



A Detection of Cosmological 21 cm Emission from CHIME in Cross-correlation with eBOSS Measurements of the Ly α Forest

Mandana Amiri¹, Kevin Bandura^{2,3}, Arnab Chakraborty^{4,5}, Matt Dobbs^{4,5}, Mateus Fandino⁶, Simon Foreman⁷, Hyoyin Gan¹, Mark Halpern¹, Alex S. Hill^{8,9}, Gary Hinshaw¹, Carolin Höfer^{1,10}, T. L. Landecker⁹, Zack Li¹¹, Joshua MacEachern¹, Kiyoshi Masui^{12,13}, Juan Mena-Parra^{14,15}, Nikola Milutinovic¹, Arash Mirhosseini¹, Laura Newburgh¹⁶, Anna Ordog^{8,9}, Sourabh Paul^{4,5}, Ue-Li Pen¹¹, Tristan Pinsonneault-Marotte¹, Alex Reda¹⁶, J. Richard Shaw¹, Seth R. Siegel^{4,5,17}, Keith Vanderlinde^{14,15}, Haochen Wang^{12,13}, D. V. Wiebe¹, and Dallas Wulf^{4,5}

The CHIME Collaboration

¹ Department of Physics and Astronomy, University of British Columbia, Vancouver, BC, Canada; tristpinsm@phas.ubc.ca

² Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, WV, USA

³ Center for Gravitational Waves and Cosmology, West Virginia University, Morgantown, WV, USA

⁴ Department of Physics, McGill University, Montreal, QC, Canada

⁵ Trottier Space Institute, McGill University, Montreal, QC, Canada

⁶ Department of Physical Sciences, Thompson Rivers University, Kamloops, BC, Canada

⁷ Department of Physics, Arizona State University, Tempe, AZ, USA

⁸ Department of Computer Science, Math, Physics, and Statistics, University of British Columbia-Okanagan, Kelowna, BC, Canada

⁹ Dominion Radio Astrophysical Observatory, Herzberg Astronomy & Astrophysics Research Centre, National Research Council Canada, Penticton, BC, Canada

¹⁰ Kapteyn Astronomical Institute, University of Groningen, Groningen, The Netherlands

¹¹ Canadian Institute for Theoretical Astrophysics, Toronto, ON, Canada

¹² MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA, USA

¹³ Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA

¹⁴ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto, ON, Canada

¹⁵ David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON, Canada

¹⁶ Department of Physics, Yale University, New Haven, CT, USA

¹⁷ Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada

Received 2023 September 7; revised 2023 November 9; accepted 2023 November 21; published 2024 February 23

Abstract

We report the detection of 21 cm emission at an average redshift $\bar{z} = 2.3$ in the cross-correlation of data from the Canadian Hydrogen Intensity Mapping Experiment (CHIME) with measurements of the Ly α forest from eBOSS. Data collected by CHIME over 88 days in the 400–500 MHz frequency band ($1.8 < z < 2.5$) are formed into maps of the sky and high-pass delay filtered to suppress the foreground power, corresponding to removing cosmological scales with $k_{\parallel} \lesssim 0.13 \text{ Mpc}^{-1}$ at the average redshift. Line-of-sight spectra to the eBOSS background quasar locations are extracted from the CHIME maps and combined with the Ly α forest flux transmission spectra to estimate the 21 cm–Ly α cross-correlation function. Fitting a simulation-derived template function to this measurement results in a 9σ detection significance. The coherent accumulation of the signal through cross-correlation is sufficient to enable a detection despite excess variance from foreground residuals ~ 6 –10 times brighter than the expected thermal noise level in the correlation function. These results are the highest-redshift measurement of 21 cm emission to date, and they set the stage for future 21 cm intensity mapping analyses at $z > 1.8$.

Unified Astronomy Thesaurus concepts: Cosmology (343); H I line emission (690); Ly α forest (980)

1. Introduction

Emission from the hyperfine transition of neutral hydrogen (H I) at 21 cm can be used to efficiently map the large-scale structure (LSS) of the Universe over most of its history, a technique known as hydrogen intensity mapping. In this approach, H I contained in galaxies or the intergalactic medium is detected in aggregate—integrated over the relatively coarse angular resolution afforded by telescopes at radio wavelengths, where redshifted 21 cm emission is measured. Such telescopes can be designed to cost-effectively observe large sky areas at high sensitivity, and digital receivers enable broad bandwidths to be sampled with fine frequency resolution—which directly

maps to redshift, due to the monochromaticity of the signal (Peterson et al. 2009).

Over the last decade, instruments have been built with the aim of measuring this signal in the late Universe (e.g., CHIME (CHIME Collaboration et al. 2022) and MeerKAT (Jonas 2009)) as well as at higher redshift (e.g., HERA (DeBoer et al. 2017), EDGES (Bowman et al. 2008), SARAS (Patra et al. 2013), and PAPER (Parsons et al. 2010), among others). Published results from the field include upper limits (Abdurashidova et al. 2022), detections in cross-correlation (Pen et al. 2009; Chang et al. 2010; Masui et al. 2013; Anderson et al. 2018; Tramonte & Ma 2020; Li et al. 2021; CHIME Collaboration et al. 2023; Wolz et al. 2022; Cunningham et al. 2023), and a first detection in autocorrelation (Paul et al. 2023). The main challenge for all of these efforts is separating the faint 21 cm signal from the extremely bright foreground emission of our galaxy and radio point sources, which can in principle be achieved by making use of their different spectral properties, but is complicated by



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.




Figure 1. Distribution of eBOSS Ly α forest spectral samples across all quasar lines of sight, binned by their corresponding redshift. The dashed region shows their overlap with the redshift range probed by CHIME.

instrumental effects and radio-frequency interference (RFI). Until foreground separation methods mature to a point where the signal becomes dominant, cross-correlation with external surveys has proven to be an effective way of mitigating residual foregrounds and obtaining interesting scientific constraints, as demonstrated by the recently reported detection from CHIME in cross-correlation with SDSS galaxy surveys (CHIME Collaboration et al. 2023).

The Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME Collaboration et al. 2022) is a compact interferometer composed of four cylindrical reflectors instrumented with a total of 1024 dual-polarization antennas. It is located at the Dominion Radio Astrophysical Observatory (DRAO) in Penticton, Canada, where it observes the entire sky visible at latitude $\sim 49^\circ$, operating as a driftscan telescope over 1024 channels in the 400–800 MHz band (corresponding to 21 cm emission redshifted by $2.5 > z > 0.8$). CHIME has been operating continuously since 2017.

The redshift range available to CHIME overlaps with measurements of the Ly α forest from eBOSS (du Mas des Bourboux et al. 2020), allowing for a cross-correlation analysis to be carried out in the high-redshift end of the CHIME band. At these redshifts, the quasar catalog that was stacked on in CHIME Collaboration et al. (2023) is sparser, and there is more statistical weight in the Ly α data. Ly α forest measurements are characterized by absorption of quasar light by H I clouds along the line of sight, whereas the 21 cm signal is in emission, so we expect the cross-correlation with CHIME maps to be *negative* at small separations, which provides a unique signal that is difficult to mimic any other way. This is a feature that has been highlighted by Carucci et al. (2017), who also emphasize the usefulness of such a cross-correlation to mitigate foreground contamination and break degeneracies in a power spectrum analysis.

Several studies have found evidence that various classes of galaxies are correlated with Ly α absorption by the IGM up to several tens of megaparsecs (Mukae et al. 2017; Liang et al. 2021; Momose et al. 2021), but these studies have been limited to small numbers of galaxies and/or absorption systems. A joint analysis of CHIME and eBOSS Ly α forest data will provide a huge statistical improvement on both sides, opening a new window into the connection between H I-rich galaxies and low-density H I in the surrounding IGM, and thus shedding light on the role of H I in galaxy evolution (e.g., the strength of H I inflows from the IGM onto denser systems) when the cosmic star formation rate is at its peak.

In this work, we present a detection of 21 cm emission in the cross-correlation of CHIME data and eBOSS Ly α forest measurements, but leave the interpretation of the signal (including modeling of its cosmological and astrophysical implications) for future work. We describe the data used in this analysis in Section 2 and the CHIME processing pipeline in Section 3. Section 4 explains the cross-correlation method we use. A model for the signal and its inclusion in simulations are described in Section 5. The results of the analysis are presented in Section 6, and some validation tests in Section 7. Finally, we conclude in Section 8.

2. Data

2.1. eBOSS

The Ly α forest measurements we use in our analysis are those from the eBOSS sixteenth data release (SDSS DR16; du Mas des Bourboux et al. 2020). The data products consist of fractional flux transmission along the line of sight to backlight quasars, defined as

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

where f_q is the measured flux and C_q is the unabsorbed quasar continuum for a given quasar, labeled q . Here, $\bar{F}$ is the mean transmission across all quasars. The product $\bar{F}(\lambda)C_q(\lambda)$ represents the mean expected flux that would be observed in the absence of fluctuations in the optical depth to Ly α photons, and was modeled and fit to the data as explained in du Mas des Bourboux et al. (2020). Apart from interpolating onto the CHIME frequency band (Section 4.1), no additional processing is performed on the catalog provided by eBOSS.

There are 210,005 background quasars in the eBOSS Ly α catalog, with redshifts $2.1 < z < 4$. The CHIME data go down to 400 MHz, corresponding to 21 cm radiation redshifted at $z = 2.5$, which overlaps with a total of $\sim 3 \times 10^7$ spectral samples in the eBOSS measurements across all lines of sight (see Figure 1). The eBOSS samples for a given line-of-sight spectrum are uniformly binned in log-wavelength, with channel widths ~ 0.25 – 0.4 nm. We include the full eBOSS spectra in this total and the analysis, including the regions blueward of the Lyman- β wavelength (~ 102 nm) at the quasar redshift (referred to as the Ly β region in the eBOSS papers). This part of the spectrum is treated separately in their analysis because it can include absorption from both transition levels, but because we are cross-correlating with an external data set, there is no concern of contamination from Ly β correlations. Damped Ly α systems are identified and masked in the eBOSS data products, so they do not contribute to the cross-correlation with CHIME. The quasars are distributed throughout the SDSS North and South Galactic Cap fields (NGC and SGC), which are fully contained within the sky area observed by CHIME.

2.2. CHIME

The CHIME data that was used in this work are composed of nighttime observations from 88 days spanning the calendar year 2019. This amounts to ~ 1000 hr of total observation time, distributed approximately uniformly over the sidereal day. The full frequency range observed with CHIME is 400–800 MHz, but we use only the lowest quarter of the band, i.e., 400–500 MHz or $z > 1.8$, which has overlap with the Ly α

forest measurements. This is a similar set of observations to what was used in CHIME Collaboration et al. (2023) to detect 21 cm emission in cross-correlation with SDSS galaxy catalogs (refer to that paper and citations therein for a general description of the CHIME data). However, that analysis considered data from the upper part of the frequency band, 585–800 MHz, whereas this work is the first report on the high-redshift end.

3. Processing

The processing that the CHIME data undergoes prior to cross-correlation with the Ly α forest is largely the same as described in the first detection paper (Section 3 of CHIME Collaboration et al. 2023). In this section, we summarize the main steps and highlight improvements that were made to the pipeline since the publication of that analysis, as well as aspects of the processing that had to be modified to accommodate the lower frequency band.

3.1. Sidereal Stacking Pipeline

The final product of the CHIME processing pipeline is a set of visibilities for unique baselines measuring the sky on a grid of R.A. spanning the sidereal day. To produce this “sidereal stack,” observations from individual days are collected and passed through daily processing that includes RFI flagging, gain calibration and regridding from local time to a fixed grid in sidereal time. Each day is inspected visually via a set of data-quality metrics to identify and flag days that appear corrupted. The remaining days are averaged together to form the final sidereal stack that is used in this analysis. Further details on each stage of this processing can be found in CHIME Collaboration et al. (2023).

3.1.1. Improved Flagging

Since the publication of the first CHIME detection, some improvements to the processing pipeline have been made. These include:

Narrowband gain errors—The calibration algorithms employed in the real-time pipeline attempt to identify and mask RFI-like features in the underlying data prior to estimating the frequency- and feed-dependent gains. However, these algorithms are not entirely robust to RFI, and we have found that transient RFI occasionally biases the resulting gain estimates. This is true for both the “digital gains,” which are updated semi-annually and whose purpose is to minimize quantization noise when truncating the data streams to 4 bits (real) + 4 bits (imaginary), and the “calibration gains,” which are updated daily and whose purpose is to correct instability in the analog receiver chain. The offending RFI is narrowband, usually with a bandwidth less than a single 390.625 kHz frequency channel, and results in spikes in the estimated gain as a function of frequency. These spikes are imprinted on the foregrounds when the gains are applied to the visibility data, leaking foreground power to high delays (the conjugate axis to frequency, where smooth foregrounds would otherwise be confined to low delays).

We have developed a new method to search the archived gains for these narrowband artifacts and then mask the corresponding frequencies and times in the visibility data. For each time sample in the visibility data, we calculate the product of the digital gain and calibration gain that was applied

to each feed and frequency by the real-time pipeline, average the amplitude over feeds, and then perform a search for narrowband features along the frequency axis. The search algorithm applies an aggressive, high-pass filter along the frequency axis, masks the frequency channel that is the largest outlier, reapplies the filter accounting for the newly masked data, and iterates until all unmasked frequencies are less than 6σ , where σ is an estimate of the standard deviation of the noise. This procedure masked 6% of the data that were considered for this analysis.

Decorrelation events—On rare occasions, communication errors during the redistribution of data across the correlator can cause certain data streams to become misaligned with all other data streams. This results in one quarter of the feeds decorrelating with all other feeds for 64 of the 1024 frequency channels, which persists until the data streams are realigned by restarting the correlator. The 64 decorrelated frequency channels are uniformly spaced across the band and leak significant foreground power to high delays. We now perform an automated search for these decorrelation events and exclude any frequency channel and time that is affected from further analysis. In total, 0.07% of the data considered for this analysis were masked for this reason, corresponding to a single decorrelation event that persisted for the majority of one sidereal day.

Excessively small weights—The real-time RFI flagging excises corrupt samples within the time integrations that are eventually saved to disk; for integrations where this excision is nearly total, the data saved to disk can be extremely noisy. A similar effect can occur due to packet loss or other digital errors in the X-engine. These integrations are assigned very low noise weights in subsequent analysis, but if the issues are confined to very narrow frequency bands or small subsets of baselines, the low weights were not properly incorporated in previous versions of the processing. This has been corrected using an extra flagging step in the offline data pipeline, which flags 0.5% of all time-frequency samples used in this analysis.

3.1.2. Thermal Gain Correction

Section 3.2.4 of CHIME Collaboration et al. (2023) describes a stage of the data processing that uses measurements of the outside temperature to correct common-mode thermal variations in the amplitude of the gains of the analog receiver chains. We recently discovered two errors in the resulting instrument stability as reported in CHIME Collaboration et al. (2023). First, the pre(post)-correction stability of 0.8% (0.5%) was reported as the standard deviation in fractional power, but it is actually the standard deviation in fractional voltage. Second, the software that was deployed in the pipeline mistakenly applied the inverse of the thermal correction, which amplified the thermal variations instead of mitigating them. We have fixed the software error and applied the correct factor to the data used for this analysis. As a result, the estimated stability for the analysis presented in CHIME Collaboration et al. (2023) is 2.7% (standard deviation in fractional power), while the estimated stability for this analysis is 1.0%.

3.1.3. Impact on Previously Reported Results

The changes to the CHIME processing pipeline described in this section do not result in an improved signal-to-noise ratio on the detection of cosmological 21 cm emission in cross-

correlation with the eBOSS galaxy catalogs that was reported in CHIME Collaboration et al. (2023). In fact, the noise level in the more recent revision of the processed data is found to be elevated compared to what was measured in the previous one, but this has no qualitative impact on the results reported in CHIME Collaboration et al. (2023). Investigating the impact of these changes and continued iteration on the processing pipeline is currently ongoing.

3.2. Ringmaps

The way CHIME observes the sky—a primary beam that is a few degrees wide in the east–west direction but spans the south and north horizons, illuminating a regularly spaced grid of feeds as the sky drifts overhead—is amenable to a particular map-making method, which produces “ringmaps.” At every time sample, corresponding to a specific R.A. on meridian, the north–south (NS) baselines are phased and combined so as to form a grid of synthesized beams within the instantaneous field of view. As the sky drifts overhead, each formed beam records its intensity along a ring of constant decl. The image generated in this way is called a ringmap. An example is displayed in the top panel of Figure 2. The resulting maps have a resolution of $0^\circ.1$ in R.A. and $0^\circ.1$ – $0^\circ.15$ in decl. within the region shown in the figure, which is finer than the resolution limit of the CHIME synthesized beam at these frequencies. A detailed description of this process is given in CHIME Collaboration et al. (2023).

For the cross-correlation analysis, we generate ringmaps from the sidereal stacks at each polarization, and find the angular pixel that contains the position of each quasar in the eBOSS catalog. The spectrum from this pixel is what is correlated with the corresponding $\text{Ly}\alpha$ forest spectrum. About 8% of the quasar lines of sight fall within the same CHIME pixel, and of these, around 12% are within 10 frequency channels of each other in redshift (roughly the scale up to which correlation is observed, see Section 6). There is thus very little redundancy in the lines of sight that are averaged together.

3.3. Lower-band Processing

3.3.1. Beam Calibration

When ringmaps are generated, a beam model is used to deconvolve the effect of the main lobe of the primary beam. The model that was used in CHIME Collaboration et al. (2023) was derived from data by fitting to the expected flux of a large number of radio point sources. This model has since been extended to the lower band considered in this analysis, using the same methods that were described in detail in CHIME Collaboration et al. (2023). Slices along each axis of the beam model are presented in Figure 3, for both polarizations.

3.3.2. Frequency Flagging

Although RFI is flagged algorithmically at multiple stages in the real-time and daily processing, low-level events that are fainter than the foregrounds are very hard to detect. This contamination becomes apparent after a high-pass delay filter has been applied to remove most of the foreground power, and it is strong enough to obscure the cosmological signal. We identify affected frequencies by iteratively filtering and flagging outlier frequencies until none remain. The flagging was performed by visually inspecting maps generated from the filtered data at every frequency, where the spatial information

helps identify bad cases. In the end, 47% of the data in the 400–500 MHz band were flagged in this way, in addition to persistent RFI bands that are statically masked, such that a total of 65% of the band was masked in this analysis.

We attributed this contamination to RFI, but instrumental miscalibration such as the narrowband gain errors described above could have a similar effect. As the calibration and flagging parts of the daily pipeline continue to improve, we may be able to recover some of these flagged frequencies.

3.3.3. Delay Filter

In this analysis, we suppress the foregrounds by high-pass filtering the 400–500 MHz band in delay. The filter is implemented using the DAYENU method to account for the masked regions of the frequency band (Ewall-Wice et al. 2021). We set the lower edge of the delay passband at 200 ns, which was chosen to encompass the bulk of foreground power at all declinations in the delay power spectrum, as illustrated in Figure 4. This power spectrum was evaluated using the Gibbs sampling method described in Appendix A of CHIME Collaboration et al. (2023) and estimating the variance across right ascensions 110° – 270° , a region corresponding roughly to the eBOSS NGC field, where the majority of the $\text{Ly}\alpha$ quasars are found. Shifting this threshold by 50 or 100 ns in either direction did not appear to improve the signal-to-noise ratio, but no further optimization was attempted. Filtering delays below 200 ns corresponds to removing large scales with $k_{\parallel} \lesssim 0.13 \text{ Mpc}^{-1}$, at the average redshift $\bar{z} = 2.3$ ($\sim 430 \text{ MHz}$). In frequency space, this corresponds to scales larger than 5 MHz separation. Figure 2 shows maps at a single frequency and polarization before and after the delay filter. The total power remaining in the maps after the filter is about 0.6% of the power before, and the fraction of power above 200 ns before filtering is 2.5%. The difference between these two figures is probably due to leakage from low to high delays in the power spectrum estimation step. Pixels near bright foreground features tend to remain corrupted after filtering, so we also derive a mask that excises any pixel above a 6σ threshold determined by the expected thermal noise. This threshold is what was used in CHIME Collaboration et al. (2023), and removes sufficient contamination to enable a detection of the 21 cm signal.

4. Cross-correlation Method

The cross-correlation function of the 21 cm temperature T_{21} and the $\text{Ly}\alpha$ forest relative flux transmission δ_α is defined at a given redshift as

$$\xi(\mathbf{r}, \mathbf{r}') = \langle T_{21}(\mathbf{r}) \delta_\alpha(\mathbf{r}') \rangle, \quad (2)$$

where $\mathbf{r}$ and $\mathbf{r}'$ are three-dimensional comoving positions and the angle brackets indicate an ensemble average. Assuming statistical isotropy, the correlation function only depends on the magnitude of the separation in comoving space: $\xi(\mathbf{r}, \mathbf{r}') = \xi(|\mathbf{r} - \mathbf{r}'|)$. However, measurements of these fields are located on our past lightcone—in redshift space and on the sphere—denoted by $(z, \hat{\mathbf{n}})$, so this is where we will evaluate the correlation function.

In aggregate, the line-of-sight measurements from eBOSS and CHIME contain three-dimensional information on the two fields, i.e., one could evaluate the correlations along the redshift (z) and angular ($\hat{\mathbf{n}}$) coordinates. However, CHIME’s resolving power is much greater in frequency than it is on the sky, so we expect the signal-to-noise ratio to be stronger along this axis.




Figure 2. (First) Ringmap generated from inter-cylinder baselines ($\gtrsim 22$ m east–west distance) at 476.56 MHz for the XX polarization. (Second) Ringmap generated from visibilities filtered using DAYENU with a 200 ns delay cut. (Third) The same as the second panel but with an outlier mask derived from a 6σ threshold based on the expected noise. (Fourth) A noise realization at the level expected in the CHIME data, delay-filtered in the same way as in the previous panels. In all panels, vertical dashed lines indicate the region that was used to estimate the delay power spectrum in Figure 4.

Given the significant added complexity of a 3D analysis, in this work we only consider correlations along the line of sight.

Here, we will describe the procedure used to compute an estimate of the radial cross-correlation function in redshift space, corresponding to

$$\xi(\Delta z) = \langle \int dz T_{21}^i(z) \delta_\alpha^i(z + \Delta z) \rangle_i, \quad (3)$$

where the angle brackets denote the average over individual line-of-sight observations labeled by i . We describe how redshifts are computed for each observable in Section 4.1, and how the cross-correlation is computed in Section 4.2.

4.1. Frequency Remapping

The eBOSS $\text{Ly}\alpha$ forest spectra measurements are arranged in wavelength bins uniformly spaced in $\ln \lambda$ and spanning

$\sim 360\text{--}700$ nm, whereas the CHIME data are uniformly spaced in frequency in the 400–800 MHz band. In order to correlate the two, it will be necessary to map these onto a common grid in redshift. Given that the resolution of the CHIME data is coarser by a factor of $\sim 1.2\text{--}1.5$ in the redshift range where they overlap, and that it will be convenient to work with the CHIME data directly, we choose to map the $\text{Ly}\alpha$ data onto the CHIME band. The mapping is defined by matching the frequency of $\text{Ly}\alpha$ absorption to the observed frequency of 21 cm emission for the corresponding redshift, a relationship given by the ratio of their rest wavelengths:

$$\nu' = \frac{\lambda_\alpha}{\lambda_{21}} \nu. \quad (4)$$

We use the reverse Lanczos interpolation method described in Section 3 of CHIME Collaboration et al. (2023) to regrid the $\text{Ly}\alpha$ forest spectra onto the CHIME band. The regions of the spectra that are empty (i.e., those with redshifts larger than that of the background quasar, or that are otherwise masked) are set to zero in the regridded data.

4.2. Correlation Function Estimation

We will write the discrete samples in frequency as $T_{21}[n] = T_{21}(n\Delta\nu + 400 \text{ MHz})$, with integers n labeling samples at the CHIME frequency resolution, $\Delta\nu = 390.625 \text{ kHz}$. Both the $\text{Ly}\alpha$ and CHIME spectra have noise estimates associated with every sample, so we use these to weight the samples and improve the signal-to-noise ratio. With inverse variance weights, $w = \sigma^{-2}$, we estimate the cross-correlation function as

$$\xi[n] = W^{-1} \sum_i^M \sum_j^N w_{21}^i[j] w_\alpha^i[j+n] T_{21}^i[j] \delta_\alpha^i[j+n], \quad (5)$$

where i labels lines of sight, n indexes separation in frequency, and W^{-1} is a normalization term given by

$$W = \sum_i^M \sum_j^N w_{21}^i[j] w_\alpha^i[j+n]. \quad (6)$$

Here, $M = 210005$ is the total number of lines of sight, and $N \leq 255$ is the number of samples along the frequency axis being averaged together at that separation.

5. Signal Model and Simulations

To allow us to interpret the results of this analysis and assess detection significance, we generate simulations of correlated 21 cm and $\text{Ly}\alpha$ forest measurements. These are run through the same processing pipeline and analysis as the data, in order to forward model the cross-correlation result. Although the models described below are parametric, we do not attempt to fit them to the data. Instead, the purpose of these models is to generate a reasonably realistic template to compare to our measurement and assess the S/N.

5.1. Large-scale Structure Simulation

A detailed description of the model used to generate maps of large-scale structure in the CHIME redshift range is given in Section 5 of CHIME Collaboration et al. (2023). Very briefly, a realization of Gaussian fluctuations on the lightcone is generated at the desired redshifts, given a nonlinear matter




Figure 3. Slices of the 3D beam model, from left to right, at constant hour angle, decl., and frequency. The top row is the power beam for the X polarization, and the bottom row is Y. The beam pattern has been projected onto sky coordinates HA and decl., which is why it broadens significantly near the NCP. It is worth noting that the decl. axis has been extended beyond 90° , to reflect the sensitivity of the beam to regions of the sky on either side of the NCP.




Figure 4. Delay power spectrum of a CHIME ringmap evaluated across right ascensions between 110° and 270° , at fixed decl. The color scale represents power in units of noise power. Vertical dashed lines show the lower bounds of the DAYENU filter passband. Significant contamination exists at all delays. Some of the horizontal features can be attributed to bright point sources in the maps, but the vertical lines are of unknown origin.

power spectrum evaluated using the halo model prediction from Mead et al. (2021).

5.2. CHIME Signal Model

The H I linear bias, Fingers of God effect, and 21 cm brightness temperature are modeled to generate a map of 21 cm temperature. See Section 5 of CHIME Collaboration et al. (2023) for the model definitions. In this work, we fix the parameters to the fiducial values defined there. Synthetic CHIME observations are derived from these maps and directly substituted for the real data in the cross-correlation analysis.

5.3. Ly α Signal Model

In order to model the Ly α forest measurements, we need a prescription for converting the matter density fluctuation field $\delta_m(z, \hat{n})$ generated by our LSS simulation into an optical depth

to Ly α photons. A commonly used and straightforward approach is to model the mildly nonlinear baryon field as log-normally distributed (Bi & Davidsen 1997) and the Ly α absorption by the fluctuating Gunn–Peterson approximation (FGPA) (Farr et al. 2020).

The log-normal transform is performed so as to preserve the variance of the original field σ_m^2 :

$$1 + \delta_{\text{LN}}(z, \hat{n}) = \exp \left[\delta_m(z, \hat{n}) - \frac{\sigma_m^2(z)}{2} \right]. \quad (7)$$

This step also ensures that the density is strictly non-negative.

The optical depth τ is proportional to the number density of H I atoms, which the FGPA assumes is related to the density by a power law:

$$\tau(z, \hat{n}) = \tau_0(z)(1 + \delta_{\text{LN}}(z, \hat{n}))^\alpha. \quad (8)$$

We adopt the parameter values

$$\tau_0(z) = 0.3 \left(\frac{1+z}{3.4} \right)^{4.5}, \quad \alpha = 1.6, \quad (9)$$

quoted in Cieplak & Slosar (2016) and Seljak (2012).

These steps are illustrated in the top three panels of Figure 5 for a realization of LSS along a given line of sight.

Rather than use the FGPA, it might be more self-consistent to derive the optical depth from the H I density that was modeled for the 21 cm temperature directly. This is expected to be part of a future effort to more carefully model the Ly α signal and perhaps derive constraints on the model parameters. We emphasize that, for the purpose of this work, the simplest approach was taken. However, we will show that a simple FGPA-based model is sufficient to describe our measurements at the current signal-to-noise level, so we leave more detailed modeling to future analyses.

5.4. Mock Ly α Forest Catalog

From the Ly α fractional flux transmission simulation, we extract individual spectra along the lines of sight at a number of




Figure 5. Fluctuations along a given line of sight in redshift, at subsequent steps of the procedure for generating simulated Ly α forest spectra. From top to bottom: the matter fluctuations from the LSS simulation; the result of the log-normal transform; the Ly α fractional flux transmission; the final spectrum masked and with noise added to emulate the real data.

sky locations, chosen to be those of the quasars in the eBOSS catalog. Each spectrum is also masked so that regions of missing data—redshifts larger than the backlight quasar and instrumental masks—match those in the corresponding eBOSS data. Gaussian random noise is optionally added to the simulated spectra at a level consistent with the noise variance recorded alongside the data. This results in a mock catalog with exactly the same quasar sample as the data, and thus the same sensitivity to cross-correlation with CHIME, but with a synthetic Ly α forest signal. The bottom panel of Figure 5 shows a mock spectrum generated in this way.

It must be noted that these simulations neglect a number of effects that would be necessary to include in order to produce realistic simulations of the eBOSS spectra (and which *are* included in more sophisticated simulations, e.g., Farr et al. (2020)). Among the most significant are

1. Redshift-space distortions,
2. Metal and Lyman- β line absorption,
3. High column density absorbers (which include damped Ly α absorber systems), and
4. Biases due to the modeling of the quasar spectra continuum.

A careful accounting of these will be necessary for future work that will aim to constrain the physics of the Ly α forest.

6. Results

The cross-correlation functions of CHIME and the eBOSS Ly α forest are presented in Figure 6, where the two polarizations have been combined to approximate total intensity (they are shown separately in Figure 7). It should be noted that the cross-polar response of the telescope beam has not been calibrated, so this approximation may include leakage from polarized emission as well. An estimate of the contamination from noise and residual foregrounds was derived by cross-correlating the CHIME data against permutations of the Ly α forest spectra, i.e., exchanging the spectrum measurements between lines of sight so that they are uncorrelated with the CHIME observations but maintain exactly the same distribution on the sky (this is described in more detail below). The measurements show a clear excess over the background

around the zero-lag bin, with the negative sign characteristic of this correlation. Also shown in this figure is a signal template derived from the result of cross-correlating the simulated data sets described in the previous section. Only the amplitude of the template was fit to the data, as explained below, and it appears to be broadly consistent with the measurement.

6.1. Ly α Forest Permutations

In order to characterize the background fluctuations in the cross-correlation estimates, we would like to generate a catalog of Ly α forest measurements with the same properties as the eBOSS one but uncorrelated with the CHIME data. Noise, foreground residuals, and any other systematic effects that are not correlated with the signal will remain present in the resulting correlation functions, and we can measure their power. We do this by drawing random permutations of the Ly α forest spectra, while keeping the sample of line-of-sight directions and quasar redshifts fixed. One thousand such permutations were generated, and correlation functions were evaluated for all of them. The standard deviation over permutations at every spectral separation is shown as the shaded region in Figure 6. It is worth noting that correlations between frequency separations were not considered (i.e., the covariance was assumed to be diagonal).

It should also be noted that this procedure keeps the selection of lines of sight in the CHIME data fixed, such that if by chance it were to include data that deviate significantly from the rest, this issue would be present in every permutation and potentially bias the estimated variance. To check for this, we generate another set of permutations where the catalog positions have also been rotated by $\sim 3^\circ$ (approximately one CHIME primary beamwidth at these wavelengths) in a random direction. We find that the measured variance in the background fluctuations from these rotated permutations is indistinguishable from the unrotated set within the sample noise. Thus, we do not expect that the variance computed from the unrotated permutations is subject to significant bias due to the fixed set of lines of sight.

6.2. Detection Significance

To assess the significance of the detection, we derive a signal template from the simulations. Synthetic CHIME and Ly α forest data were generated without adding noise and run through the same filtering and cross-correlation analysis. The resulting correlation functions, the dashed lines in Figure 7, have a shape that is compact in frequency separation, a result of the delay filter removing larger scales. These were fit to the cross-correlation measurements by varying a single amplitude parameter and minimizing the χ^2 . Finding the $\Delta\chi^2 = 1$ bounds in parameter space provides an estimate for the standard deviation of the amplitude constraint. The fitted template and residuals are presented in Figure 6 for the combined polarizations. We quote the significance of the detection as the amplitude of the template in units of its standard deviation from the fit in Table 1.

The fitted amplitude of the template to match the measured signal is a factor of ~ 4 , i.e., the measurement is about 4 times brighter than the simulation. Given that the CHIME instrument model used in the simulations is a simple approximation—in particular, the model for the beam response, which is essential




Figure 6. (Top) Cross-correlation function of CHIME and eBOSS Ly α forest data, for the combined X and Y polarizations. An estimate of the standard deviation of the background, as described in Section 6.1, is plotted as a shaded region. The black dashed line is a template derived from simulations with an amplitude fit to the data. (Bottom) Residuals normalized by the estimated background level.




Figure 7. Measured cross-correlation functions for each polarization, along with the corresponding template fit to that polarization.

in order to accurately normalize the signal amplitude—we are not confident in interpreting this as a discrepancy between the physics of the model and the observations. However, future work that refines both the instrumental calibration and modeling should allow a useful amplitude to be constrained.

7. Validation

In this section, we describe checks that were performed to validate the detection and background estimate.

7.1. Impact of Selection Function on Catalog Permutations

To estimate the level of background fluctuations in the correlation function measurements, random permutations of the Ly α forest catalog were drawn and correlated against the CHIME data, as described in the previous section. By construction, this method preserves the distribution of background quasars on the sky and in redshift separately, but if there exist correlations between these two axes in the selection function, they will be erased by our procedure. We check for

Table 1
Detection Significance of Correlation Function Template Fits

Measurement	χ^2	χ^2_ν	Significance
X pol	217.1	0.85	6.4 σ
Y pol	295.8	1.16	6.5 σ
Combined	238.5	0.94	9.1 σ

Note. The number of degrees of freedom used to compute χ_ν is 254.

the presence of such a correlation by binning the background quasars by their angular position onto a HEALPix (Górski et al. 2005) grid of “nside” 32 and computing the average redshift in each bin. The standard deviation across non-empty bins is $\Delta z < 0.08$, or about 3% of the average redshift. For comparison, we shuffled the positions of the quasars and repeated the binning calculation, resulting in a typical standard deviation of just under 3%. There does not appear to be more variation in the redshifts with location on the sky than would be expected for an uncorrelated distribution, and we conclude that selection effects are unlikely to be important.

7.2. Even/Odd Days

For this test, the full set of days of CHIME data that went into the sidereal stack were divided into two sets according to an even/odd split in chronological order, and each was stacked in the same way as the full set (see Section 2). The 21 cm and foreground signals are the same on every sidereal day, so they are expected to be common to the even and odd stacks, but thermal noise and RFI should be uncorrelated between the two. By taking the difference (and dividing by a factor of two), the common foreground residuals should cancel and what is left can be compared to our expectation for thermal noise. Any remaining excess over the noise can be attributed to day-to-day variations, e.g., RFI or changes to the instrumental response that were not entirely captured in the daily calibration.




Figure 8. Correlation function evaluated on the difference of even and odd days, with the standard deviation derived from noise realizations as a shaded region.

The correlation functions evaluated on the even/odd difference are shown in Figure 8, along with the expected noise level. Not surprisingly, there is no visible excess around the zero-lag bin, which was confirmed by fitting the signal template to this data, resulting in amplitude parameters that are consistent with zero within 1σ . We note that the background fluctuations appear to have a magnitude in excess of the noise. To make this comparison more clearly, the correlation function was evaluated for 1000 permutations of the $\text{Ly}\alpha$ forest spectra for the even–odd difference, as well as for a CHIME noise realization. In all cases, the same set of permutations was used. Figure 9 shows the variance across permutations for these two cases along with the regular data. The ratio between the variance in the regular analysis and the even–odd difference indicates the contribution of residual foregrounds to the background fluctuations in the correlation function. These account for about twice as much variance as the even–odd difference, which in turn is about 3–5 times as large as the noise. However, one should note that the variance is skewed by significant non-Gaussian tails in the distributions of the pixel values.

Figure 10 shows histograms of the pixels in the filtered and masked maps across all frequencies that were used in the upper and lower bands, within a region of the sky restricted roughly to the eBOSS NGC field. In the even–odd difference of the lower band, the distribution is closer to Gaussian for values $\lesssim 2\sigma$, but large tails remain. The histograms of the upper band are much closer to Gaussian for $\lesssim 3\sigma$, especially in the even–odd difference.

The analysis of the difference of even and odd days in the lower band suggests the following:

1. There is a large excess of variance in the maps above the thermal noise (by a factor ~ 6 – 10).
2. This excess is not fully accounted for by features that are repeated day to day, like celestial foregrounds (the variance is only $\sim$ halved in the even–odd difference).
3. The excess above the noise is non-Gaussian, with large tails in the distribution of map pixels.
4. The distribution is significantly less Gaussian in the lower band than in the upper band, especially in the even–odd difference.




Figure 9. Variance computed across correlation functions estimated on 1000 permutations of the $\text{Ly}\alpha$ forest spectra. In blue is the full CHIME data, orange is the difference of even and odd days, and green is a noise realization. It should be noted that the minimum does not occur at zero separation, because the maximal overlap of the highest weighted regions in the CHIME and eBOSS spectra is not achieved there. Also, the full range of frequency separations is shown here, whereas previous plots were cropped to better show the signal.

A possible explanation for this is that the lower band is more contaminated than the upper band by unflagged RFI—which does not repeat day to day and contributes in a non-Gaussian-distributed way. Another contributing effect could also be errors in the daily calibration, which would have a similar signature. Both of these factors have been observed in other internal CHIME analyses of the lower band. Identifying the source of this contamination and removing it will likely require further investigation of the individual days, before they are averaged into a sidereal stack. While these residuals are concerning and finding their origin in order to remove them will enable more sensitive analyses, their presence has not impeded the detection in cross-correlation with the $\text{Ly}\alpha$ forest reported in this work.

8. Conclusion

This work reports on the first detection of 21 cm emission at redshifts $z > 1.5$, obtained by cross-correlating CHIME data from 88 days in the 400–500 MHz band with the DR16 eBOSS measurements of the $\text{Ly}\alpha$ forest. To do so, we extended the analysis methods described in CHIME Collaboration et al. (2023) to the lower quarter of the CHIME frequency band, and developed a cross-correlation method to combine this data with the $\text{Ly}\alpha$ forest as a tracer of large-scale structure.

This demonstrates that CHIME is able to detect emission from cosmological H I in the entire redshift range available to it, but it also highlights the difficulty of separating this signal from the extremely bright foregrounds. Even in cross-correlation and after aggressive delay filtering, excess variance from foregrounds and other sources is found to be about 6–10 times above thermal noise. This level of contamination nevertheless allows for a $\sim 9\sigma$ detection. The shape of the measured correlation function is dominated by the effect of the foreground filter, and although we have not attempted to interpret it, we have found that a simple physical model agrees with the measurement well enough to use this model to




Figure 10. Histogram of pixel values in filtered and masked ringmaps across all frequencies used in (left) the lower band, 400–500 MHz, and (right) the upper band, 585–800 MHz. The pixels are normalized to the expected noise standard deviation, and a Gaussian curve with a matching amplitude is plotted in orange for comparison. Rows show the two polarizations, and columns show the difference of even and odd days and the average of all days. The effect of the outlier mask can be seen as the sharp cutoff at 6σ .

quantify the detection significance. The amplitude of the cross-correlation should in principle tell us something about the physics of the HI that sources the Ly α forest and its relationship to the high-HI-density systems primarily probed with 21 cm intensity mapping, but interpreting its measurement will require more sophisticated modeling of the Ly α and 21 cm signals. This will be the subject of future work.

Improvements to the CHIME calibration and analysis pipelines, as well as the inclusion of several more years of data that have yet to be processed, will lead to future improvements in the constraining power of the analysis. Significant residuals in day-to-day jackknives also suggest that the lower band considered in this work is more contaminated than the upper band that was reported on previously, and may especially benefit from improved RFI flagging. In the short term, this might reduce the number of frequencies that need to be flagged and allow for a less restrictive delay cut, enabling access to information from scales larger than the $k_{\parallel} \gtrsim 0.13 \text{ Mpc}^{-1}$ that remain in this analysis. These improvements will position CHIME’s data set as a powerful source of cross- and autocorrelation measurements.

Acknowledgments

We thank Meiling Deng for her substantial contributions to the instrument design and analysis of CHIME.

We thank the Dominion Radio Astrophysical Observatory, operated by the National Research Council Canada, for gracious hospitality and expertise. The DRAO is situated on the traditional, ancestral, and unceded territory of the Syilx Okanagan people. We are fortunate to live and work on these lands.

CHIME is funded by grants from the Canada Foundation for Innovation (CFI) 2012 Leading Edge Fund (Project 31170), the CFI 2015 Innovation Fund (Project 33213), and by contributions from the provinces of British Columbia, Québec, and Ontario. Long-term data storage and computational support for analysis is provided by WestGrid, SciNet, and the Digital Research Alliance of Canada, and we thank their staff for flexibility and technical expertise that has been essential to this work.

Additional support was provided by the University of British Columbia, McGill University, and the University of Toronto.

CHIME also benefits from NSERC Discovery Grants to several researchers, funding from the Canadian Institute for Advanced Research (CIFAR), from Canada Research Chairs, from the FRQNT Centre de Recherche en Astrophysique du Québec (CRAQ) and from the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto, which is funded through an endowment established by the David Dunlap family. This material is partly based on work supported by the NSF through grants (2008031) (2006911) and (2006548) and by the Perimeter Institute for Theoretical Physics, which in turn is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation.

We thank the Sloan Digital Sky Survey and eBOSS collaborations for publicly releasing the Ly α forest catalogs used in this work. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSSIV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is www.sdss.org.

Software: bitshuffle (Masui et al. 2015), CAMB (Lewis et al. 2000), caput (Shaw et al. 2020c), ch_pipeline (Shaw et al. 2020b), cora (Shaw et al. 2020d), Cython (Behnel et al. 2011), draco (Shaw et al. 2020e), driftscan (Shaw et al. 2020a), hankl (Karamanis & Beutler 2021), h5py (Collette et al. 2021), HDF5 (The HDF Group 2020), HEALPix (Górski et al. 2005), healpy (Zonca et al. 2019), Matplotlib (Hunter 2007), mpi4py (Dalcin & Fang 2021), NumPy (Harris et al. 2020), OpenMPI (Gabriel et al. 2004), pandas (pandas development team 2020; McKinney 2010), peewee (Leifer 2021), SciPy (Virtanen et al. 2020), Skyfield (Rhodes 2019).

ORCID iDs

Mandana Amiri <https://orcid.org/0000-0001-6523-9029>
 Kevin Bandura <https://orcid.org/0000-0003-3772-2798>
 Arnab Chakraborty <https://orcid.org/0000-0002-7758-9859>
 Matt Dobbs <https://orcid.org/0000-0001-7166-6422>
 Mateus Fandino <https://orcid.org/0000-0002-6899-1176>
 Simon Foreman <https://orcid.org/0000-0002-0190-2271>

Hyoyin Gan  <https://orcid.org/0000-0002-4903-838X>
 Mark Halpern  <https://orcid.org/0000-0002-1760-0868>
 Alex S. Hill  <https://orcid.org/0000-0001-7301-5666>
 Gary Hinshaw  <https://orcid.org/0000-0002-4241-8320>
 Carolin Höfer  <https://orcid.org/0000-0003-4887-8114>
 T. L. Landecker  <https://orcid.org/0000-0003-1455-2546>
 Zack Li  <https://orcid.org/0000-0002-0309-9750>
 Joshua MacEachern  <https://orcid.org/0000-0001-8064-6116>
 Kiyoshi Masui  <https://orcid.org/0000-0002-4279-6946>
 Juan Mena-Parra  <https://orcid.org/0000-0002-0772-9326>
 Nikola Milutinovic  <https://orcid.org/0000-0001-8292-0051>
 Arash Mirhosseini  <https://orcid.org/0000-0002-2626-5985>
 Laura Newburgh  <https://orcid.org/0000-0002-7333-5552>
 Anna Ordog  <https://orcid.org/0000-0002-2465-8937>
 Sourabh Paul  <https://orcid.org/0000-0002-8671-2177>
 Ue-Li Pen  <https://orcid.org/0000-0003-2155-9578>
 Tristan Pinsonneault-Marotte  <https://orcid.org/0000-0002-9516-3245>
 Alex Reda  <https://orcid.org/0000-0001-6967-7253>
 J. Richard Shaw  <https://orcid.org/0000-0002-4543-4588>
 Seth R. Siegel  <https://orcid.org/0000-0003-2631-6217>
 Keith Vanderlinde  <https://orcid.org/0000-0003-4535-9378>
 Haochen Wang  <https://orcid.org/0000-0002-1491-3738>
 D. V. Wiebe  <https://orcid.org/0000-0002-6669-3159>
 Dallas Wulf  <https://orcid.org/0000-0001-7314-9496>

References

- Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *ApJ*, **925**, 221
 Anderson, C. J., Luciw, N. J., Li, Y. C., et al. 2018, *MNRAS*, **476**, 3382
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, **13**, 31
 Bi, H., & Davidsen, A. F. 1997, *ApJ*, **479**, 523
 Bowman, J. D., Rogers, A. E. E., & Hewitt, J. N. 2008, in AIP Conf. Ser. 1035, The Evolution of Galaxies Through the Neutral Hydrogen Window, ed. R. Minchin & E. Momjian (Melville, NY: AIP), 87
 Carucci, I. P., Villaescusa-Navarro, F., & Viel, M. 2017, *JCAP*, **2017**, 001
 Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, *Natur*, **466**, 463
 CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022, *ApJS*, **261**, 29
 CHIME Collaboration, Amiri, M., Bandura, K., et al. 2023, *ApJ*, **947**, 16
 Cieplak, A. M., & Slosar, A. 2016, *JCAP*, **2016**, 016
 Collette, A., Kluyver, T., Caswell, T. A., et al. 2021, h5py/h5py v3.5.0, Zenodo, doi:10.5281/zenodo.5585380
 Cunningham, S., Li, Y., Santos, M. G., et al. 2023, *MNRAS*, **518**, 6262
 Dalcin, L., & Fang, Y.-L. L. 2021, *CSE*, **23**, 47
 DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
 du Mas des Bourboux, H., Rich, J., Font-Ribera, A., et al. 2020, *ApJ*, **901**, 153
 Ewall-Wice, A., Kern, N., Dillon, J. S., et al. 2021, *MNRAS*, **500**, 5195
 Farr, J., Font-Ribera, A., Bourboux, H. D. M. D., et al. 2020, *JCAP*, **2020**, 068
 Gabriel, E., Fagg, G. E., Bosilca, G., et al. 2004, Recent Advances in Parallel Virtual Machine and Message Passing Interface (Berlin: Springer), 97
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Jonas, J. L. 2009, *Proc. SPIE*, **97**, 1522
 Karamanis, M., & Beutler, F. 2021, arXiv:2106.06331
 Leifer, C. 2021, peewee/peewee, latest, Github, <https://github.com/coleifer/peewee>
 Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473
 Li, L.-C., Staveley-Smith, L., & Rhee, J. 2021, *RAA*, **21**, 030
 Liang, Y., Kashikawa, N., Cai, Z., et al. 2021, *ApJ*, **907**, 3
 Masui, K., Amiri, M., Connor, L., et al. 2015, *A&C*, **12**, 181
 Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, *ApJL*, **763**, L20
 McKinney, W. 2010, in Proc. 9th Python in Sci. Conf., ed. S. van der Walt & J. Millman, 56
 Mead, A. J., Brieden, S., Tröster, T., & Heymans, C. 2021, *MNRAS*, **502**, 1401
 Momose, R., Shimasaku, K., Kashikawa, N., et al. 2021, *ApJ*, **909**, 117
 Mukae, S., Ouchi, M., Kakiichi, K., et al. 2017, *ApJ*, **835**, 281
 Pandas development team 2020, pandas-dev/pandas: Pandas, v2.2.0, Zenodo, doi:10.5281/zenodo.3509134
 Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *AJ*, **139**, 1468
 Patra, N., Subrahmanyam, R., Raghunathan, A., & Udaya Shankar, N. 2013, *ExA*, **36**, 319
 Paul, S., Santos, M. G., Chen, Z., & Wolz, L. 2023, arXiv:2301.11943
 Pen, U.-L., Staveley-Smith, L., Peterson, J. B., & Chang, T.-C. 2009, *MNRAS*, **394**, L6
 Peterson, J. B., Aleksan, R., Ansari, R., et al. 2009, *astro2010: The Astronomy and Astrophysics Decadal Survey*, **2010**, 234
 Rhodes, B., 2019 Skyfield: High precision research-grade positions for planets and Earth satellites generator, Astrophysics Source Code Library, ascl:1907.024
 Seljak, U. 2012, *JCAP*, **2012**, 004
 Shaw, J. R., Foreman, S., Nitsche, R., et al. 2020a, radiocosmology/driftscan v20.5.0, Zenodo, doi:10.5281/zenodo.5949193
 Shaw, J. R., Masui, K., Hincks, A. D., et al. 2020b, chime-experiment/ch_pipeline v20.10.0, Zenodo, doi:10.5281/zenodo.5846379
 Shaw, J. R., Masui, K., Nitsche, R., et al. 2020c, radiocosmology/caput v20.10.0, Zenodo, doi:10.5281/zenodo.5846375
 Shaw, J. R., Nitsche, R., Foreman, S., & Kefala, A. 2020d, radiocosmology/cora v20.5.0, Zenodo, doi:10.5281/zenodo.5949195
 Shaw, R., Nitsche, R., Siegel, S. R., et al. 2020e, radiocosmology/draco v20.10.0, Zenodo, doi:10.5281/zenodo.5828373
 The HDF Group 2020, Hierarchical Data Format 5 v1.12.0, <https://github.com/HDFGroup/hdf5>
 Tramonte, D., & Ma, Y.-Z. 2020, *MNRAS*, **498**, 5916
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
 Wolz, L., Pourtsidou, A., Masui, K. W., et al. 2022, *MNRAS*, **510**, 3495
 Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298