0.137	0.074	0.074
PCA	0.069	0.039	0.039	0.026	0.026

where f is the observed flux and C_q is the LyCAN prediction of the unabsorbed continuum. This $N_{\text{QSO}} \times N_z$ matrix is our measurement of the mean transmission per quasar in each redshift bin. We convert this to an optical depth matrix according to $\tau = -\ln\langle F \rangle$ and compute the 3σ -clipped mean τ_{eff} per redshift bin. These raw measurements are listed in Table 3.

We quantify our statistical uncertainties with 10,000 bootstrap realizations to compute the covariance matrix. In this calculation, we randomize the quasars in the sample and compute the 3σ -clipped mean optical depth per redshift bin across all objects in the sample. There is covariance in our measurements because most quasars contribute to the optical depth measurement at several redshifts. The correlation matrix is shown in Figure 13 and the diagonal of the covariance matrix σ_{stat} is listed in Table 3.

We repeat this analysis with the DESI Y5 mock spectra, once with the LyCAN continuum predictions and once with the true continua from the mocks. The correlation matrices for each of these data sets are also shown in Figure 13. The mock quasars do not span the full redshift range of the data, although the binning in redshift is the same. While the measurements on DESI Y1 data show correlations over several bins that is

consistent with the redshift range spanned by individual forests, there are only some modest correlations in the immediately adjacent redshift bin for both of the analyses on mocks. This is likely because LyCAN performs better on DESI mocks than DESI data.

5.2. Corrections and Systematic Error

Our measurement of the total optical depth described in Section 5.1 includes other absorption features in addition to Ly α , most notably metal lines and optically thick absorbers. There are also small biases in the continuum predictions, which are illustrated in Figure 12. We correct the total optical depth measurement for both the continuum biases and other absorption features to determine the evolution of the optical depth of the Ly α forest. We first use the DESI Y5 mock data set to calibrate our corrections for the redshift bias in the LyCAN continuum predictions. Figure 14 shows our effective optical depth measurements for the mock data set based on the LyCAN predicted continua and the mock truth continua. We attribute the difference to the continuum bias, which may be due to the systematic selection of higher luminosity quasars at higher redshifts. The top right-hand panel in Figure 12 shows that LyCAN tends to slightly underpredict the true continuum




Figure 11. The dispersions of continuum predictions for DESI Y1 quasars. The top panels show three ranges of SNR, while the bottom panels show three ranges of redshift. The solid line in each panel corresponds to the median prediction for that SNR or redshift range, and the shaded regions indicate the 1 and 2σ regions.




Figure 12. Systematic biases of LyCAN predictions on DESI Y5 mock spectra. The top panels show the FFE as functions of SNR and z , while the bottom panels show the AFPE (see Equation (3)) as functions of SNR and z . The solid line indicates the median (A)FFE, and the shaded regions show the 1 and 2σ regions. LyCAN tends to perform best on higher SNR, low- z spectra, and is more likely to overpredict the true continuum at higher z . This is a small effect ($\sim 2\%$ at highest z).




Figure 13. Correlation matrices for optical depth measurements in DESI Y1 with LyCAN continuum predictions, Y5 mock spectra with LyCAN predictions, and the same sample of Y5 mock spectra with their true continua. These matrices do not include corrections for uncertainties due to continuum bias, metals, or optically thick absorbers.

Table 3
Effective Optical Depth Measurements τ_{eff} and Associated Uncertainties

z_c	Raw ^a	Bias-corrected ^b	Final ^c	$\sigma_{\tau_{\text{eff}}, \text{stat}}$ ^d	$\sigma_{\tau_{\text{eff}}, \text{tot}}$ ^e
2.05	0.159	0.171	0.147	0.001	0.012
2.15	0.174	0.183	0.158	0.001	0.012
2.25	0.201	0.210	0.179	0.001	0.015
2.35	0.224	0.232	0.200	0.001	0.016
2.45	0.251	0.258	0.226	0.001	0.016
2.55	0.267	0.272	0.235	0.001	0.018
2.65	0.305	0.307	0.268	0.002	0.019
2.75	0.333	0.333	0.292	0.002	0.020
2.85	0.363	0.360	0.316	0.003	0.021
2.95	0.394	0.388	0.342	0.003	0.022
3.05	0.431	0.423	0.373	0.003	0.023
3.15	0.468	0.459	0.410	0.003	0.023
3.25	0.515	0.505	0.455	0.004	0.022
3.35	0.562	0.552	0.498	0.005	0.025
3.45	0.599	0.590	0.527	0.006	0.030
3.55	0.654	0.646	0.579	0.007	0.032
3.65	0.713	0.706	0.638	0.008	0.031
3.75	0.771	0.764	0.694	0.009	0.032
3.85	0.848	0.841	0.770	0.010	0.033
3.95	0.908	0.901	0.830	0.016	0.034
4.05	0.935	0.927	0.854	0.019	0.036
4.15	1.011	1.003	0.928	0.023	0.039

Notes. z_c refers to the redshift at the center of the bin.

^a Raw τ_{eff} measured directly from LyCAN continuum predictions.

^b Raw measurement corrected for redshift-dependent continuum bias in LyCAN.

^c Final measurement of τ_{eff} : bias-corrected measurement corrected for metal line absorption according to J. Schaye et al. (2003) and optically thick absorbers according to G. D. Becker et al. (2013).

^d Statistical uncertainty on τ_{eff} estimated from our covariance matrix.

^e Total statistical and systematic uncertainty (Section 5.2).

at lower redshifts and slightly overpredict the continuum at higher redshifts. This effect is small ($\sim 1\%$ in FFE) and is consistent with what is shown in Figure 14—the underpredictions at lower redshifts lead to less measured absorption, while the overpredictions at high redshifts lead to greater measured absorption than truth. Since the mock quasar samples in Figure 14 are identical, the difference in continuum solely determines the observed discrepancy in measured optical depth.




Figure 14. Top panel: effective optical depth measurements for the DESI Y5 mock data set based on truth continua (gray points) or LyCAN continuum predictions (red crosses). These are raw measurements and are not corrected for metals or optically thick absorbers. The error bars represent the statistical uncertainties. Bottom panel: the residual between the true and predicted measurements with statistical error bars.

Because the FFE is a fractional error, it has an additive effect on the effective optical depth,

$$\tau_{\text{eff,measured}} = \tau_{\text{eff,true}} + \delta F, \quad (7)$$

where $\tau_{\text{eff,measured}}$ is the measured optical depth with LyCAN continuum predictions, $\tau_{\text{eff,true}}$ is the true optical depth, and δF is the FFE (i.e., the nonabsolute version of the AFFE; see Equation (3)). We therefore use Equation (7) with the 3σ -clipped mean FFE per $z_{\text{Ly}\alpha}$ bin to correct for this redshift bias and extrapolate to $z_{\text{Ly}\alpha} > 3.6$ from the highest redshift measurements in our mock sample. We also compute the uncertainty on the mean FFE, σ_{bias} , to quantify the uncertainty




Figure 15. Measurements of the effective optical depth and associated power-law fit with DESI Y1 data using LyCAN continuum predictions. We find the best agreement with measurements based on high-resolution spectra at lower redshifts ($z < 3$) where continuum fitting errors are minimal (C.-A. Faucher-Giguère et al. 2008; P. Gaikwad et al. 2021). Our total uncertainty error bars are dominated by the systematic error due to the metal correction scheme, which is not reflected in some other results in the literature (e.g., G. D. Becker et al. 2013; B. Liu & R. Bordoloi 2021) (see Section 5.2).

in this correction. Our bias-corrected τ_{eff} measurements are reported in Table 3. The uncertainties on this correction are negligible ($\sim 10^{-4}$) compared to other sources of uncertainty, and therefore not listed.

After correcting our DESI Y1 optical depth measurements for the LyCAN continuum bias, we apply corrections due to absorption from metals and optically thick absorbers. C.-A. Faucher-Giguère et al. (2008) presented additive corrections due to metal absorption lines per redshift bin using the J. Schaye et al. (2003) correction scheme. They also calculated a systematic error for this metal correction by examining the deviation between the J. Schaye et al. (2003) and D. Kirkman et al. (2005) metal correction schemes. We adopt the metallicity correction from J. Schaye et al. (2003) and the associated systematic errors $\sigma_{\tau_{\text{eff,metal}}}$ from C.-A. Faucher-Giguère et al. (2008). The systematic error due to the metal correction scheme dominates our total error budget.

G. D. Becker et al. (2013) presented a functional form for the flux decrement due to optically thick absorbers ($N_{\text{HI}} \geq 10^{17.2} \text{ cm}^{-2}$) that is dominated by super Lyman limit systems (SLLSs; $10^{19.0} \leq N_{\text{HI}} < 10^{20.3} \text{ cm}^{-2}$) and DLAs ($N_{\text{HI}} \geq 10^{20.3} \text{ cm}^{-2}$). They found that the flux decrement is equal to $F_T = 1 - 0.0045[(1+z)/3]^{3.0}$. We apply this to our effective optical depth measurements per redshift bin by subtracting $\tau_T = -\ln(F_T)$ from our metal-corrected measurements. The authors did not report an associated uncertainty for this correction, so this is not included in our total error budget.

Our total uncertainties include statistical and systematic contributions, which we add in quadrature. These are $\sigma_{\tau_{\text{eff,tot}}} = \sqrt{\sigma_{\tau_{\text{eff,stat}}}^2 + \sigma_{\tau_{\text{eff,bias}}}^2 + \sigma_{\tau_{\text{eff,metal}}}^2}$. As stated earlier, $\sigma_{\tau_{\text{eff,bias}}}$ has a negligible effect on the total error budget, while $\sigma_{\tau_{\text{eff,metal}}}$

dominates it. The $\sigma_{\tau_{\text{eff,tot}}}$ values per redshift bin are given in Table 3.

We fit our measurements with a power law of the form $\tau(z) = \tau_0(1+z)^\gamma$. While we used the raw covariance matrix to quantify our statistical uncertainties (see Section 5.1), we computed a corrected covariance matrix to use as an input to the fitting procedure. This corrected covariance matrix was generated via bootstrapping as in Section 5.1 from a corrected per-quasar optical depth matrix. The corrections we used are not unique to quasars and only change with redshift bin. We find $\tau_0 = (2.46 \pm 0.14) \times 10^{-3}$ and $\gamma = 3.62 \pm 0.04$. We report our bias-corrected and fully corrected measurements alongside the raw measurements in Table 3. Our fully corrected measurements with total uncertainty error bars $\sigma_{\tau_{\text{eff,tot}}}$ and the associated power-law fit are shown in Figure 15.

5.3. Comparison with the Literature

Previous measurements of the optical depth have either used a relatively small number of high-resolution, high-SNR spectra or a large number of more moderate resolution spectra with generally lower SNR. The high-resolution studies offer the promise of a direct measurement of the continuum between the absorption features, particularly at lower redshifts, while large samples of moderate resolution spectra offer greater statistical power. C.-A. Faucher-Giguère et al. (2008) performed a direct measurement using 86 $z > 2$ high-resolution quasar spectra. In contrast, G. D. Becker et al. (2013) used 6065 moderate resolution SDSS DR7 spectra for a relative measurement of the mean transmitted flux as a fraction of its value at $z \leq 2.5$, with the assumption that the mean unabsorbed quasar continuum does not evolve with redshift. They used results from C.-A. Faucher-Giguère et al. (2008) to convert this to an

absolute measurement. V. Kamble et al. (2020) used 40,035 SDSS DR12 quasar spectra to perform a similar relative measurement while also allowing for spectral diversity. They found a higher opacity at $z \lesssim 3$ compared to G. D. Becker et al. (2013). More recently, B. Liu & R. Bordoloi (2021) used their feedforward neural network-based continuum predictions on ~ 3200 SDSS DR16 quasar spectra to perform an absolute measurement of the effective optical depth. P. Gaikwad et al. (2021) used a sample of 103 high-resolution, high-SNR quasar spectra from KODIAQ DR2 to directly measure the evolution of the mean flux, and found results in agreement with C.-A. Faucher-Giguère et al. (2008). Additionally, J. Ding et al. (2024) presented a new measurement of the effective Ly α forest opacity based on 27,008 SDSS DR14 quasar spectra using a Markov chain Monte Carlo analysis. The authors found a systematic offset toward lower opacity values compared to V. Kamble et al. (2020).

Some previous works reported evidence for a decrease in the effective optical depth at $z \simeq 3.2$. Such a feature could be a signature of He II reionization because it could signal a rise in the IGM gas temperature at the end of helium reionization (see e.g., T. Theuns et al. 2002; C.-A. Faucher-Giguère et al. 2008; G. D. Becker et al. 2013). M. Bernardi et al. (2003) found a sudden $\sim 10\%$ reduction in the effective optical depth at this redshift, which was later corroborated by C.-A. Faucher-Giguère et al. (2008) despite substantial differences in their data sets and analysis methods. However, others have argued from a theoretical perspective that such a sharp feature should not be observed (e.g., J. S. Bolton et al. 2009). Some temperature evolution measurements (e.g., G. D. Becker et al. 2011; M. Walther et al. 2019) also suggest that IGM heating due to He II reionization occurred over a more prolonged period, which would not produce a sharp feature in the evolution of the effective optical depth. However, more recent temperature measurements from P. Gaikwad et al. (2021) provide evidence for a more rapid He II reionization. Additional measurements of the mean transmitted flux by G. D. Becker et al. (2013) revealed no evidence for a departure from a smooth power-law description. Later works (e.g., V. Kamble et al. 2020; B. Liu & R. Bordoloi 2021; J. Ding et al. 2024) also reported no evidence for this $z \simeq 3.2$ feature.

Our results for the evolution of the mean optical depth show a smooth power-law relationship of the form $\tau_{\text{eff}}(z) = \tau_0(1+z)^\gamma$, where $\tau_0 = (2.46 \pm 0.14) \times 10^{-3}$ and $\gamma = 3.62 \pm 0.04$, and no evidence for a dip at $z \simeq 3.2$. Our approach should be sensitive to this potential feature, as our systematic errors vary smoothly with redshift. We find the closest agreement with measurements based on high-resolution spectra (e.g., C.-A. Faucher-Giguère et al. 2008; P. Gaikwad et al. 2021), particularly at the lower redshift end. High-resolution spectra are subject to smaller errors in the continuum fitting process, and continuum errors are minimal at lower redshifts where Ly α forest absorption is reduced. This agreement indicates that LyCAN can accurately predict the underlying quasar continuum. We also measure a lower optical depth at higher redshifts ($z \gtrsim 3.5$) compared to previous works, which may indicate a stronger ionizing background at these redshifts. Furthermore, we find good agreement with B. Liu & R. Bordoloi (2021), whose measurements were also corrected for metals according to J. Schaye et al. (2003), over our whole redshift range. This points to the self-consistency of deep learning methods for quasar continuum prediction. Our

measurements are shown alongside others in Figure 15. The error bars for the G. D. Becker et al. (2013) and B. Liu & R. Bordoloi (2021) measurements are likely underestimated because they do not include the systematic error in the correction for metal line absorption.

6. Summary

We have developed LyCAN, a CNN-based approach to predict the underlying quasar continuum in the Ly α forest region strictly based on quasar spectra on the red side of the Ly α emission line. We have shown that LyCAN performs extremely well on all testing data and outperforms the PCA- and NMF-based predictive models developed for this work using the same training data. LyCAN also uses the flux on the red side of Ly α as input without any smoothing or additional preprocessing steps, while the PCA- and NMF-based models are more sensitive to noise and the presence of other features in the quasar spectra, such as metal line absorption.

We began our analysis with a sample of 48 archival COS observations. We generated NMF components for this sample and fit these components to a large sample of high-SNR (median r -band SNR ≥ 5) DESI Y1 spectra. The NMF coefficients for 10 of the COS spectra were very inconsistent with the coefficient distribution for the DESI quasars, and we therefore eliminated them from the sample. We used the remaining set of 38 low-redshift COS spectra to generate a sample of 40,000 unique synthetic spectra with varying levels of noise and emission line properties modeled after COS and DESI Y1 observations. We supplemented these with a sample of 40,000 DESI Y5 mock spectra, and randomly divided the resulting data set of 80,000 spectra into training, testing, and validation subsets.

We found the optimal CNN architecture through a pseudo-random search of hyperparameters. We optimized the number of convolutional layers, number of filters, and kernel size per layer, as well as the number of fully connected layers and number of nodes per layer. Our final architecture was chosen as the option that minimized both the median AFFE score (Equation (4)) and 1σ range score (Equation (5)). The LyCAN architecture features six hidden layers, including three convolutional layers and three dense layers.

We used the LyCAN continuum predictions on DESI Y1 quasars to perform an absolute measurement of the evolution of the effective optical depth with the largest sample to date. Our results show close agreement with previous measurements based on high-resolution quasar spectra at lower redshifts (e.g., C.-A. Faucher-Giguère et al. 2008; P. Gaikwad et al. 2021) where those methods are most able to measure the quasar continuum between the absorption features. This demonstrates that LyCAN can accurately predict the unabsorbed continuum. We find that the mean optical depth varies smoothly with redshift and we do not find evidence for a feature at $z \simeq 3.2$ that could be associated with He II reionization. Due to our large sample size of 83,635 quasars and LyCAN's ability to accurately predict the underlying continuum, we expect that this is the most accurate measurement of the effective optical depth evolution to date.

Our analysis is limited by the quality of our training set and its application to DESI quasars. The median AFFE on synthetic data sets is likely to be more of a lower bound on the uncertainty associated with real DESI data. More UV spectra of high-luminosity, low-redshift quasars could substantially

improve the empirical training set. Additionally, our correction procedure for the observed redshift bias in LyCAN predictions assumes that the DESI mock spectra are representative of the real Y1 data. This is a common assumption in large surveys such as DESI because mocks are used to quantify sources of systematic error. Future studies of spectral diversity in the quasar sample would help to better quantify this source of systematic error. Furthermore, since LyCAN is a CNN, continuum predictions are deterministic and do not enable the quantification of uncertainties on real DESI data. Improvements on this, as well as direct comparison to other continuum prediction models in the literature, are left to future work.

LyCAN was primarily developed to be an alternative method to `picca` to measure the flux transmission field $\delta_q(\lambda)$ of the Ly α forest and measure the Ly α auto-correlation and Ly α -QSO cross correlation functions. We plan to use LyCAN continuum predictions on DESI Y1 quasars to compute the flux transmission field in the Ly α forest region in a future paper, and use these to compute cosmological parameters. This alternative approach will complement the standard analysis, especially because the measurements of $\delta_q(\lambda)$ with LyCAN will not have the same biases of the mean and spectral slope that distort the correlation function. This will help retain information on large scales for the one-dimensional (C. Ravoux et al. 2023; N. G. Karaçaylı et al. 2024) and three-dimensional (M. L. A. Karim et al. 2023; R. de Belsunce et al. 2024) power spectrum, and may be especially valuable for future work to measure the full shape of the auto-correlation function (A. Cuceu et al. 2023b, 2023a).

Acknowledgments

W.T. and P.M. acknowledge support from the United States Department of Energy, Office of High Energy Physics under Award Number DE-SC-0011726. We thank Joe Antognini, David Weinberg, and Chris Hirata for helpful discussions, and Nicolas Lehner for sharing information about COS observations of quasars.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions: <https://www.desi.lbl.gov/collaborating-institutions>. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

Data Availability

The data points for each figure are available at doi:[10.5281/zenodo.11106581](https://doi.org/10.5281/zenodo.11106581). The COS observations used for generation of synthetic spectra can be accessed via doi:[10.17909/ae35-e630](https://doi.org/10.17909/ae35-e630).

Facilities: HST (COS), Mayall (DESI).

Software: Numpy (C. R. Harris et al. 2020), Scipy (P. Virtanen et al. 2020), Matplotlib (J. D. Hunter 2007), Keras (F. Chollet et al. 2018), TensorFlow (M. Abadi et al. 2016), LineTools (J. X. Prochaska et al. 2017), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), Scikit-Learn (F. Pedregosa et al. 2011), NonnegMFPy (G. Zhu 2016), Spectres (A. C. Carnall 2017), Corner (D. Foreman-Mackey 2016).

ORCID iDs

Wynne Turner  <https://orcid.org/0009-0008-3418-5599>
 Paul Martini  <https://orcid.org/0000-0002-0194-4017>
 Naim Göksel Karaçaylı  <https://orcid.org/0000-0001-7336-8912>
 S. Ahlen  <https://orcid.org/0000-0001-6098-7247>
 D. Brooks  <https://orcid.org/0000-0002-8458-5047>
 A. de la Macorra  <https://orcid.org/0000-0002-1769-1640>
 A. Dey  <https://orcid.org/0000-0002-4928-4003>
 K. Fanning  <https://orcid.org/0000-0003-2371-3356>
 J. E. Forero-Romero  <https://orcid.org/0000-0002-2890-3725>
 S. Gontcho A Gontcho  <https://orcid.org/0000-0003-3142-233X>
 A. X. Gonzalez-Morales  <https://orcid.org/0000-0003-4089-6924>
 J. Guy  <https://orcid.org/0000-0001-9822-6793>
 H. K. Herrera-Alcantar  <https://orcid.org/0000-0002-9136-9609>
 S. Juneau  <https://orcid.org/0000-0002-0000-2394>
 T. Kisner  <https://orcid.org/0000-0003-3510-7134>
 A. Kremin  <https://orcid.org/0000-0001-6356-7424>
 M. Landriau  <https://orcid.org/0000-0003-1838-8528>
 L. Le Guillou  <https://orcid.org/0000-0001-7178-8868>
 A. Meisner  <https://orcid.org/0000-0002-1125-7384>
 R. Miquel  <https://orcid.org/0000-0002-6610-4836>
 J. Moustakas  <https://orcid.org/0000-0002-2733-4559>
 J. Nie  <https://orcid.org/0000-0001-6590-8122>
 G. Niz  <https://orcid.org/0000-0002-1544-8946>
 C. Poppett  <https://orcid.org/0000-0003-0512-5489>
 F. Prada  <https://orcid.org/0000-0001-7145-8674>
 M. Rezaie  <https://orcid.org/0000-0001-5589-7116>
 E. Sanchez  <https://orcid.org/0000-0002-9646-8198>
 E. F. Schlafly  <https://orcid.org/0000-0002-3569-7421>
 D. Schlegel  <https://orcid.org/0000-0002-5042-5088>
 H. Seo  <https://orcid.org/0000-0002-6588-3508>
 D. Sprayberry  <https://orcid.org/0000-0001-7583-6441>
 G. Tarlé  <https://orcid.org/0000-0003-1704-0781>
 H. Zou  <https://orcid.org/0000-0002-6684-3997>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, arXiv:1603.04467
 Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, *ApJS*, 249, 3
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Baldwin, J. A. 1977, *ApJ*, **214**, 679
- Bautista, J. E., Busca, N. G., Guy, J., et al. 2017, *A&A*, **603**, A12
- Becker, G. D., Bolton, J. S., Haehnelt, M. G., & Sargent, W. L. W. 2011, *MNRAS*, **410**, 1096
- Becker, G. D., Bolton, J. S., Madau, P., et al. 2015, *MNRAS*, **447**, 3402
- Becker, G. D., Hewett, P. C., Worseck, G., & Prochaska, J. X. 2013, *MNRAS*, **430**, 2067
- Bernardi, M., Sheth, R. K., SubbaRao, M., et al. 2003, *AJ*, **125**, 32
- Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, *AJ*, **154**, 28
- Blomqvist, M., du Mas des Bourboux, H., Busca, N. G., et al. 2019, *A&A*, **629**, A86
- Bolton, J. S., Haehnelt, M. G., Viel, M., & Springel, V. 2005, *MNRAS*, **357**, 1178
- Bolton, J. S., Oh, S. P., & Furlanetto, S. R. 2009, *MNRAS*, **396**, 2405
- Busca, N. G., Delubac, T., Rich, J., et al. 2013, *A&A*, **552**, A96
- Carnall, A. C. 2017, arXiv:1705.05165
- Chollet, F., et al., 2018 Keras: The Python Deep Learning library, Astrophysics Source Code Library, ascl:1806.022
- Cole, S., Percival, W. J., Peacock, J. A., et al. 2005, *MNRAS*, **362**, 505
- Cuceu, A., Font-Ribera, A., Martini, P., et al. 2023a, *MNRAS*, **523**, 3773
- Cuceu, A., Font-Ribera, A., Nadathur, S., Joachimi, B., & Martini, P. 2023b, *PhRvL*, **130**, 191003
- Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, *ApJ*, **864**, 143
- Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, *AJ*, **151**, 44
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
- de Belsunce, R., Philcox, O. H. E., Irsic, V., et al. 2024, *MNRAS*, **533**, 3756
- de Sainte Agathe, V., Balland, C., du Mas des Bourboux, H., et al. 2019, *A&A*, **629**, A85
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2016a, arXiv:1611.00036
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2022, *AJ*, **164**, 207
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024a, *AJ*, **167**, 62
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024b, *AJ*, **168**, 58
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024c, arXiv:2404.03000
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024d, arXiv:2404.03001
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024e, arXiv:2404.03002
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016b, arXiv:1611.00037
- Ding, J., Madau, P., & Prochaska, J. X. 2024, *MNRAS*, **532**, 2082
- du Mas des Bourboux, H., Rich, J., Font-Ribera, A., et al. 2020, *ApJ*, **901**, 153
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, **142**, 72
- Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. 2005, *ApJ*, **633**, 560
- Farr, J., Font-Ribera, A., Bourboux, H. D. M. D., et al. 2020, *JCAP*, **2020**, 068
- Faucher-Giguere, C.-A. 2020, *MNRAS*, **493**, 1614
- Faucher-Giguere, C.-A., Prochaska, J. X., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, **681**, 831
- Font-Ribera, A., Kirkby, D., Busca, N., et al. 2014, *JCAP*, **2014**, 027
- Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
- Gaikwad, P., Srianand, R., Haehnelt, M. G., & Choudhury, T. R. 2021, *MNRAS*, **506**, 4389
- Gordon, K. D., Clayton, G. C., Decleir, M., et al. 2023, *ApJ*, **950**, 86
- Greig, B., Mesinger, A., McGreer, I. D., Gallerani, S., & Haiman, Z. 2017, *MNRAS*, **466**, 1814
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Herrera-Alcántar, H. K., Muñoz-Gutiérrez, A., Tan, T., et al. 2023, arXiv:2401.00303
- Huang, L., Croft, R. A. C., & Arora, H. 2021, *MNRAS*, **506**, 5212
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jalan, P., Khaire, V., Vivek, M., & Gaikwad, P. 2024, *A&A*, **688**, A126
- Kamble, V., Dawson, K., du Mas des Bourboux, H., Bautista, J., & Scheinder, D. P. 2020, *ApJ*, **892**, 70
- Karaçaylı, N. G., Font-Ribera, A., & Padmanabhan, N. 2020, *MNRAS*, **497**, 4742
- Karaçaylı, N. G., Martini, P., Guy, J., et al. 2024, *MNRAS*, **528**, 3941
- Karaçaylı, N. G., Padmanabhan, N., Font-Ribera, A., et al. 2022, *MNRAS*, **509**, 2842
- Karim, M. L. A., Armengaud, E., Mention, G., et al. 2023, *JCAP*, **2024**, 088
- Khaire, V., & Srianand, R. 2019, *MNRAS*, **484**, 4174
- Khaire, V., Srianand, R., Choudhury, T. R., & Gaikwad, P. 2016, *MNRAS*, **457**, 4051
- Khaire, V., Walther, M., Hennawi, J. F., et al. 2019, *MNRAS*, **486**, 769
- Kingma, D. P., & Ba, J. 2017, arXiv:1412.6980
- Kirkman, D., Tytler, D., Suzuki, N., et al. 2005, *MNRAS*, **360**, 1373
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, in *Advances in Neural Information Processing Systems*, ed. F. Pereira et al., 25 (Curran Associates, Inc.), 1097
- Kulkarni, G., Keating, L. C., Haehnelt, M. G., et al. 2019, *MNRAS*, **485**, L24
- Lee, K.-G., Suzuki, N., & Spergel, D. N. 2012, *AJ*, **143**, 51
- Liu, B., & Bordoloi, R. 2021, *MNRAS*, **502**, 3510
- McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2001, *ApJ*, **562**, 52
- McQuinn, M. 2016, *ARA&A*, **54**, 313
- Meiksin, A. A. 2009, *RvMP*, **81**, 1405
- Nasir, F., Gaikwad, P., Davies, F. B., et al. 2024, *MNRAS*, **534**, 1299
- Nayak, P., Walther, M., Gruen, D., & Adiraju, S. 2024, *A&A*, **689**, A153
- Pâris, I., Petitjean, P., Rollinde, E., et al. 2011, *A&A*, **530**, A50
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Peeples, M., Tumlinson, J., Fox, A., et al. 2017, *The Hubble Spectroscopic Legacy Archive*, Tech. Instrument Science Report, *COS 2017-4*
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *ApJ*, **517**, 565
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13
- Prochaska, J. X., Tejos, N., Crighton, N., et al. 2017, *linetools/linetools: Third Minor Release, v0.3*, Zenodo, doi:10.5281/zenodo.1036773
- Puchwein, E., Haardt, F., Haehnelt, M. G., & Madau, P. 2019, *MNRAS*, **485**, 47
- Ramírez-Pérez, C., Pérez-Ràfols, I., Font-Ribera, A., et al. 2024, *MNRAS*, **528**, 6666
- Rauch, M., Miralda-Escudé, J., Sargent, W. L. W., et al. 1997, *ApJ*, **489**, 7
- Ravoux, C., Abdul Karim, M. L., Armengaud, E., et al. 2023, *MNRAS*, **526**, 5118
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, **116**, 1009
- Schaye, J., Aguirre, A., Kim, T., et al. 2003, *ApJ*, **596**, 768
- Schaye, J., Theuns, T., Rauch, M., Efstathiou, G., & Sargent, W. L. W. 2000, *MNRAS*, **318**, 817
- Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, **737**, 103
- Slosar, A., Iršič, V., Kirkby, D., et al. 2013, *JCAP*, **2013**, 026
- Sun, Z., Ting, Y.-S., & Cai, Z. 2023, *ApJS*, **269**, 4
- Suzuki, N., Tytler, D., Kirkman, D., O'Meara, J. M., & Lubin, D. 2005, *ApJ*, **618**, 592
- Theuns, T., Bernardi, M., Frieman, J., et al. 2002, *ApJ*, **574**, L111
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Walther, M., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2019, *ApJ*, **872**, 13
- Weinberg, D. H., Mortonson, M. J., Eisenstein, D. J., et al. 2013, *PhR*, **530**, 87
- Zaldarriaga, M., Hui, L., & Tegmark, M. 2001, *ApJ*, **557**, 519
- Zhu, G. 2016, arXiv:1612.06037