



CHILES. VIII. Probing the Evolution of Average H I Content in Star-forming Galaxies over the Past 5 Gyr

Downloaded from: <https://research.chalmers.se>, 2025-06-16 18:18 UTC

Citation for the original published paper (version of record):

Luber, N., Pisano, D., van Gorkom, J. et al (2025). CHILES. VIII. Probing the Evolution of Average H I Content in Star-forming Galaxies over the Past 5 Gyr. *Astrophysical Journal*, 985(2). <http://dx.doi.org/10.3847/1538-4357/adc715>

N.B. When citing this work, cite the original published paper.



CHILES. VIII. Probing the Evolution of Average HI Content in Star-forming Galaxies over the Past 5 Gyr

Nicholas Luber¹, D. J. Pisano², J. H. van Gorkom¹, Julia Blue Bird³, Richard Dodson⁴, Hansung B. Gim⁵, Kelley M. Hess^{6,7}, Lucas R. Hunt³, Danielle Lucero⁸, Martin Meyer⁴, Emmanuel Momjian³, and Min S. Yun⁹

¹Department of Astronomy, Columbia University, Mail Code 5247, 538 West 120th Street, New York, NY 10027, USA; nicholas.m.luber@gmail.com

²Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

³National Radio Astronomy Observatory, P.O. Box O, Socorro, NM 87801, USA

⁴International Centre for Radio Astronomy Research (ICRAR), University of Western Australia, 35 Stirling Highway, Crawley, WA 6009, Australia

⁵Department of Physics, Montana State University, P.O. Box 173840, Bozeman, MT 59717, USA

⁶Department of Space, Earth and Environment, Chalmers University of Technology, Onsala Space Observatory, 43992 Onsala, Sweden

⁷ASTRON, The Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands

⁸Department of Physics, Virginia Tech, 850 West Campus Drive, Blacksburg, VA 24061, USA

⁹Department of Astronomy, University of Massachusetts, Amherst, MA 01003, USA

Received 2024 September 4; revised 2025 March 25; accepted 2025 March 27; published 2025 May 27

Abstract

Utilizing the COSMOS HI Large Extragalactic Survey (CHILES) data set, we investigate the evolution of the average atomic neutral hydrogen (HI) properties of galaxies over the continuous redshift range $0.09 < z < 0.47$. First, we introduce a simple multistep, multiscale imaging and continuum subtraction process that we apply to each observing session. These sessions are then averaged onto a common uv -grid and run through a Fourier filtering artifact mitigation technique. We then demonstrate how this process results in science quality data products by comparing to the expected noise and image-cube kurtosis. This work offers the first-look description and scientific analysis after the processing of the entire CHILES database. These data are used to measure the average HI mass in four redshift bins, out to a redshift 0.47, by separately stacking blue cloud (near-UV, $\text{NUV-}r = -1$ to 3) and red sequence ($\text{NUV-}r = 3-6$) galaxies. We find little-to-no change in gas fraction for the total ensemble of blue galaxies and make no detection for red galaxies. Additionally, we split up our sample of blue galaxies into an intermediate stellar mass bin ($M_* = 10^{9-10} M_\odot$) and a high stellar mass bin ($M_* = 10^{10-12.5} M_\odot$). We find that in the high-mass bin galaxies are becoming increasingly HI poor with decreasing redshift, while the intermediate-mass galaxies maintain a constant HI gas mass. We place these results in the context of the star-forming main sequence of galaxies and hypothesize about the different mechanisms responsible for their different evolutionary tracks.

Unified Astronomy Thesaurus concepts: Interstellar atomic gas (833); Late-type galaxies (907); Radio interferometry (1346)

1. Introduction

The observation and characterization of the atomic neutral hydrogen (HI) content of galaxies as a function of redshift is crucial for the development of a fully predictive theoretical model of galaxy evolution. Neutral hydrogen has been studied extensively in the local Universe in a range of galaxy types and environments. In clusters of galaxies, the densest environments, spiral galaxies, have been shown to be HI deficient (M. P. Haynes & R. Giovanelli 1986; J. M. Solanes et al. 2001; L. Cortese et al. 2011) and have compressed HI morphologies, or extraplanar HI gas, as a result of ram pressure stripping from the intracluster medium (J. E. Gunn & J. R. I. Gott 1972; J. D. P. Kenney et al. 2004; A. Chung et al. 2009; B. M. Poggianti et al. 2017). In galaxy groups (those collections of galaxies that are less numerous and sometimes less dense than clusters), the neutral gas content of galaxies can be affected by interactions between group members and processing by the intergroup medium (V. A. Kilborn et al. 2009; K. M. Hess & E. M. Wilcots 2013). In low-density regions, or cosmic voids, galaxies are shown to have HI

morphologies consistent with local optical–HI size scaling relationships for spiral and dwarf galaxies (K. Kreckel et al. 2012). Additionally, the HI phase has been characterized for massive elliptical galaxies (W. K. Huchtmeier 1994; R. Morganti et al. 2006), as well as a range of dwarf galaxy morphologies and stages of interaction (R. A. Swaters et al. 2002, 2009; S. Pearson et al. 2016).

Despite the wealth of HI observations for galaxies in different environments and different morphological types in the local Universe, it has been extremely challenging to directly measure HI emission in individual galaxies beyond $z > 0.1$. Some work has been done beyond this redshift, for example, the HI detection of a cluster galaxy at $z = 0.18$ in M. A. Zwaan et al. (2001), and the HIGHz Arecibo survey, which measured massive galaxies in the redshift range $0.17 < z < 0.25$ (B. Catinella et al. 2008; B. Catinella & L. Cortese 2015). The advent of broadband correlator backends made it possible to cover instantaneously a large frequency range. The new opportunity to probe a large volume at once made it worthwhile to do the deep integrations required to detect galaxies at higher redshift. The first project to do a blind survey of a large volume, probing two clusters and the large-scale structure (LSS) in which they are embedded in, was Blind Ultra Deep HI Extragalactic Survey (M. Verheijen et al. 2007; A. R. Gogate et al. 2020), where they sought to detect HI in individual



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.

galaxies around previously known clusters. Using the upgraded Giant Metrewave Radio Telescope, there have been detections of stacked H I emission at $z \approx 0.35$ (A. Bera et al. 2019), and $z = 0.7\text{--}1.45$ (A. Chowdhury et al. 2020, 2021), and characterizations of the $M_{\text{H I}}\text{--}M_*$ scaling relationship at $z \approx 0.35$ (A. Bera et al. 2023a). These measurements have indicated very little evolution in H I gas content at $z \approx 0.35$, but have found a significantly increased amount of H I at $z = 0.7\text{--}1.45$. In contrast, the MIGHTEE-H I project (N. Maddox et al. 2021), making use of the MeerKAT radio telescope found evolution of the $M_{\text{H I}}\text{--}M_*$ scaling relationship at $z \approx 0.35$ (F. Sinigaglia et al. 2022) and in follow-up studies that make use of a combined MIGHTEE and CHILES data cube (A. Bianchetti et al. 2025). These measurements represent the current frontier of understanding for the redshift evolution of the H I content of galaxies.

Further observations of neutral hydrogen across cosmic time are critical in understanding the evolution of galaxies in the Universe. Observationally, we know that the star formation rate density has decreased by a factor of 3 from a redshift $z = 0.5$ to the present day (A. M. Hopkins & J. F. Beacom 2006; P. Madau & M. Dickinson 2014), but we don't know what causes the decline, as we know very little about the evolution of the gas, which is the fuel for star formation. Theoretically, it has been shown that the dominant mechanism by which galaxies gain their gas, hot versus cold mode accretion, depends on both local environment and redshift with cold mode accretion most likely still observable at $z = 0$ in small galaxies in voids (D. Kereš et al. 2005, 2009).

The COSMOS H I Large Extragalactic Survey¹⁰ (CHILES) is a 1000 hr integration on a single pointing in the COSMOS field with the Karl G. Jansky Very Large Array¹¹ (VLA) in the B configuration. CHILES observes a frequency range of 960–1420 MHz, allowing for the direct detection of galaxies with an $M_{\text{H I}}$ of $3 \times 10^{10} M_\odot$, out to a redshift of 0.47. The first 178 hr of CHILES data have been successfully used to perform several different science verification studies. X. Fernández et al. (2016) presented the highest-redshift direct detection of H I in emission at $z = 0.376$ and found a galaxy with an unusually high star formation rate for its stellar mass. Note that I. Heywood et al. (2024) obtained a very low upper limit for the same galaxy, which will be discussed in J. Blue Bird et al. (2025, in preparation). K. M. Hess et al. (2019) demonstrated that even in a frequency range severely affected by radio frequency interference (RFI), there are stretches where galaxies can be detected with reliable extended morphologies. J. Blue Bird et al. (2020) investigated the spin alignment of nearby galaxies with respect to the orientation of the cosmic web filaments, as traced by the Discrete Persistent Source Extractor (T. Sousbie 2011) algorithm and customized by N. Luber et al. (2019) for the CHILES survey. In addition, in R. Dodson et al. (2022), we developed a successful imaging routine using the first 178 hr of CHILES data that utilizes a local sidereal time-dependent uv -based modeling and subtraction for sources that are beyond the half-power point of the VLA primary beam. This procedure properly accounts for the hour-angle dependence of the VLA primary beam and allows for the most accurate modeling of the flux density of sources beyond the

half-power point. In order for this routine to use this highly precise approach, large computational resources, such as those available through Amazon Web Services and the computing framework introduced in C. Wu et al. (2017), are required (R. Dodson et al. 2016).

The calibration and imaging of the CHILES data are complicated by the large volume (approximately 220 TB of raw data), and long duration (6 yr¹²) while requiring regular checks on the quality of the calibration and imaging. Additionally, our observations extend beyond the protected bands for radio astronomy, resulting in some of the bandwidth being significantly affected by RFI. As more data were calibrated and we integrated deeper, low-level image artifacts, such as residual sidelobes from a source 3@ from the field center, were found to be affecting the science field. As a result, it became apparent that a rapid turnaround to assess data and image quality was necessary in order to make informed decisions about how to proceed with data processing. In order to address these issues, here we introduce a quick and efficient imaging routine, that produces science-quality images. The technique we introduce in this work does not make all of the appropriate considerations that R. Dodson et al. (2022) do, but produces data cubes with a rapid turnaround time that can be used to assess data quality, as well as for scientific pursuits, and offers a suitable alternative for imaging large interferometric databases. As such, this work does not supersede that presented in R. Dodson et al. (2022), but offers an alternative methodology that can be used for scientific analysis while the full database subtraction using the technique in R. Dodson et al. (2022) can be completed.

This paper is structured as follows: In Section 2, we discuss the CHILES data, and the imaging techniques utilized to produce the cube used in our analysis, and go on to describe the ancillary data in the COSMOS field that we used in Section 3. In Section 4, we outline the procedure for H I and continuum stacking implemented for our analysis, and in Section 5, we summarize the results of the stacks we conduct. In Section 6 we discuss our findings within the broader context of previous literature results, and in Section 7, we summarize our findings. In this work, we assume a standard flat cold dark matter cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_\Lambda = 0.7$, and $\Omega_M = 0.3$.

2. CHILES Data and Imaging

2.1. The CHILES Database

The CHILES observations were taken with the VLA in its second most extended configuration (B). This configuration has a maximum baseline length of $\approx 11 \text{ km}$, corresponding to a synthesized beam of major axis $5''.5\text{--}8''.5$ at 1420 and 960 MHz, respectively, which in astrophysical terms corresponds to a physical resolution of $0.6\text{--}47.8 \text{ kpc}$ for redshifts $z = 0.005\text{--}0.48$. Assuming a velocity width of 150 km s^{-1} , our 3σ H I mass detection limit at $z = 0.48$ is $3.0 \times 10^{10} M_\odot$. This allows us to simultaneously have detailed H I maps at the lowest redshifts, and be able to identify individual galaxies at the highest redshifts. The observations use 60 of the 64 available correlator baseline board pairs to achieve the desired 30,720 channels spread across 15 spectral windows each of 32 MHz of bandwidth, in each observing session. To minimize loss of sensitivity at the edges of the band of each spectral window, we used frequency dithering.

¹⁰ VLA project 13B-266.

¹¹ The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

¹² Beginning in 2013 November and ending in 2019 April.

Table 1
Summary of Observations

Property	Value
Target Pointing	10 ^h 01 ^m 24 ^s +02°21′ 00″ ^a
Observation Dates	2013/10/25–2019/04/11 ^b
Total Observing Time	1062 hr ^{c,d}
Bandpass and Flux Density Scale	3C286
Complex Gain Calibrator	J0943-0819
Correlator Integration Time	8 s
Spectral Window Setup	15 × 32 MHz
Channel Width	15.625 kHz
Continuous Frequency Coverage	960–1420 MHz
Synthesized Beam	5″5–8″5 ^e

Notes.

^a In the J2000 coordinate standard.

^b The observations were taken over five different epochs, corresponding to five consecutive VLA-B configurations. This results in the observations being taken over different parts of the year and having a mix of both day- and nighttime observations.

^c The total integration time is split among the five observing epochs with 188, 220, 190, 237, and 192 hr, respectively.

^d Individual sessions have total observation times ranging from 1–8 hr, with a mean time of 4.75 hr, and approximately 45% of observations being 6 hr.

^e The range of synthesized beam major axis length. The size increases within this range with decreasing frequency.

Specifically, we shifted the edge of the lowest frequency in steps of 5 MHz, and used three different settings per epoch, and different settings across the five epochs. This resulted in almost uniform sensitivity across the 460 MHz band. A summary of the properties of the observations, including the pointing, observing dates, instrumental setup and calibrators, can be found in Table 1.

The frequency range covered by CHILES is impacted by a variety of RFI sources of varying magnitudes occurring at different frequencies. The lowest frequencies (approximately 960–1080 MHz) remain relatively RFI free, with the only exception being a narrow spike of RFI at 1030 MHz due to aeronautical ground-to-air radar. Similarly, the highest frequencies (approximately 1300–1420 MHz) are relatively RFI free with the exception of several narrow spikes due to GPS satellites and aeronautical radar. The frequency range in the middle (1080–1300 MHz) is the most severely affected range, with several different sources of RFI. This includes RFI from aeronautical radar, GLONASS, and GPS satellites, and communications from the nearby airport in Albuquerque.

The data are calibrated with a modified version of NRAO’s calibration pipeline for the VLA data using Common Astronomical Software Application (CASA) version 5.3.0 (D. J. Pisano et al. 2025, in preparation) at the West Virginia University High Performance Computing Spruce Knob facility.¹³ In total, 856 hr have successfully been processed through the calibration pipeline, with some data excluded due to offline processing issues.

2.2. Imaging

2.2.1. Continuum Subtraction

The production of a usable H I spectral line cube requires all radio continuum to be subtracted from it for H I analysis. There

are three basic ways to subtract the continuum in CASA (J. P. McMullin et al. 2007): *imcontsub*, *uvcontsub*, and *uvsub*. They are described as follows:

1. *Uvcontsub*. This is a purely visibility-domain-based continuum-subtraction method in which a polynomial is fit, per baseline, to the visibilities, and then subtracted off. Ideally, this method would use only the H I line-free emission channels in the fit, but this requires them to be known a priori.
2. *Imcontsub*. This is a purely image-domain-based continuum-subtraction method in which a polynomial is fit to each pixel in the image domain and then subtracted off. Ideally, this method would use only the H I line-free emission channels in the fit, but this requires them to be known a priori.
3. *Uvsub*. This method is a hybrid imaging and visibility method. In this method a continuum image is made, and then the Fourier transform of the CLEAN components is subtracted from the visibilities.

Each method has benefits and limitations associated with it, and oftentimes, combinations of the methods are used to achieve optimal image cubes. For a summary of these tasks, and their strengths and weaknesses, as implemented in the Astronomical Image Processing System (E. W. Greisen 2003), see T. J. Cornwell et al. (1992).

For the CHILES data, the use of *uvcontsub* is not optimal due to the fact that the implementation in CASA is not effective in removing sources far from the field center. Normally, this is not an issue, but in the CHILES data, we are able to detect sources beyond the first null of the VLA primary beam due to our excellent sensitivity. As a result, the use of *uvcontsub* results in residual image plane artifacts. The use of *imcontsub* is complicated by the fact that in order for it to be useful, we would need to image and CLEAN a very large image for the whole database, which is extremely computationally expensive. Additionally, the sources are affected by a standing wave, with an approximately 17 MHz period, due to a central blockage affecting aperture efficiency (P. Jagannathan et al. 2018). In order to account for this, we would need to use an extremely high-order polynomial that could very well subtract real emission. Additionally, both *imcontsub* and *uvcontsub* could slightly bias our H I source parameters; as this is a blind survey, we do not a priori know where the H I emission is and include it in spectral fits. For the stronger sources, we can correct for this later by subtracting a fit that excludes the H I channels.

This leaves us with the use of *uvsub*, which we employ at two different spatial resolutions, and in three different steps. Furthermore, rather than commit large amounts of computing resources to this effort, we grid our data, in *uv*-space, and average each session on to a common grid, using the CASA task *msuvbin*, described in K. Golap & E. Momjian (2016). This allows us to reduce our entire 220 session database to one common *uv*-grid that can then be used in the subsequent final H I cube creation. Below, we present the three-step radio continuum subtraction, entirely contained within CASA 5.6, using three iterations of *uvsub* and an iteration of *imcontsub*. The radio continuum subtraction and data combination method we implement for the CHILES data is done on a per-session basis for the continuum subtraction, and then combined in the *uv*-plane. It is performed in the following manner:

¹³ Computational resources were provided by the WVU Research Computing Spruce Knob HPC cluster, which is funded in part by NSF EPS-1003907.

1. We split off the target data from a full single session uv data set, and spectrally average the data by a factor of 4, to a resolution of 62.5 kHz (13.2 km s^{-1} at $z = 0$). This channel width is chosen because it is a compromise between the required sensitivity for H I detection, and the velocity resolution required to properly deduce kinematics of our H I detected galaxies.
2. We produce a cube spanning the entire bandwidth, with 2 MHz channels to properly sample the standing wave spectral response, for the inner 1.14° of the field, at the native angular resolution, approximately $5''$. For this image, we use a robust value of 0.5, as implemented in the `tclean` task of CASA, and 512 w-projection planes to correct for the effects of a noncoplanar array and sky curvature for the production of widefield image (T. J. Cornwell et al. 2005). We implement a targeted CLEAN, using a mask over known continuum sources that was created using a continuum image generated by the CHILES data. We CLEAN the sources within the mask down to $400 \mu\text{Jy}$, corresponding to the approximately 1σ level for a 4 hr session. The Fourier transform of the CLEAN components are then subtracted from the uv data using `uvsub`. This subtracts all of the in-field continuum sources.
3. We produce a cube spanning the entire bandwidth, with 2 MHz channels to properly sample the standing wave spectral response, but this time for a field with diameter 2.84° , and using a Gaussian uv -taper that creates a resolution element of $25''$. For this tapered image, we use 1024 w-projection planes. We CLEAN the sources within the mask down to $300 \mu\text{Jy}$, corresponding to the approximately 1σ level for a 4 hr session at this tapered resolution. The Fourier transform of the CLEAN components are then subtracted from the uv data using `uvsub`. This subtracts the strong out-of-field continuum sources.
4. We, for the last time, create an image cube with a 2 MHz channel width across the entire bandwidth, over a small region, approximately $8.6'$, covering a strong source, 4C-00.37, approximately 3° south of our field with an observed flux of 1.5 mJy at our lowest frequencies, at the native resolution. The data are phase-shifted to the position of this problematic source, and for this small image, we use no w-projection planes and robust value of 0.5. We provide a CLEAN mask over the known continuum source and CLEAN down to the $400 \mu\text{Jy}$, corresponding to the approximately 1σ level, within this mask. The Fourier transform of the CLEAN components are then subtracted from the uv data using `uvsub`. This subtracts the last remaining detectable continuum source in our field.
5. We perform a final round of RFI flagging using the statistical flagging task in CASA, `rflag`, at the 5σ level on the fully `uvsubbed` data. We then average the data for all databases on a common uv -grid, performing an on-the-fly Doppler correction before the combination, at the desired channel width, and image size, using the CASA task `msuvbin`, in the LSRK velocity reference frame. We leave out data with $|u| < 50 \text{ m}$ to eliminate horizontal stripes in the image caused by RFI and imaging artifacts from the projected baselines distributions. This uv -gridded

database is then imaged with `tclean` with the specified pixel and image size and from the `msuvbin` step, and any other required imaging parameters, such as uv -tapering and spectral averaging. This step produces an image cube that spans the entire frequency bandwidth with the required specifics.

6. We run an implementation of `imcontsub`, with a linear fit, that is run over the entire bandwidth of the cube. After this, we do any image post-processing necessary, such as that in Section 2.2.2 in order to produce a final data product.

Steps 1–4 are done with a single script and, on average, have a wall time (the elapsed time taken by computing) of 20 hr and require 58 GB of memory. Using the WVU HPC Thorny Flat,¹⁴ we can run 10 sessions simultaneously, which translates into a total wall time of 18.5 days to get all 222 sessions through steps 1–4. Step 5 takes approximately 1 hr for every session that is being averaged in the uv -plane, which, for the CHILES data, translates to approximately 9.5 days, and requires, on average, 256 GB of memory. Step 6 is completed in approximately 12 hr, and utilizes 48 GB of memory. Therefore, from beginning to completion, for the entire CHILES database, this procedure takes 28.5 days, and requires resources ranging from 48 GB of memory to 256 GB of memory. The final data product, an image cube of 1200×1200 pixels and 3680 channels, is approximately 20 GB in size.

In Figure 1, we show examples of steps 1–4 for a single 2 MHz channel, at 960 MHz, for a single 6 hr observing session. In the first panel, we show an image that represents a simple fast Fourier transform (FFT) of the data with no deconvolution or continuum subtraction. This image shows significant sidelobes from sources both within, and outside of, the 20% point of the primary beam, indicated by the black circle. In the second panel, we show the CLEANed image described in step 2, in which the in-field sources from this image are subtracted. The sidelobes to the extreme east and extreme west of the image are still present, as the sources they come from are outside this field. In panel three, we show an image after the CLEANed image described in step 3, where the sources beyond the 20% point of the primary beam have been removed with `uvsub`. The image created in this step is, in actuality, 4 times the size of the image in this panel. However, we display a zoomed-in image to just encompass the science field (the inner quarter of the image produced in step 3), to show that the sidelobes have been removed, and additionally to show a low-level north–south ripple across the center of the image caused by the source 3° to the south. Lastly, in panel four, we show a pre-subtraction image from this observing run of the source that is subtracted in step 4. This image shows how, quite remarkably, a source 3° away from field center can still be detected, and thus causes sidelobes that affect the science field. As this panel shows, this source is able to be successfully modeled, CLEANed, and subsequently subtracted, thus successfully mitigating any systematic offsets from the resulting residual sidelobes.

These panels illustrate how each step is necessary in the reduction of the prevalence of imaging artifacts. An important reminder is that this is just a single session, the full CHILES database is more than a factor 10 times more sensitive than this

¹⁴ Computational resources were provided by the WVU Research Computing Thorny Flat HPC cluster, which is funded in part by NSF OAC-1726534.

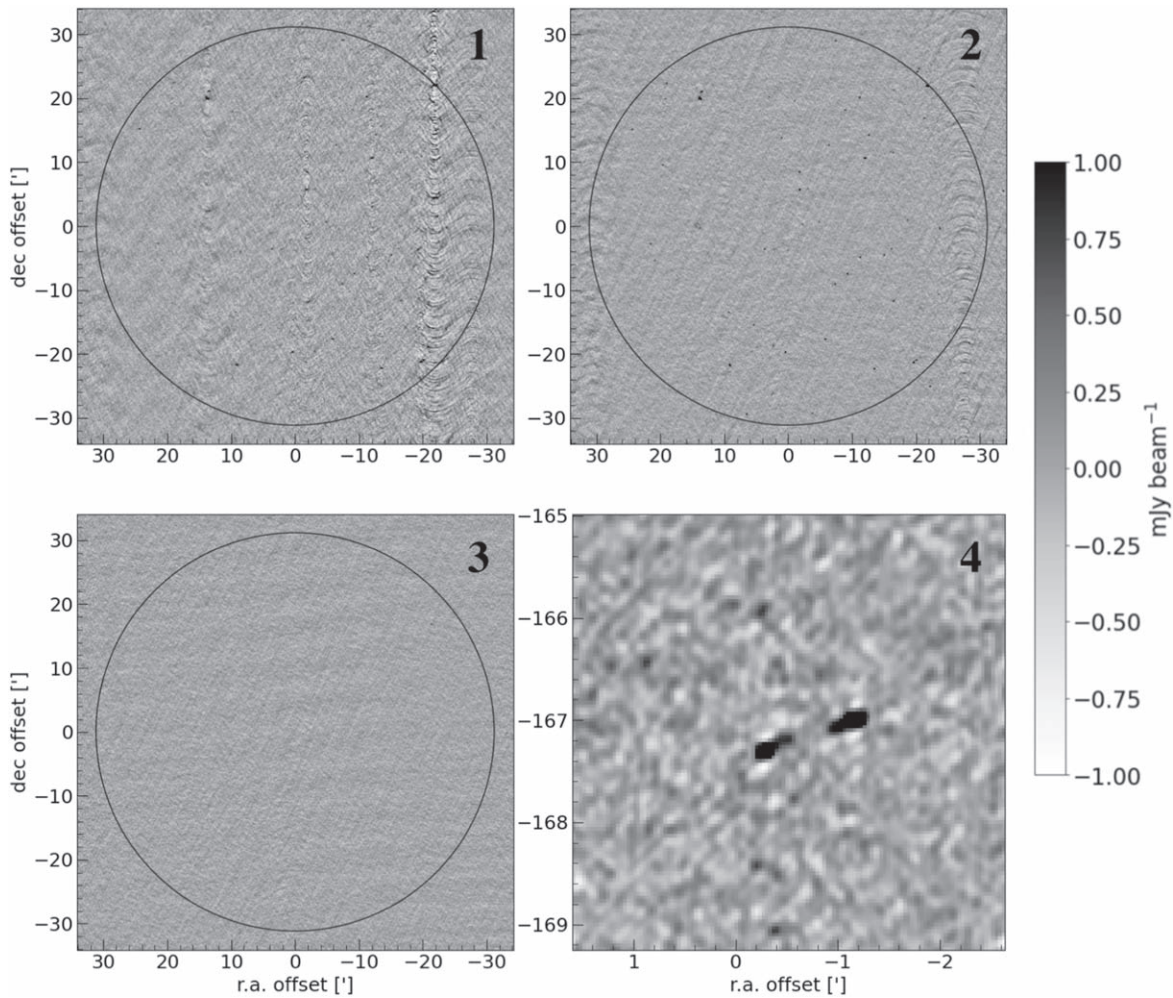


Figure 1. An example of the continuum-subtraction routine on a single 6 hr observation run, for a single channel of width 2 MHz, centered on 961 MHz. In panel 1, we show an unCLEANed image of the science field, in panels 2–4, we show the CLEANed pre-subtraction images produced in each of those steps, whose Fourier components are then subtracted from the visibilities using *uvsub*. The black circle corresponds to the 20% point of the primary beam of the VLA at these frequencies, and the gray scale is common for all images, and illustrated by the gray bar to the right. For each image, the x - and y -axes are defined by the angular offset from the CHILES field phasecenter. This changes for some images, as images 1, 2, and 3 are of field center, while image 4 is centered on a problematic source, and images 1, 2, and 4 have $2''$ pixels and image 3 has $5''$ pixels.

session. This means that any imaging artifacts seen in these images will have appeared at a much more significant level using the full database, had we not corrected for them with this routine. Additionally, in the above points, we determine the CLEAN threshold based upon the number of visibilities in a 4 hr session, the average length of a CHILES session. We use this same threshold for all databases, independent of total time, as observations are never more than a factor of 2 shorter or longer than 4 hr. For shorter observations, this may cause an increase in run time due to the lower threshold, in terms of signal-to-noise ratio, and for the longer observations, it may cause some residual continuum to remain in the image. Only the latter will affect the image quality and is mitigated by the last image-plane subtraction in step 6.

2.2.2. Additional RFI Mitigation

In the spectral line cubes produced after our imaging pipeline, there are still some lingering effects of RFI and imperfect calibration. Both of these effects are seen in the image-domain as broad stripe-like features that stretch across the images. As a result of their similarity, and the convenience

of Fourier transforms, we approach these problems with a uv -plane mitigation technique. Any feature that appears with broad image-wide characteristics will appear as localized features in a Fourier transform of the image. The uv -space RFI mitigation technique we use is an implementation of the method introduced in VLA Memo 198 (K. Golap & E. Momjian 2016), and further explored in S. Sekhar & R. Athreya (2018), optimized for use with the CHILES data. We use this adaptation on the spectral line data produced following the completion of the six steps outlined in Section 2.2.1. The adaptation we use is completely native to Python, only requires the NumPy package, is done on a per-channel basis, and is described as follows:

1. Take an FFT of the image plane data.
2. Rotate each quadrant in uv -space such that u and $v = 0$ are in the lower-left corner, and average the four quadrants. There is some redundancy as quadrants I and III, and II and IV are identical, but this does not affect the results, as the threshold is chosen after the averaging and optimized to remove artifacts and conserve source flux.

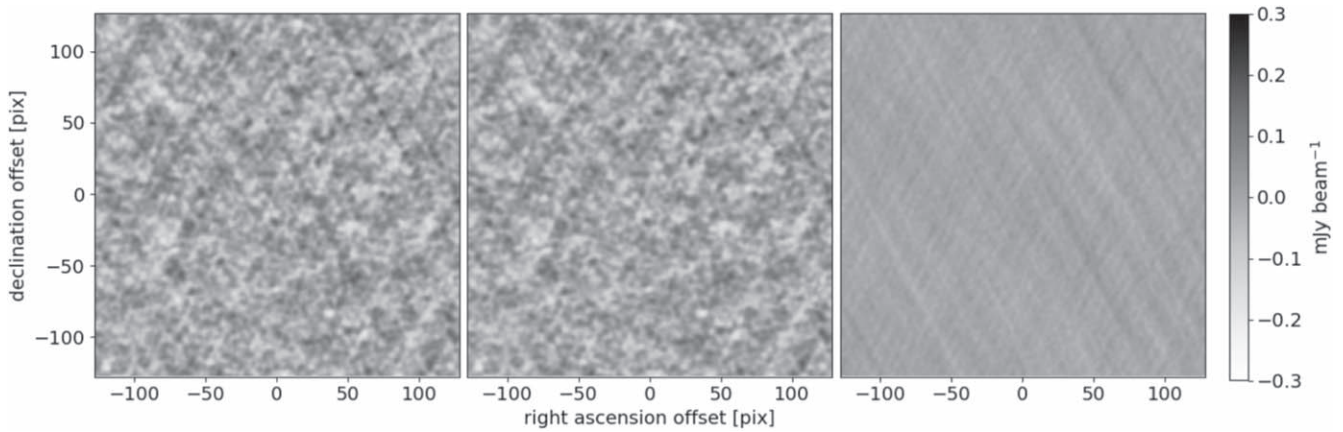


Figure 2. Left panel: a single channel input, at 1153.25 MHz, before the flagging routine, with an rms noise of $51.5 \mu\text{Jy beam}^{-1}$. Middle panel: the same data after the flagging routine has been run, with an rms noise of $49.3 \mu\text{Jy beam}^{-1}$. Right panel: the difference between the images before and after the flagging routine. Each image has a common gray scale, ranging from -0.3 to 0.3 mJy.

This optimization is verified by investigating a number of thresholds and selecting the lowest threshold that has no effect on the flux of our brightest and most extended source.

3. Bin the data by baseline length in increments of 10 m, calculate the mean and standard deviation in each bin, and if the value in a cell is $>2\sigma$ above the mean for the radial profile at the corresponding baseline length of the cell, set the point equal to zero.
4. FFT the uv -space data back into the image-domain.

This technique has been applied to the final cube, representing the aggregate of all CHILES data, whose production is described in Section 2.2.1. The technique was parallelized on a per-channel basis, and can be run on the WVU HPC, Thorny Flat, in approximately 4 hr, requiring 48 GB of memory, for a 20 GB image cube of 3680 channels and 1200×1200 spatial pixels. In Figure 2, we show the effect of this uvbin flagging method on the CHILES data. The rightmost panel of this figure shows the residual features that this technique successfully eliminates from the image. The residual features in this image are various diagonal stripes running across the image. The diagonal stripes are likely caused by a combination of two reasons: low-level RFI that the automated flagging missed as part of the calibration pipeline, and artifacts from the imperfect removal of sources at, and beyond, the half-power point of the VLA primary beam. The removal of both of these effects is expected, as they can both be similarly characterized as large-scale variations across the image. As a result, they appear similarly in the uv -plane and are successfully excised by our technique.

The rms noise in the pre-flagging image is $51.5 \mu\text{Jy beam}^{-1}$ and is slightly reduced in the post-flagging image to $49.3 \mu\text{Jy beam}^{-1}$. The artifacts visible in the difference image that were successfully removed lie at, or below, the 1σ level, and in the two images from before and after the technique, it is hard to detect the effect of the technique by eye. However, it is crucial to remove these artifacts, as they can have a profound impact on source structure and characteristics. For example, a 0.5σ ripple would introduce a $\pm 16.7\%$ error in flux density on a 3σ detection, if the detection lies on a peak or trough of the ripple.

2.3. Final Data Quality

In order to validate the output of the procedure outlined in Sections 2.2.1 and 2.2.2, we can compare the empirical results to that of expectations. If they are similar, it is an indication that the continuum subtraction has been successful. In Figure 3, we show the range of all pixel values (top row) and the ratio of measured to theoretical noise (third row) for all channels. The theoretical noise is calculated using the frequency dependent VLA System Equivalent Flux Density (SEFD) and the number of unflagged visibilities used in the cube creation. Strictly speaking, this calculation is only correct for an image cube created with naturally weighted visibilities, while our cube is made with a briggs robust value of 0.5. However, our choice of a briggs robust parameter has a minimal, $< \approx 5\%$, effect on the measured noise in an image cube (D. S. Briggs 1995). Interestingly, in the third panel, the ratio occasionally dips to slightly below 1. We theorize that this is due to the fact that the SEFD used in the theoretical calculation is not a 100% accurate measurement. This is because there is some natural variation of SEFD from antenna to antenna. As a result, we only consider a deviation in the ratio of 30% or greater as significant, as any deviation lower than that could be a result of these somewhat inaccurate uses of SEFD measurements. There is some deviation greater than this threshold at our lowest frequencies, but we believe this is due to the fact that these frequencies represent the extreme low end of the VLA L-band receivers where the SEFD is least well behaved.

Simply having noise that is comparable to theory, however, is not the only necessary indicator for science-ready data. An outstanding issue for deep radio interferometric surveys is the presence of non-Gaussian noise. Non-Gaussian noise is caused by calibration/deconvolution errors, RFI, and subsequent incomplete/incorrect continuum subtraction, which can significantly affect the scientific results of the data. To inspect this issue in the CHILES data, we show the kurtosis value (the fourth moment of a Gaussian distribution) as a function of frequency in the bottom row of Figure 3. The horizontal red line at a kurtosis value of three in the bottom row of Figure 3 indicates the expectation value for perfectly Gaussian data, with the expected variation on three, for CHILES-like data, being 0.004. Values of kurtosis that are greater than or less than three indicate that the distribution has a greater or fewer number of extreme values than expected, respectively.

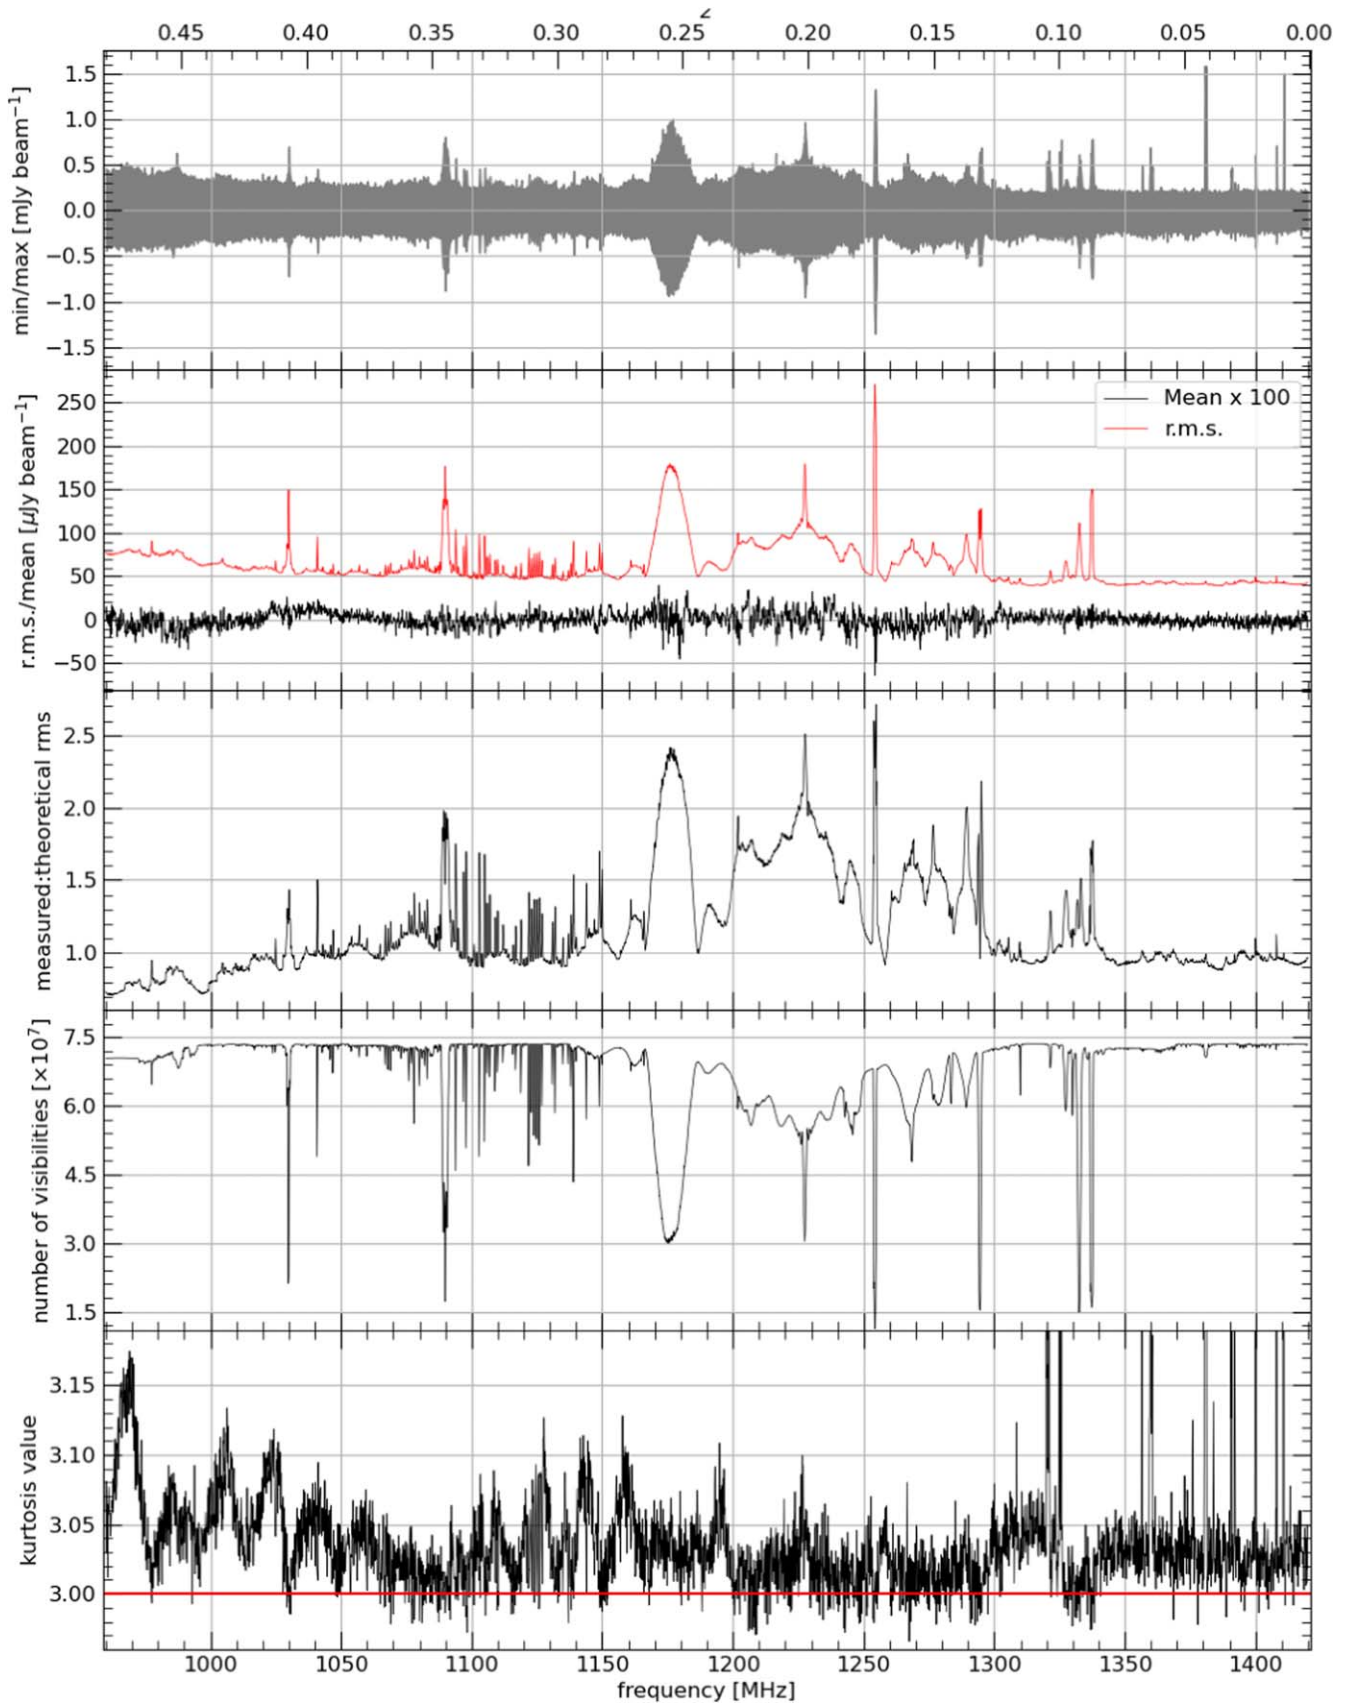


Figure 3. As a function of frequency (enumerated on the bottom axis with corresponding redshifts enumerated on the top), we show the (top row) range of all pixel values, the (second row) rms noise and mean pixel value multiplied by 100, the (third row) ratio of measured noise to theoretical noise, the (fourth row) number of unflagged visibilities used in imaging, and the (bottom row) kurtosis values of the CHILES data, following the processing outlined in Section 2.

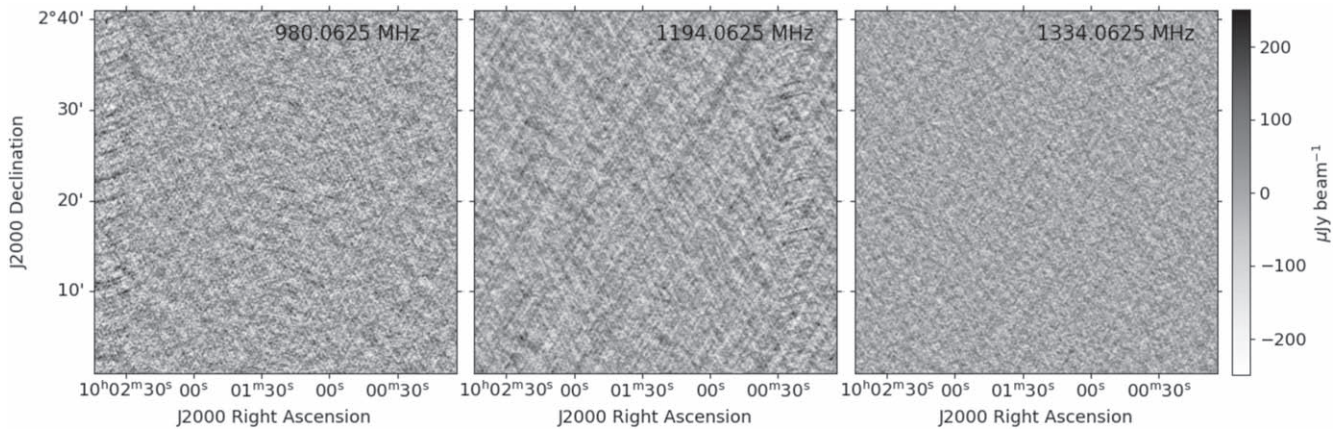


Figure 4. Three channel maps located at 980.0625 (left panel), 1194.0625 MHz (middle panel), and 1334.0625 MHz (right panel) pulled from the final cube used in our analysis. Each image has a common gray scale, ranging from -0.25 to 0.25 mJy. These represent exemplary channels of intermediate quality (left panel), poor quality, less than 10% of the data (middle panel), and exceptional quality (right panel).

(K. P. Balanda & H. L. MacGillivray 1988). The expectation for this final CHILES image cube, with perfect calibration and RFI excision, should be that the visibilities have a Gaussian distribution, except for the regions for which there is H I emission, and thus a perfectly Gaussian image in emission-free channels. The sharp spike features seen in the bottom row of Figure 3 above 1300 MHz can be attributed to known bright H I detections. The flux from the H I emission, as well as the bright sidelobe patterns, injects those image planes with pixel values that significantly deviate from Gaussian noise, and present as sharp image-plane features.

Taken in its totality, Figure 3 offers some insight into the complicated nature of the CHILES data. The kurtosis values indicate that the noise in the data departs from Gaussianity, in several frequency ranges, most notably in 960–1060, 1100–1200, and 1300–1320 MHz, with the largest deviations at the lowest frequencies. However, the noise ratio only significantly increases in the range 1150–1300 MHz. From these two facts, we can surmise that the data in the frequency range 1150–1300 MHz suffers strongly from residual RFI and imaging artifacts, but that the entire bandwidth suffers from non-Gaussian noise on a lower level. This second point is key in understanding this data set, as it emphasizes the fact that we must treat each H I detection and H I stacking result with scrutiny to ensure it is not an imaging artifact, or a combination of imaging artifacts, respectively.

The statistical overview presented by the panels in Figure 3 offer a comprehensive overview of the data, and in Figure 4, we complement this by presenting three channel maps from our final cube. Specifically, in Figure 4, we present three channel maps at three different representative frequencies. We show the data at 980.0625 MHz, which is an example of a channel that still has some continuum-subtraction residuals on high values of R.A., 1194.0625 MHz, which is an example of a channel that has noticeable RFI and continuum residuals across the entire image, and 1334.0625 MHz, which shows an ideal channel with very little artifacts. A significant portion of our data, especially at frequencies above 1300 MHz appear as immaculate as our example channel at 1334.0625 MHz. We estimate that in the range 960–1300 MHz, no more than 10% of the channels are as bad as the channel shown at 1194.0625 MHz and thus not suitable for science, with that example channel being among the most extreme examples of a nonideal channel. Note that even in the worst range,

1150–1300 MHz, useful scientific results can be obtained (K. M. Hess et al. 2019; K. M. Hess et al. 2025, in preparation). The channel maps, in general, show data of the type of quality that would be suggested by the plots in Figure 3 at the specific frequencies. In Figure 5, we show an example of two detections in our 856 hr cube. In the left panel, we see an example of a detection at 1216 MHz, which is in the range most heavily contaminated by RFI, and in the right panel, we show an example a channel in our most pristine frequency range. These images demonstrate the suitability of our final image for detecting H I, even in frequency ranges most affected by RFI.

The above-mentioned imaging and RFI mitigation techniques have been applied to all successfully calibrated data (856 hr). The data quality is slightly less than what we hoped for with a typical rms noise of $50 \mu\text{Jy beam}^{-1}$ per 125 kHz channel as opposed to the proposed $50 \mu\text{Jy beam}^{-1}$ per 31.25 kHz channel. This and its consequences will be discussed in a forthcoming paper (J. H. van Gorkom et al. 2025, in preparation). However, our results in Figure 3 demonstrate that we achieve reasonable values of measured noise based on the actual numbers of visibilities used in our cube production.

3. Ancillary Data

Throughout this work, we make use of the extensive multiwavelength data, and subsequently derived galaxy properties, made publicly available by the COSMOS collaboration (N. Scoville et al. 2007). From the public COSMOS data, we use the COSMOS2008 ID as well as the rest-frame NUV-*r* color, and derived stellar masses, calculated with the stellar population model of G. Bruzual & S. Charlot (2003) and the initial mass function of G. Chabrier (2003), as presented in C. Laigle et al. (2016). In addition to the COSMOS data, we also utilize data from the Galaxy and Mass Assembly (GAMA) survey (S. P. Driver et al. 2011). Specifically, we use the GAMA collaboration’s reprocessing of COSMOS multiwavelength data in order to derive the highest possible confidence redshifts (L. J. M. Davies et al. 2015). For the H I stacking done in this work, it is critical to have a priori redshifts with errors similar to or less than the typical H I line width in the stack, in order to ensure that we are able to correctly identify the expected line center of the H I detection, and not smear the average spectrum (N. Maddox et al. 2013). As a result, in all input galaxies for the H I stacks in this paper, we only use

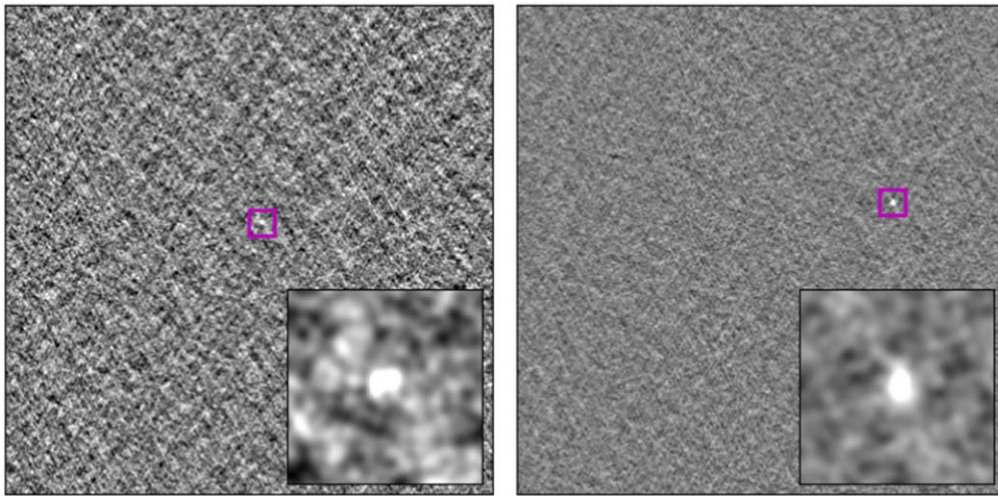


Figure 5. An example of two channels that contain emission from a previously published CHILES H I detection. Each panel represents the entire CHILES field of view, and the zoomed insets in the bottom-right corners are the $2'$ region around the detection, corresponding to the magenta boxes. Each image has a common gray scale, ranging from -0.25 to 0.25 mJy. Left panel: this is a detection at 1216.25 MHz, originally presented in K. M. Hess et al. (2019) using 178 hr of the CHILES data. Right panel: this is a detection at 1356.5 MHz, originally presented in J. Blue Bird et al. (2020), using 178 hr of the CHILES data.

galaxies with the highest-confidence redshifts, as specified in L. J. M. Davies et al. (2015).

In addition to publicly available COSMOS data, we make use of the radio continuum data from the CHILES Continuum/Polarization Survey (CHILES Con Pol). CHILES Con Pol is a 1000 hr 1.4 GHz wideband full polarization radio continuum deep field that was observed commensurate with the CHILES H I deep field, utilizing the four baseline board pairs of the VLA correlator that were not used by CHILES. Further details on the calibration, imaging, and data quality can be found in C. A. Hales & N. Stephenson (2019) and N. M. Luber et al. (2025). We utilize this data in order to estimate star formation rates, in the methodology described below in Section 4.2.

4. Stacking Methodology

4.1. Neutral Hydrogen Emission

In order to produce an accurate measurement of the average neutral hydrogen content of a sample of galaxies, we perform a cubelet stack with the following specifications (for a detailed description of cubelet stacking, see Q. Chen et al. 2021a; for a demonstration of its capabilities, see Q. Chen et al. 2021b):

1. Extract a cubelet centered in position and frequency on a source. We chose to extract cubelets with 64 pixels in R. A. and decl., and 89 frequency channels, corresponding to a cube with dimensions $128'' \times 128'' \times 11.125$ MHz. At the lowest and highest redshifts, we use in our stack, 0.09 and 0.47, this corresponds to physical dimensions of $236 \text{ kpc} \times 236 \text{ kpc} \times 2565 \text{ km s}^{-1}$ and $764 \text{ kpc} \times 764 \text{ kpc} \times 3474 \text{ km s}^{-1}$, respectively.
2. Perform a test on the cubelet to see if the noise is sufficiently Gaussian. We define this as having no more than 3 times the expected number of pixels greater than $\pm 3\sigma$ for a normal distribution. Cubelets that do not meet this criterion are not used in the subsequent stacking. These rejected cubelets account for approximately 15% of the total number of galaxies within the sample.

3. Perform the appropriate primary beam correction by multiplying the entire cubelet by the primary beam correction factor for the predicted central frequency and pixel of the source, and multiply the data by the distance (in Mpc) squared to convert the cube into line luminosity units.
4. Produce an inverse-variance-weighted average stacked cube. In order to calculate the inverse variance weight, we use the cube before it is converted into line luminosity units. This is to avoid down-weighting the sources at the higher redshifts in the probed range. For a description of the different possible distance weights and their effects, see J. Delhaize et al. (2013).
5. Produce an inverse-variance-weighted average point-spread function (PSF) for the channels used in the stack. This is done extracting PSF cubelets in the same frequency ranges as the data cubelets and averaging them using the same weights as in the previous step.
6. Perform a Hogbom CLEAN (J. A. Högbom 1974) on the stacked cube using the average stacked PSF on the channels with expected emission. This is done by inspecting the cube and creating a mask over the identified emission. We then clean the flux in this mask down to the 1σ level.
7. Fit a second-order polynomial to the channels identified as emission free and subtract it off. This is to remove any remaining imaging artifacts that may have been prevalent in the stack.

We chose to do a cubelet stacking, as opposed to spectral stacking, due to the fact that CHILES does resolve sources throughout our redshift range, and in order to produce an accurate H I measurement, it is essential to CLEAN the stacked detection (Q. Chen et al. 2021a). The emission of individual galaxies in the stack cannot be cleaned due to the low signal to noise. As the stack raises the signal to noise, it becomes essential to CLEAN the image. The H I mass for the stacked detection is calculated by summing the flux in the region with emission. We define the H I mass error as the standard

Table 2
Stacking Results from CHILES

Redshift	NUV- <i>r</i> (mag)	N_{gal}	Average H I Mass ($10^9 M_{\odot}$)	Average SFR ($M_{\odot} \text{ yr}^{-1}$)	Average Stellar Mass ($10^9 M_{\odot}$)	Significance ^a	Relevant Figures
0.43	−1 to 3	308	2.8 ± 0.9	0.37 ± 0.02	11.4	3.6	5 and 7
0.37	−1 to 3	315	2.7 ± 0.9	0.24 ± 0.02	9.1	3.7	5 and 7
0.30	−1 to 3	527	1.1 ± 0.4	0.21 ± 0.01	8.5	2.2	5 and 7
0.15	−1 to 3	166	1.7 ± 1.8	0.14 ± 0.02	8.1	2.5	5 and 7
0.43	3–6	96	<1.3	...	50.6	...	6 and 7
0.37	3–6	134	<0.6	...	46.8	...	6 and 7
0.30	3–6	175	<0.4	...	35.3	...	6 and 7
0.13	3–6	58	<0.5	...	25.1	...	6 and 7
0.42	−1 to 3	142	3.2 ± 1.2	0.66 ± 0.04	27.4	3.6	8 and 9
0.32	−1 to 3	155	1.9 ± 0.6	0.51 ± 0.01	25.8	3.4	8 and 9
0.16	−1 to 3	33	1.1 ± 0.5	0.42 ± 0.05	29.4	2.1	8 and 9
0.41	−1 to 3	340	1.4 ± 0.6	0.19 ± 0.02	3.5	5.5	8 and 9
0.31	−1 to 3	511	1.4 ± 0.5	0.12 ± 0.01	3.5	3.9	8 and 9
0.15	−1 to 3	135	1.1 ± 0.4	0.07 ± 0.02	3.2	3.8	8 and 9

Notes. The redshift ranges, color ranges, average galaxy properties, significances, and relevant figures for the stacked detections presented in this work. Above the horizontal line is for the stacks that consider all galaxies with stellar masses $>10^9 M_{\odot}$, and below the line are shown the results when we divide the blue galaxy sample into two stellar mass bins. When calculating average stellar masses and average star formation rates, we only include the galaxies used in the H I stack and use the same weighting scheme as in the H I stack.

^a We define the significance of the H I stacked detection by dividing the peak flux in the spectra by the rms noise of the spectra, as measured in the line-free channels. These spectra correspond to the integrated H I profiles shown in Figures 7 and 10.

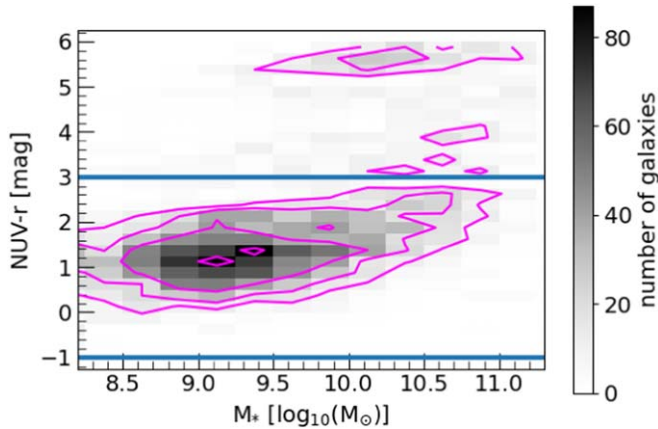


Figure 6. The rest-frame NUV-*r* color as a function of stellar mass for all galaxies within the CHILES field. Both the color scale and the contours indicate the number of galaxies with each respective rest-frame NUV-*r* color and stellar mass combination. The solid blue lines indicate the boundaries of our blue color cutoff.

deviation of a distribution of total fluxes in 100 emission-free regions from the H I stack cubelet.

In principle, our stacked measurement of average H I mass could be affected by unaccounted for galaxies that may lie within the beam of an input galaxy used in the stack. This effect is referred to as source confusion. In M. G. Jones et al. (2016), the authors modeled the effect of stacked source confusion for several H I surveys, including CHILES, and they found that confusion introduces an error at the level of approximately $2 \times 10^8 M_{\odot}$. This is approximately an order of magnitude less than the H I masses we report, and a factor of 2–9 less than the error associated with the stacks, thus making it an insignificant source of error on our reported H I measurements. In addition, our high-resolution observations have the potential to spatially smear the stacked detection. This,

and the fact that the stacked galaxies are initially unCLEANed, could have the effect of systematically lowering the measured H I mass.

4.2. Radio Continuum Emission

Radio continuum emission at 1.4 GHz is dominated by synchrotron emission, which mostly arises from the effects of recent star formation or active galactic nuclei (AGNs). As is the case with radio spectral line data, one can stack radio continuum data for a sample of galaxies that are most likely to have star formation as the cause for radio continuum emission and derive an average flux density measurement and star formation rate (R. L. White et al. 2007). In order to calculate average star formation rates for the galaxies used in the H I stack, appropriately not including the galaxies that were rejected from the H I stack, we stack the radio continuum data from the CHILES Continuum Polarization Survey and use the scaling relationship between 1.4 GHz radio flux and star formation rate in E. J. Murphy et al. (2011). This technique allows us to calculate a statistical star formation rate using the same techniques as our H I mass measurements, ensuring that our reported values have been produced in a self-consistent manner.

The manner in which we stack radio continuum is parallel to that of the H I stacking with a couple of key differences. First, as radio continuum is a single frequency intensity measurement, our stack comprises of 2D images, as opposed to the 3D cubelets from the H I stack. Second, in order to measure the parameter of interest, the rest 1.4 GHz emission, we scale the radio continuum flux of each source by using the following equation:

$$F_{1.4 \text{ GHz, rest}} = F_{1.4 \text{ GHz, obs}} \times (1 + z)^{\alpha} \quad (1)$$

where $F_{1.4 \text{ GHz, obs}}$ is the observed flux in the map, and α is taken to be -0.75 , a value typical for radio continuum due to star formation (M. S. Yun & C. L. Carilli 2002). Lastly,

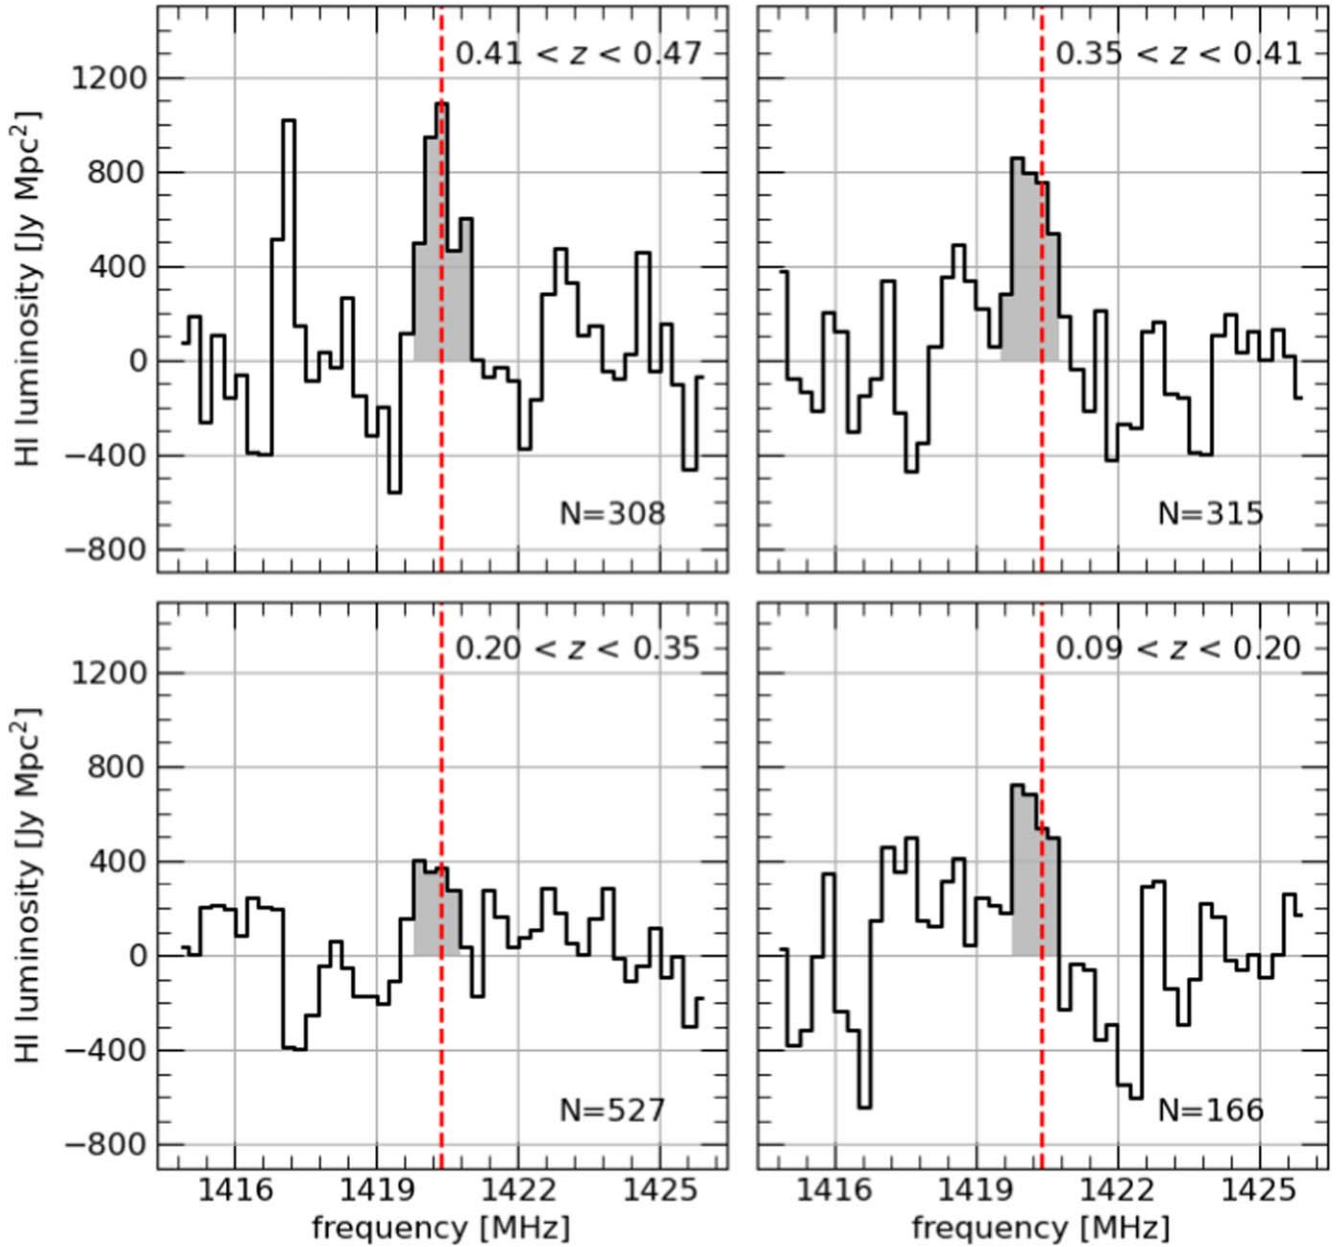


Figure 7. The stacked H I spectra for blue galaxies in four redshift ranges continuous from $0.09 < z < 0.47$. In each panel, we indicate the redshift range (top right), number of galaxies included in the spectrum (bottom right), the rest frequency of H I (red dashed line), and the region integrated over for the H I mass calculation (gray shaded region).

contrary to the H I data, the radio continuum image was CLEANED very deeply, a $5 \mu\text{Jy beam}^{-1}$ threshold (which is 3.5 times the rms noise level of $1.4 \mu\text{Jy beam}^{-1}$). Our stack is, by necessity, a mix of cleaned and uncleaned sources, but since almost all continuum sources are unresolved at this resolution, we think this is an acceptable strategy. Other than those two aspects, each step in the stack is identical, as well as the weighting schemes.

Using our stacked image, we can then measure the central flux and scale this value to calculate star formation rates. This is done by leveraging the far-IR (FIR)–radio correlation and the subsequent relationship between FIR flux and star formation rates (M. S. Yun et al. 2001; E. F. Bell et al. 2003). Specifically, we use the following equation (Equation (17) in

E. J. Murphy et al. 2011):

$$\left(\frac{\text{S.F.R.}}{M_{\odot} \text{yr}^{-1}} \right) = 6.35 \times 10^{-29} \left(\frac{L_{1.4 \text{ GHz}}}{\text{erg s}^{-1} \text{Hz}^{-1}} \right) \quad (2)$$

for our calculation of average star formation rate. One limitation of our procedure is that we do not reject radio flux that could be arising from AGNs. We find that approximately 40% of galaxies show at least some level of AGN activity, meaning that the radio flux is not necessarily dominated by AGNs, but rather AGN activity is present, which could systematically elevate our measured star formation rates (H. B. Gim et al. 2025). Additionally, our flux measurements could be limited by source confusion or blending, which would systematically raise our

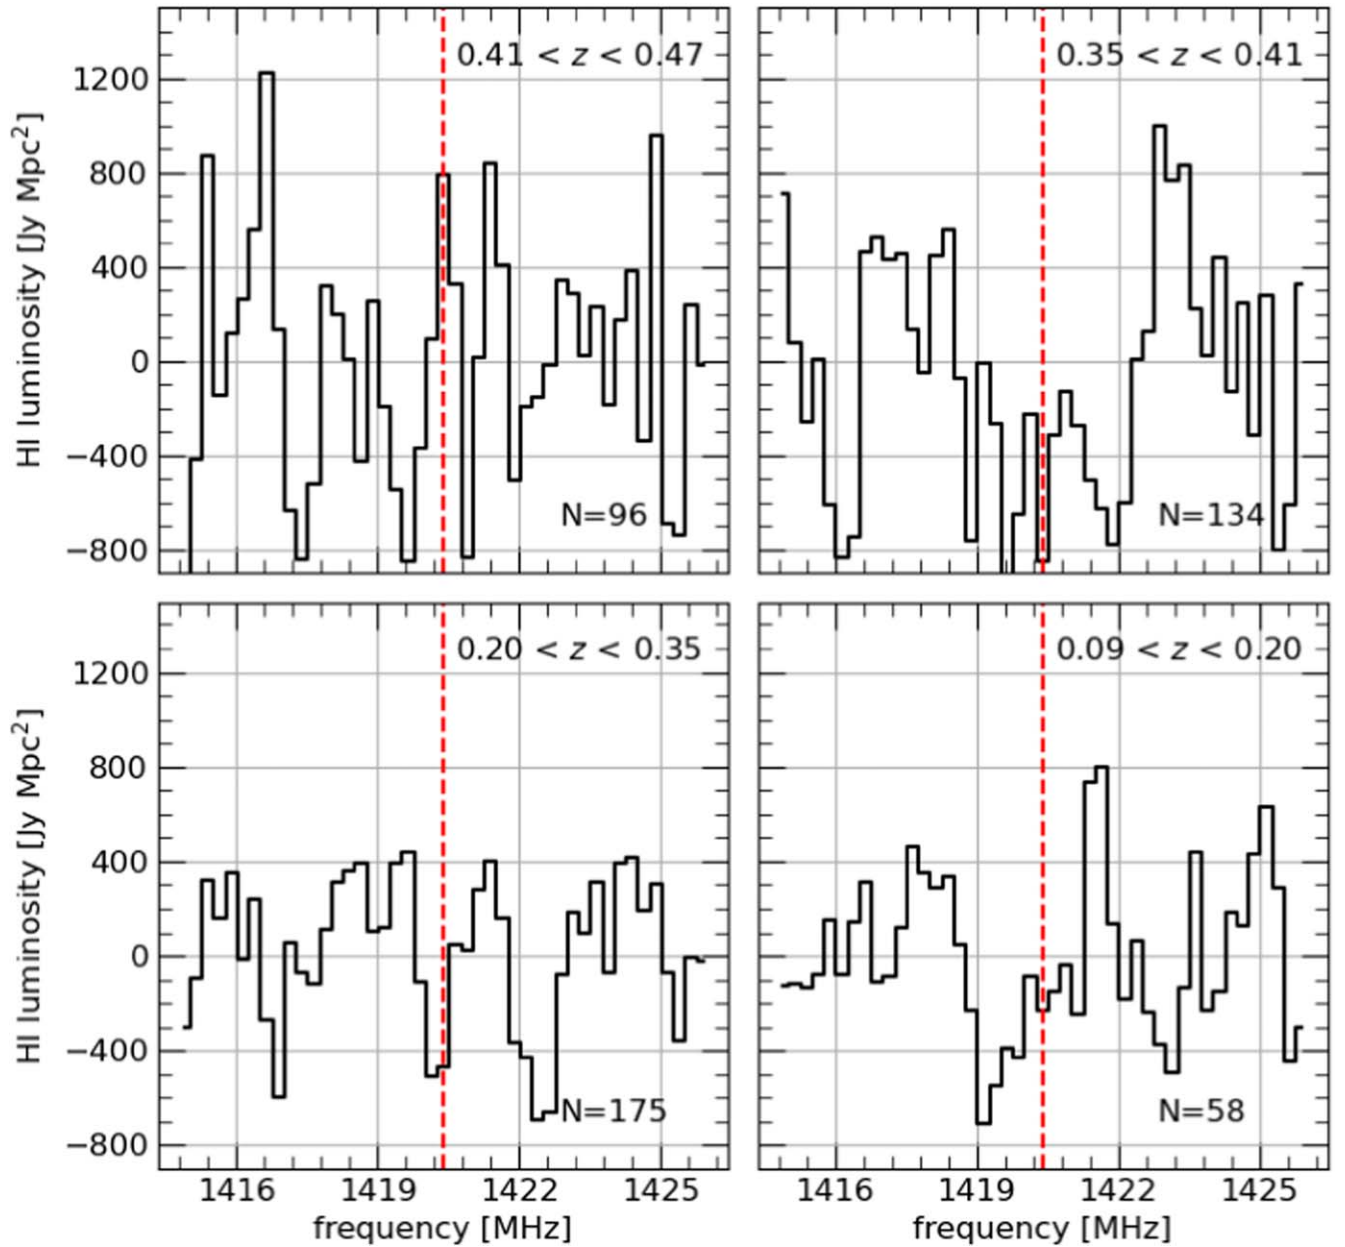


Figure 8. The stacked H I spectra for red galaxies in four redshift ranges continuous from $0.09 < z < 0.47$. In each panel, we indicate the redshift range (top right), number of galaxies included in the spectrum (bottom right), and the rest frequency of H I (red dashed line).

measurements, but are unlikely to be a major source of error due to the high sensitivity and resolution of CHILES Con Pol.

5. Results

The wide bandwidth and sub- $100 \mu\text{Jy beam}^{-1}$ noise of the CHILES database processed in the manner described in Section 2, combined with the H I stacking procedure outlined in Section 4.1, allow us to directly probe the average H I content of galaxies out to a redshift of 0.47, in contiguous redshift bins. All results presented in the following section are summarized in Table 2.

5.1. Blue versus Red Galaxies

To begin our investigation, we probe the evolution of average H I content for galaxies in four redshift bins, chosen to

be approximately equal in redshift space and number of galaxies and without any a priori selection based on LSS, spanning $0.09 < z < 0.47$. Specifically, in each bin, we compare blue galaxies, defined as having a rest-frame NUV- r color in the range -1 to 3 , and red galaxies, defined as having a rest-frame NUV- r color in the range 3 – 6 , and stellar masses in the range $10^{9-12.5} M_{\odot}$. In Figure 6, we demonstrate how this color criterion separates the blue sequence versus the red sequence galaxies, allowing us to properly separate and investigate the evolution of the two different populations. Additionally, we choose our stellar mass criteria because COSMOS is complete for these stellar masses for our redshift bins (I. Davidzon et al. 2017).

In Figure 7, we show the stacked spectra for the blue galaxies, and in Figure 8 we show the stacked spectra for the red galaxies. For each of the blue galaxies, we detect and are

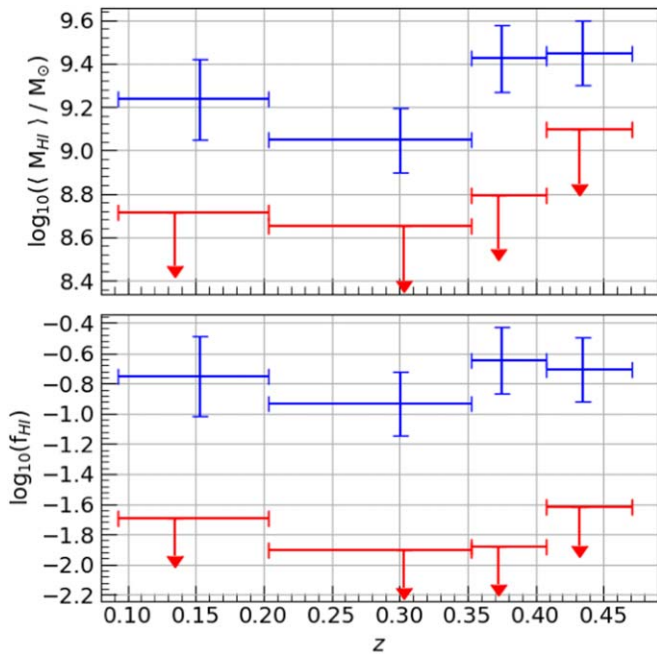


Figure 9. The stacked H I mass (top panel) and H I gas fraction (bottom panel) as a function of redshift, for the blue galaxies (blue markers) and red galaxies (red markers). In all plots, the x -axis error bars correspond to the redshift bin boundaries, and the y -axis error bars indicate the errors calculated from the H I measurements. Galaxies that do not have an H I measurement are plotted as upper limits and are shown with downward-pointing arrows.

able to determine an H I mass. However, for the red galaxies, each of our stacks resulted in a nondetection, and instead, an estimate for the H I mass upper limit. To produce each H I spectrum, we integrate over the equivalent of a 60 kpc box at the average redshift in each bin. In each panel, the gray shaded region indicates the channels identified to have emission. This identification is determined by finding the first channels at both positive and negative velocities, with respect to rest-frame line center, that cross the expected 1σ value, and including all channels that lie between. These channels are then used to calculate the average H I mass.

With our stacked measurements in hand, we can now investigate the differences in H I content for galaxies over the past 5 Gyr in red and blue galaxies. In Figure 9, we show the H I mass (top panel) and H I gas fraction measurements (for the blue galaxies) and upper limits (for the red galaxies) as a function of redshift. The most obvious result from these measurements, is that the blue galaxies have a higher H I gas content, and that this relationship persists across the entire redshift range that we probe. Second, the H I content of blue galaxies with our stellar mass criterion do not significantly evolve in this redshift range. Each H I mass and H I gas fraction is reasonably within the error of the other measurements, which indicates no significant change in the average H I gas content of massive blue galaxies in this redshift range. However, as a note, the H I measurement for blue galaxies in the second-lowest redshift bin appears to scale with the noise. This is not surprising, given that this is among our lowest significance detection, and suggests that there may be some systematics affecting the data here. Acknowledgment of this caveat further restricts us from over-interpreting these results and claiming any evolution.

5.2. The Effect of Stellar Mass

The findings of the aforementioned section prompted a probe into the effect that mass may have on the redshift evolution of the H I content of blue galaxies. To explore this relationship, we stack blue galaxies, defined as having an NUV- r color in the range -1 to 3 , in three redshift bins spanning $0.09 < z < 0.47$, in two different mass bins. For our two different mass bins, we selected two mass bins with cutoffs of stellar masses in the ranges $10^{9-10} M_{\odot}$ and $10^{10-12.5} M_{\odot}$ as our intermediate- and high-mass mass bins, respectively. In Figure 10, we present the stacked H I spectra, calculated in the same way as outlined above, for the blue galaxies in the high-mass bin (top panel) and intermediate-mass bin (bottom panel). Additionally, in Figure 11, we present the stacked radio continuum maps for the blue galaxies in the high-mass bin (top panel) and intermediate-mass bin (bottom panel). The H I stacks, specifically the lowest-redshift high-mass bin and highest-redshift low-mass bin, are redshifted with respect to predicted line center. Both of these detections suffer from either negative features on the blueshifted side, low signal-to-noise, or comparable positive peaks. This does raise the possibility that these detections may be affected by systematic errors altering derived measurements. Additionally, low signal-to-noise ratio measurements suffer a larger uncertainty in the true measured line center, further obfuscating the interpretation of measurements of a velocity offset or width. As a result, below we concentrate specifically on trends across all redshift bins in order to gain an understanding on what these measurements mean.

To investigate the possible differences in redshift evolution for our two stellar mass bins, in Figure 12, we present the H I mass, gas fraction, star formation rate, and specific star formation rate for the high-mass bin (green markers) and intermediate-mass bin (black markers) as a function of redshift. In these plots, we define f_{HI} as the ratio of H I mass to stellar mass, and specific star formation rate as the ratio of star formation rate to stellar mass. These figures make clear that there is a clear difference in the redshift evolution of galaxies in these different stellar mass ranges.

To begin, let's focus on the redshift evolution of the average H I gas content. In the top panel of Figure 12, we see that for the intermediate-mass bin, there is little to no evolution of the H I gas mass—specifically, only an increase of a factor of 1.3 that is not of statistical significance. For the high-mass galaxies, we see a more noticeable increase of a factor of 3.0, with the highest-redshift bin measurement beyond the statistical errors of the lowest-redshift bin, indicating a significant difference. For the H I gas fraction, there is a similar significant increase at the level of a factor of 3.0 for the massive galaxies, and no significant change in the intermediate-mass galaxy bin. These results show a clear decrease in H I mass for the most-massive galaxies, and no trend for intermediate-mass galaxies, in the redshift range, $z = 0.09$ – 0.47 .

As for the star formation rate evolution, we see the inverse relationship. The higher-mass galaxies have consistently higher star formation rates, as is generally found for galaxies on the star formation main sequence (A. Renzini & Y.-j. Peng 2015). However, the relative redshift evolution of these galaxies shows some interesting variation. For the intermediate-mass galaxies, we see an increase of star formation rate and specific star formation rate by factors of 2.9 and 2.6, respectively. For the high-mass galaxies, we see an increase of star formation

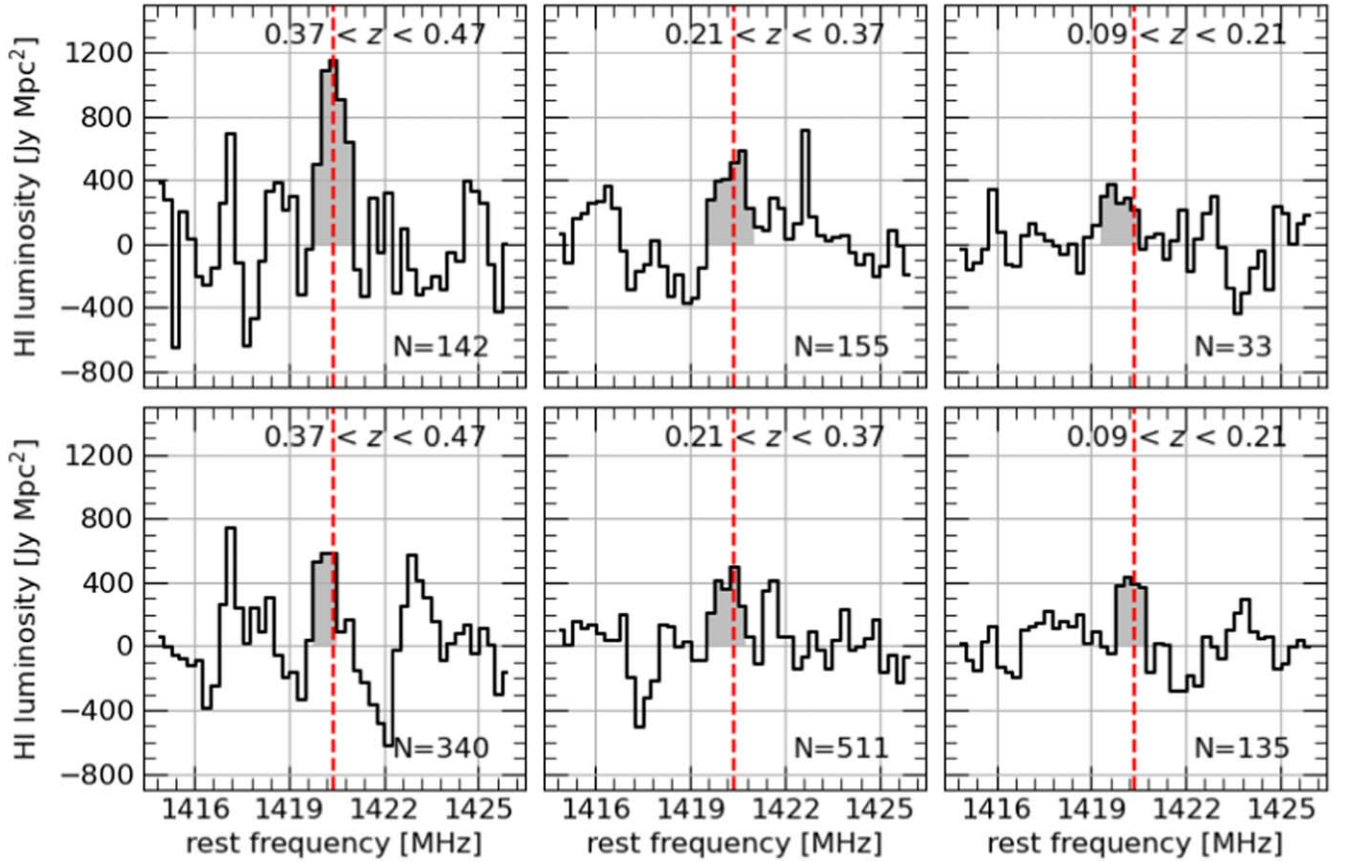


Figure 10. The stacked H I spectra for blue galaxies in three redshift ranges continuous from $0.09 < z < 0.47$, in the stellar mass range $10^{10-12.5} M_{\odot}$ (top panel) and $10^{9-10} M_{\odot}$ (bottom panel). In each panel, we indicate the redshift range (top right), number of galaxies included in the spectrum (bottom right), the rest frequency of H I (red dashed line), and the region integrated over for the H I mass calculation (gray shaded region).

rate and specific star formation rate by factors of just 1.5 and 1.6, respectively.

6. Discussion

Our results demonstrate that the redshift evolution of the H I mass content of our sample of galaxies is not scaling with the redshift evolution of the star formation rate for these galaxies. This difference suggests that different physical mechanisms could be responsible for the regulation of star formation rate for the different stellar masses we probe. The two clearest candidates for these mechanisms are some inefficiency in the H I to H₂ conversion, or some inefficiency in the ability of galaxies to replenish their H I gas reservoirs via accretion. In order to understand the specific role that H I has in the evolution of galaxies, we will place our results within the context of the star formation main sequence of galaxies.

The star-forming main sequence of galaxies is the relationship between the star formation rate and stellar masses of star-forming galaxies, and is often used as a benchmark by which one can study the evolution of galaxies (J. S. Speagle et al. 2014; N. Scoville et al. 2017; M. Ramatsoku et al. 2020; A. Bera et al. 2023b). For our study, we will look in particular at the relationship between H I gas fraction and specific star formation rate. In order to compare our results to well-known calibrations from the local Universe, we will compare our measurements to those in B. Catinella et al. (2018), where they present fits to this relationship using the rich database of H I masses, stellar masses, and star formation rates, as well as

many other derived quantities, from the xGASS survey. In Figure 13, we show the H I gas fraction as a function of specific star formation rate for the galaxies discussed in Section 5.2, where the colors correspond to the two different stellar mass bins, and the marker size increases with increasing redshift. The results from B. Catinella et al. (2018) are shown as the blue shaded region, where the dark-shaded region corresponds to the fitted relationship, with errors as reported in B. Catinella et al. (2018), and the light-shaded region corresponds to 0.5 dex from the mean, which is the characteristic scatter around the average value.

Interestingly, the results in Figure 13 are suggestive of different redshift evolution for these two samples of galaxies, albeit an evolution still within the scatter of the star-forming main sequence. The intermediate-mass galaxies in the highest-redshift bin are well within the expected value of the star-forming main sequence, and at decreasing redshift are becoming increasingly H I rich, as evident by the fact that they are moving above the star-forming main sequence. For the high-mass galaxies, we find that they similarly start well within the scatter of the star-forming main sequence, but at decreasing redshift are becoming increasingly H I poor, as is evident by their moving below the star-forming main sequence. We point out that all of our data points lie within the scatter of the relationship presented in B. Catinella et al. (2018), and that the interesting difference between the two mass bins lies in the differing gradients in redshift between the mass bins. This difference in gradient suggests that both of these samples

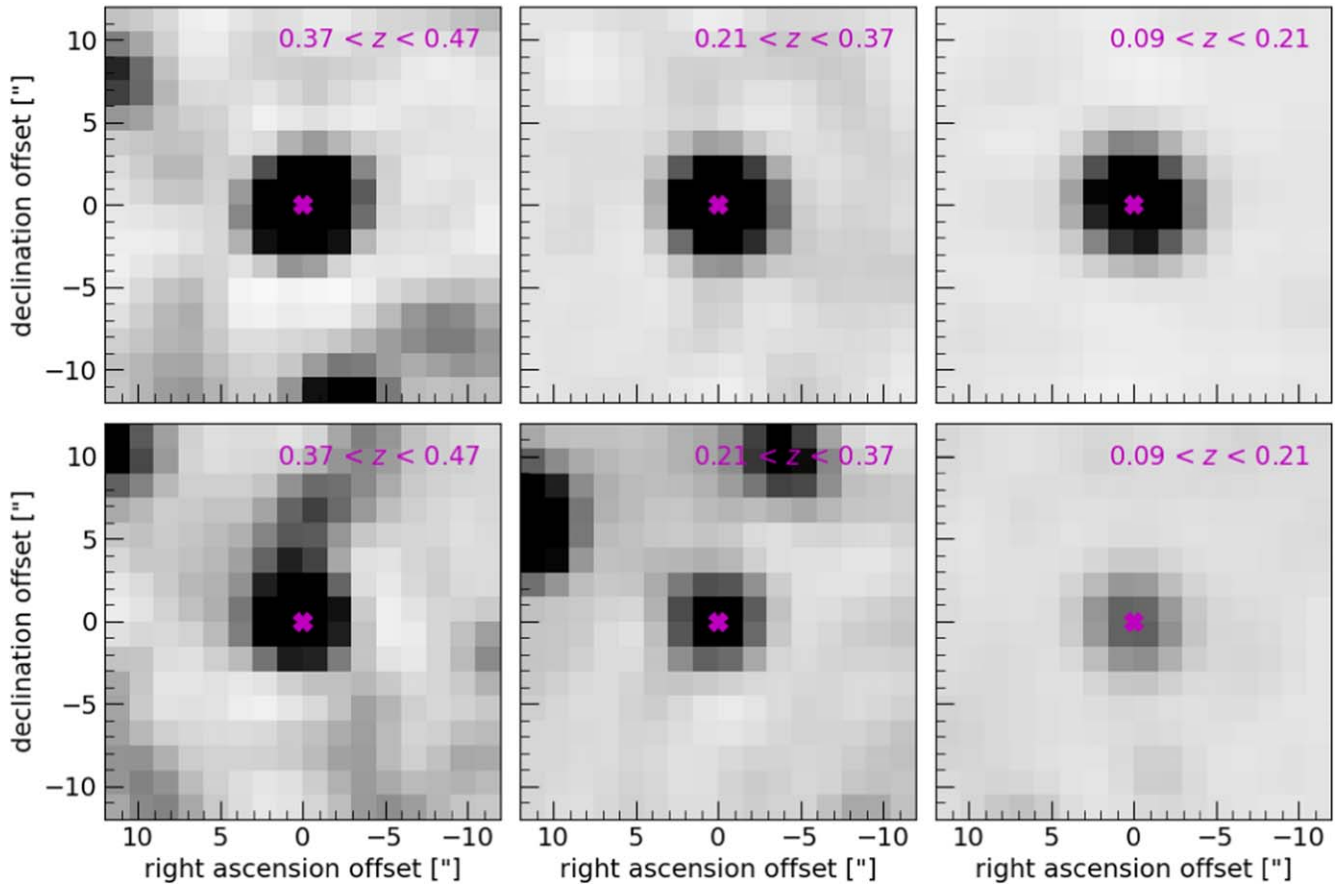


Figure 11. The stacked radio continuum for blue galaxies in three redshift ranges continuous from $0.09 < z < 0.47$, in the stellar mass range $10^{10-12.5} M_{\odot}$ (top panels) and $10^{9-10} M_{\odot}$ (bottom panels). In each panel, we indicate the redshift range (top right), and the color scale runs from -4 to $20 \text{ Jy Mpc}^2 \text{ beam}^{-1}$ (top) and -3 to $8 \text{ Jy Mpc}^2 \text{ beam}^{-1}$ (bottom). The magenta cross in the center indicates the average optical position for each galaxy in the stack.

appear to be evolving in redshift, but in a manner opposite of one another.

First, let's discuss the redshift evolution of the intermediate-mass galaxies. Figures 12 and 13 illustrate that the blue galaxies in this mass range are maintaining their supply of cold atomic gas, but the star formation rate is nonetheless declining. These two facts taken in conjunction illustrate that inefficient accretion from the circumgalactic and intergalactic media are not responsible for the decline in star formation rate. Rather, it must be that the neutral hydrogen is not being effectively converted into molecular gas and subsequently formed into stars. For the high-mass galaxies, Figures 12 and 13 demonstrate that the blue galaxies in this mass range are failing to replenish their cold gas reservoirs. This suggests that these massive galaxies are not efficiently accreting gas from the circumgalactic and intergalactic media and subsequently not providing the raw fuel for conversion into molecular hydrogen and subsequently stars. However, for these high-mass galaxies, it could also be the case that the H I to H_2 conversion becomes less efficient, as is our proposed scenario for the intermediate-mass galaxies. However, our observations do not allow us to properly discriminate between these two effects for our results.

Several recent studies have used stacking techniques to study the evolution of H I mass content evolution out to these intermediate redshifts (e.g., A. Bera et al. 2019, 2023a; A. Bianchetti et al. 2025); note that A. Chowdhury et al. (2022) studied H I content at redshifts out to 1.3. Additionally, in A. Bera et al. (2023b) and A. Chowdhury et al. (2023)

studied the gas accretion rate and molecule formation for redshifts out to 0.35 and 1.3, respectively. In both of these studies, they used the respective H I mass to stellar mass scaling relationship they derive at different redshifts, known scaling relationships involving molecular gas content, and the assumption that galaxies evolve along the main sequence. Interestingly, A. Chowdhury et al. (2023) found that for their sample of galaxies, it is a combination of lower net gas accretion and inefficient atomic to molecular conversion that drives the decrease in star formation rate between redshift 1 and 0. In contrast, A. Bera et al. (2023b) found that their redshift range, 0.35–0, is solely the inefficient conversion of atomic to molecular gas. A. Bianchetti et al. (2025) derived the $M_{\text{H I}}-M_*$ relationship for exclusively star-forming galaxies at intermediate redshifts, and found that for this sample of galaxies, it seems to be the case that H I is not being accreted quickly enough to replenish their gaseous reservoirs. Within this context, our results do not prove or disprove either scenario, but instead show there is a difference in the evolution of H I content of galaxies of different masses, and different physical scenarios can explain this evolution. Therefore, in order to confidently discuss the evolution of galaxies, we need to separate them by stellar mass and star formation rate in order to form a coherent picture.

As we begin to have meaningful statistical results of average H I content out to higher redshifts, there are several different manners in which more observational data can help us understand the answer to our posed question. First, and most

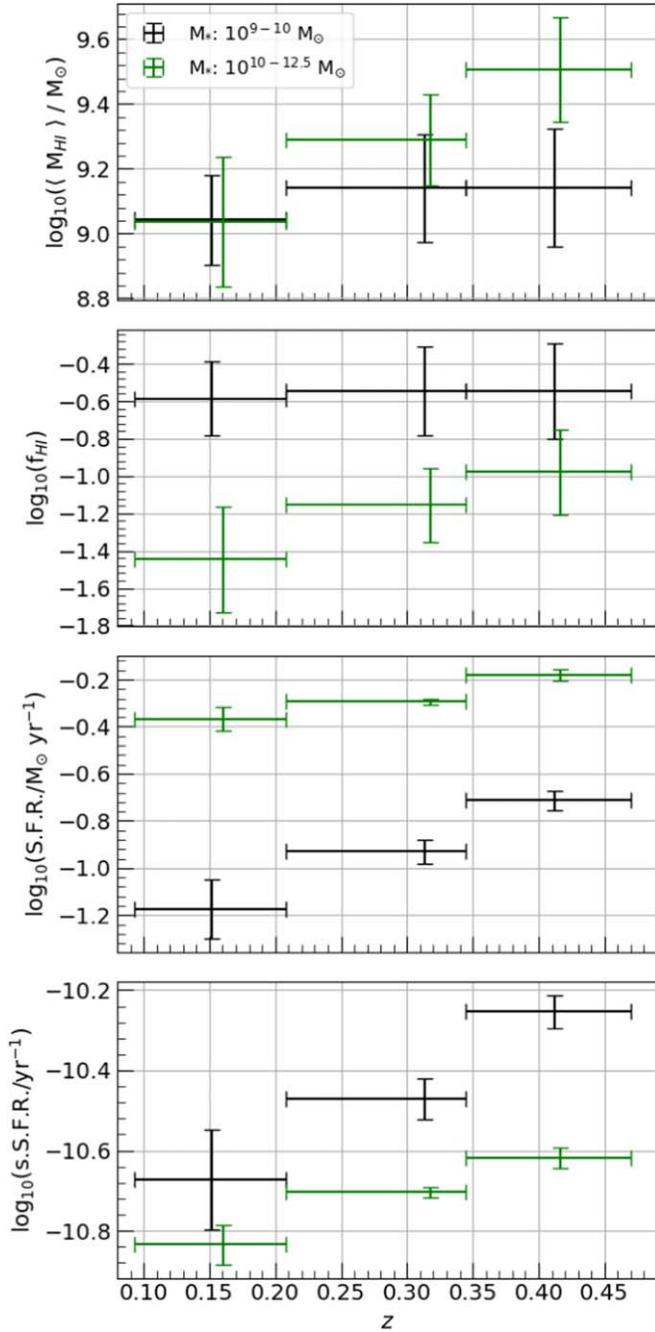


Figure 12. The derived properties for blue galaxies in three redshift ranges continuous from $0.09 < z < 0.47$, in the stellar mass ranges $10^{10-12.5} M_{\odot}$ (green) and $10^{9-10} M_{\odot}$ (black).

obviously, is the direct detection of neutral atomic and molecular gas content in individual galaxies across redshift space. Although current surveys, such as CHILES, Looking At the Distant Universe with the MeerKAT Array (LADUMA; S. Blyth et al. 2016), and MIGHTEE-HI (N. Maddox et al. 2021), will be able to contribute to the literature on this front, they will lack the numbers of galaxies necessary to provide a robust statistical basis to prove these claims. Such direct detections in neutral hydrogen will likely have to wait until the era of the Square Kilometer Array. Additionally, large-scale dust mass and H_2 surveys will complement these upcoming SKA surveys by contributing to the total knowledge of the

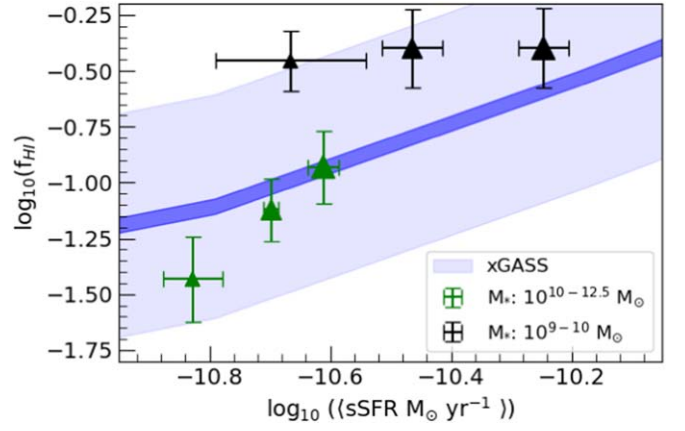


Figure 13. The H I gas fraction as a function of specific star formation rate for the galaxies in the CHILES field (individual points with error bars) and the galaxies in the xGASS survey (blue shaded regions). For the CHILES galaxies, the colors correspond to the two stellar mass ranges $10^{10-12.5} M_{\odot}$ (green) and $10^{9-10} M_{\odot}$ (black), and the sizes of the markers correlate with redshift, where larger marker size means higher redshift. The dark-blue shaded region corresponds to the fit, and error on the fit, for the H I gas fraction as a function of specific star formation rate. The light-blue shaded region indicates the typical scatter around this relationship, as reported in B. Catinella et al. (2018).

baryon cycle for these galaxies. However, this study demonstrates that future surveys should have galaxies of all masses, as the baryonic physics taking place within them that is responsible for the decline in star formation may differ between masses.

7. Summary

In this work, we have utilized the full CHILES survey to investigate the evolution of the average H I content of galaxies over the continuous redshift range $0.09 < z < 0.47$. In order to do this, we introduced a fast imaging pipeline that utilizes a per-observing session multistep and multiscale uv -plane based subtraction. We then demonstrate the applicability and success of uv -plane based gridding and averaging, which we perform per-observing epoch, to greatly reduce our total visibility data volume that is imaged. An image including all of the processed CHILES data then is passed through a gridded uv -plane RFI mitigation technique, which successfully removes large-scale artifacts from our data. The effect of this procedure is then quantified through discussion of kurtosis and noise values, which allows us to conclude that the CHILES data have been successfully processed through this fast imaging pipeline and are suitable for science.

Using the calibrated, imaged, and continuum subtracted CHILES data, as well as ancillary data from the COSMOS collaboration and CHILES Con Pol, we measure average H I mass, H I gas fraction, and star formation rate using stacking techniques. The results of these measurements are summarized as follows:

1. For the total ensemble of massive blue galaxies, which we define as galaxies with stellar masses in the range $10^{9-12.5} M_{\odot}$, and an NUV- r color in the range -1 to 3 , we find little-to-no evolution in the H I content for galaxies in the redshift range $0.09 < z < 0.47$.
2. For the massive red galaxies, which we define as galaxies with stellar masses in the range $10^{9-12.5} M_{\odot}$, and an NUV- r color in the range $3-6$, we make no H I mass

detection at any redshift, indicating that it is not the case that red galaxies are becoming significantly H I rich at increasing redshift such that they would be detected.

3. We separated our sample blue galaxies into two mass bins, an intermediate-mass bin, composed of galaxies with stellar masses in the range $10^9\text{--}10^9 M_\odot$, and a high-mass bin, composed of galaxies with stellar masses in the range $10^{10\text{--}12.5} M_\odot$. For the intermediate stellar mass bin, we find no significant evolution in the H I gas content. However, the high stellar mass bin shows significant increase in H I content for increasing redshifts.
4. We place our results into the context of the star-forming main sequence of galaxies and hypothesize that for the intermediate-mass galaxies, the decline in star formation rate is caused by an inefficient conversion of atomic to molecular hydrogen. For the high-mass galaxies, we find that H I is not being accreted onto the disks of these galaxies, which could be the cause of a decline in star formation rate, as well as an inefficient conversion of atomic to molecular hydrogen

The results presented here add to the recent findings of multiple groups using multiple different telescopes to understand the evolution of H I at intermediate redshifts using stacking techniques (e.g., A. Bera et al. 2019, 2023a; A. Chowdhury et al. 2020; A. Bianchetti et al. 2025). As new deep long-integration H I surveys are being undertaken (e.g., LADUMA; S. Blyth et al. 2016), future studies will be able to more properly fill in the redshift space of $z = 0.5\text{--}2$, and add to the results from A. Chowdhury et al. (2020), allowing for us to fully appreciate the role of H I in the evolution of galaxies. Additionally, next-generation telescopes will provide the spatial resolution necessary in order to understand the sundry physical processes that are governing the star formation rate evolution of the Universe.

Acknowledgments

We gratefully acknowledge the thorough review and helpful comments from the anonymous referee that added to the

clarity, demonstration of results, and scientific discussion, all of which served to greatly improved this paper.

We greatly appreciate valuable insights from Nick Scoville on an earlier draft of this work concerning the evolution of the total ISM. Additionally, we are grateful to Kumar Golap for helpful instructions on the usage of msuvbin, and to Bradley Frank for helpful conversations on the parallelization of tclean. Lastly, we thankfully acknowledge Guillermo Avendano Franco and Nate Garver-Daniels for their invaluable assistance with the use of the WVU HPCs.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. The CHILES survey was partially supported by a collaborative research grant from the National Science Foundation under grant Nos. AST-1412843, 1412578, 1413102, 1413099, and 1412503.

D.J.P. gratefully acknowledges support by the South African Research Chairs Initiative of the Department of Science and Technology and the National Research Foundation. J.B.B. was a Jansky Fellow of the National Radio Astronomy Observatory.

Facilities: EVLA (R. A. Perley et al. 2011).

Software: astropy (Astropy Collaboration et al. 2013), CASA (J. P. McMullin et al. 2007).

Appendix Stacked Measurements

In Figure A1, we present plots indicating the reliability for each of our stacked detections. Specifically, we show the stacked spectrum overlaid on 100 blank sky spectra, and a comparison of how the noise in our stacked spectra integrates down, compared to the expected $N^{-0.5}$, where N is the number of input galaxies for the stack. These figures demonstrate that, on average, features in the spectra are representative of global per-channel image artifacts, as opposed to purely serendipitous or spurious features, and that the noise for our stacks broadly agrees with expectations. The spectra here are organized by their order of appearance in the paper.

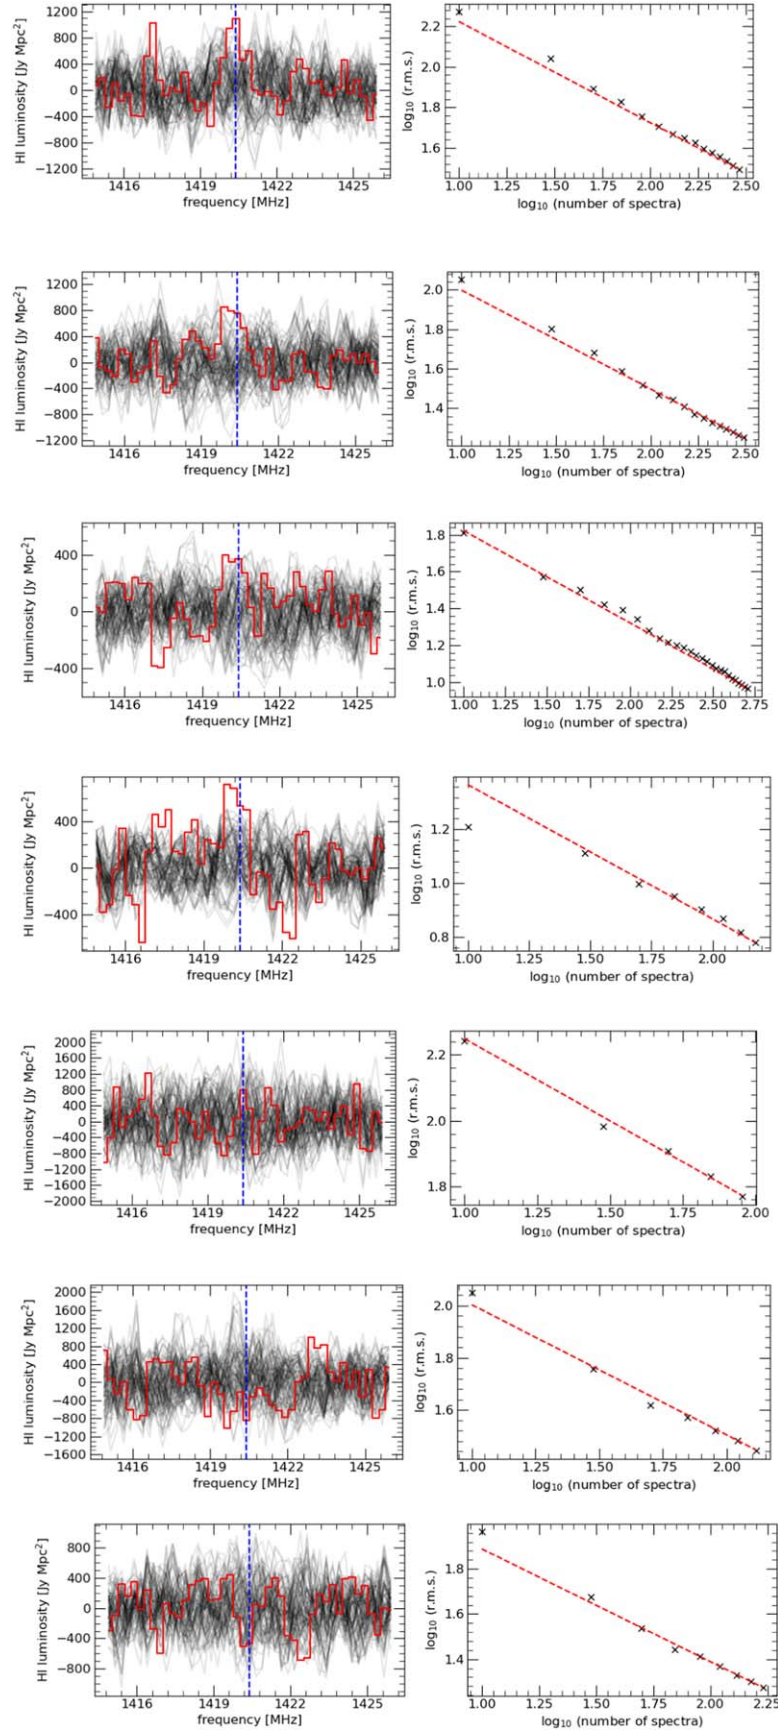


Figure A1. For each of the 14 stacked spectra summarized in Table 2, we present: (left) the stacked H I spectrum (red), 100 reference stacked spectra taken over randomized locations in the stacked cube (black), and the H I line rest frequency (blue dashed); and (right) the observed rms noise (black markers) and expected rms noise, assuming an $N^{-0.5}$ noise integration and the final stacked noise as the anchoring point (red dashed).

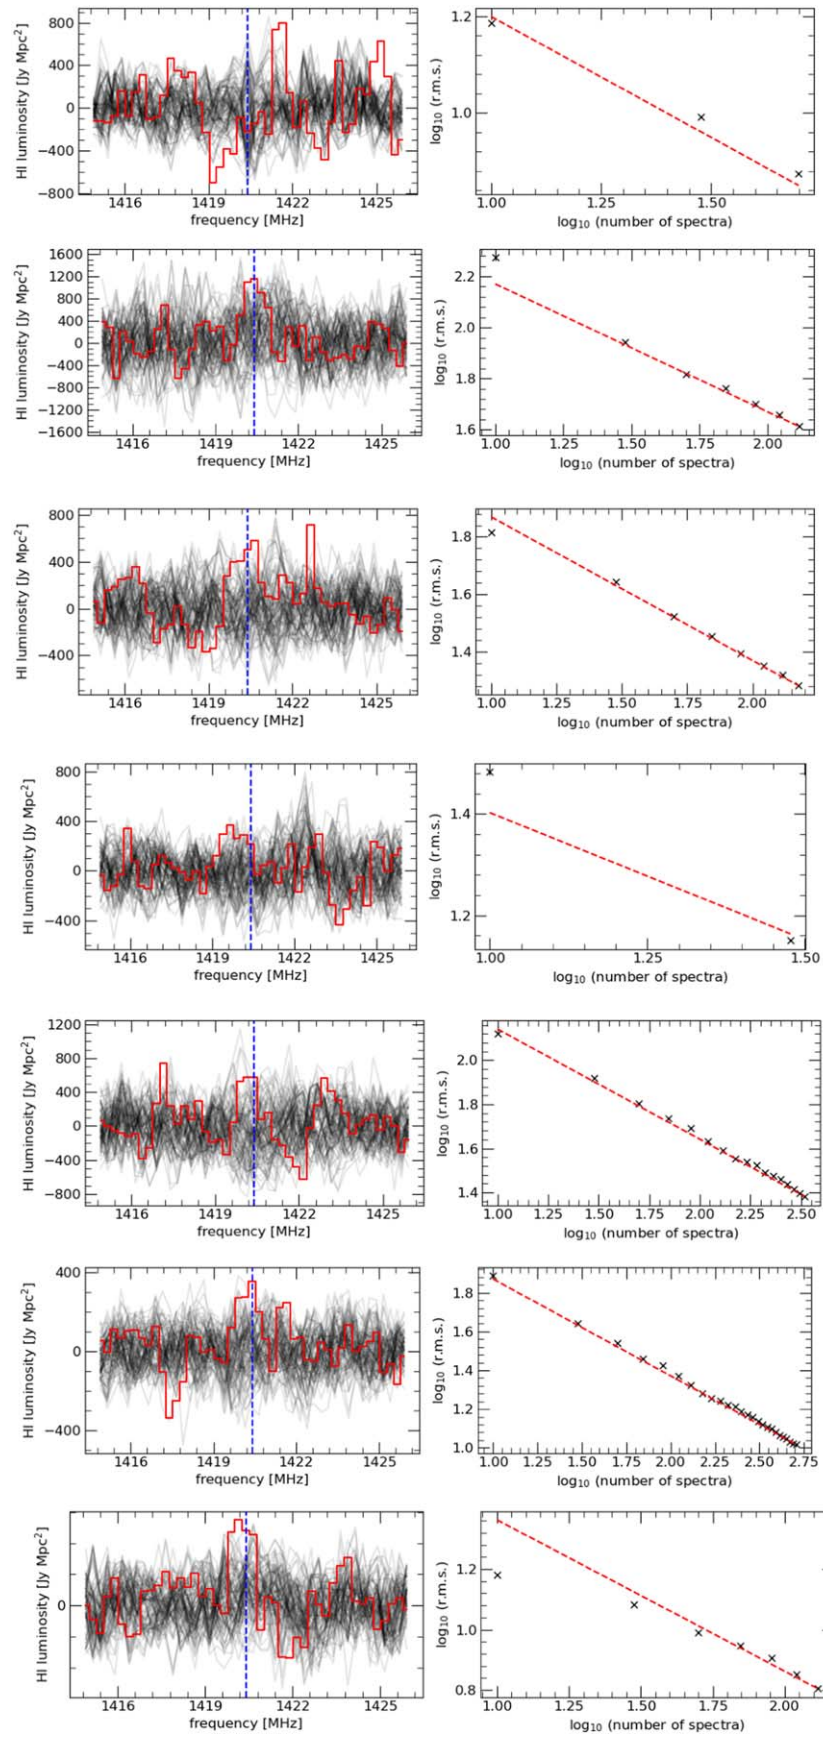











Figure A1. (Continued.)

ORCID iDs

D. J. Pisano  <https://orcid.org/0000-0001-7996-7860>
 J. H. van Gorkom  <https://orcid.org/0000-0002-7679-9344>
 Richard Dodson  <https://orcid.org/0000-0003-0392-3604>
 Hansung B. Gim  <https://orcid.org/0000-0003-1436-7658>
 Kelley M. Hess  <https://orcid.org/0000-0001-9662-9089>
 Lucas R. Hunt  <https://orcid.org/0000-0001-8587-9285>
 Danielle Lucero  <https://orcid.org/0000-0002-8799-3054>
 Emmanuel Momjian  <https://orcid.org/0000-0003-3168-5922>
 Min S. Yun  <https://orcid.org/0000-0001-7095-7543>

References

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Balanda, K. P., & MacGillivray, H. L. 1988, *Am. Stat.*, **42**, 9
- Bell, E. F., McIntosh, D. H., Katz, N., & Weinberg, M. D. 2003, *ApJS*, **149**, 289
- Bera, A., Kanekar, N., Chengalur, J. N., & Bagla, J. S. 2019, *ApJL*, **882**, L7
- Bera, A., Kanekar, N., Chengalur, J. N., & Bagla, J. S. 2023a, *ApJL*, **950**, L18
- Bera, A., Kanekar, N., Chengalur, J. N., & Bagla, J. S. 2023b, *ApJL*, **956**, L15
- Bianchetti, A., Sinigaglia, F., Rodighiero, G., et al. 2025, *ApJ*, **982**, 82
- Blue Bird, J., Davis, J., Luber, N., et al. 2020, *MNRAS*, **492**, 153
- Blyth, S., Baker, A. J., Holwerda, B., et al. 2016, in *Proc. MeerKAT Science: On the Pathway to the SKA (Trieste: PoS)*, **4**
- Briggs, D. S. 1995, PhD Thesis, New Mexico Institute of Mining and Technology
- Bruzual, G., & Charlot, S. 2003, *MNRAS*, **344**, 1000
- Catinella, B., & Cortese, L. 2015, *MNRAS*, **446**, 3526
- Catinella, B., Haynes, M. P., Giovanelli, R., Gardner, J. P., & Connolly, A. J. 2008, *ApJL*, **685**, L13
- Catinella, B., Saintonge, A., Janowiecki, S., et al. 2018, *MNRAS*, **476**, 875
- Chabrier, G. 2003, *PASP*, **115**, 763
- Chen, Q., Meyer, M., Popping, A., & Staveley-Smith, L. 2021a, *MNRAS*, **502**, 2308
- Chen, Q., Meyer, M., Popping, A., et al. 2021b, *MNRAS*, **508**, 2758
- Chowdhury, A., Kanekar, N., & Chengalur, J. N. 2022, *ApJL*, **941**, L6
- Chowdhury, A., Kanekar, N., & Chengalur, J. N. 2023, *ApJL*, **958**, L29
- Chowdhury, A., Kanekar, N., Chengalur, J. N., Sethi, S., & Dwarakanath, K. S. 2020, *Natur*, **586**, 369
- Chowdhury, A., Kanekar, N., Das, B., Dwarakanath, K. S., & Sethi, S. 2021, *ApJL*, **913**, L24
- Chung, A., van Gorkom, J. H., Kenney, J. D. P., Crowl, H., & Vollmer, B. 2009, *AJ*, **138**, 1741
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2005, in *ASP Conf. Ser. 347, Astronomical Data Analysis Software and Systems XIV*, ed. P. Shopbell, M. Britton, & R. Ebert (San Francisco, CA: ASP), **86**
- Cornwell, T. J., Uson, J. M., & Haddad, N. 1992, *A&A*, **258**, 583
- Cortese, L., Catinella, B., Boissier, S., Boselli, A., & Heinis, S. 2011, *MNRAS*, **415**, 1797
- Davidzon, I., Ilbert, O., Laigle, C., et al. 2017, *A&A*, **605**, A70
- Davies, L. J. M., Driver, S. P., Robotham, A. S. G., et al. 2015, *MNRAS*, **447**, 1014
- Delhaize, J., Meyer, M. J., Staveley-Smith, L., & Boyle, B. J. 2013, *MNRAS*, **433**, 1398
- Dodson, R., Momjian, E., Pisano, D. J., et al. 2022, *AJ*, **163**, 59
- Dodson, R., Vinsen, K., Wu, C., et al. 2016, *A&C*, **14**, 8
- Driver, S. P., Hill, D. T., Kelvin, L. S., et al. 2011, *MNRAS*, **413**, 971
- Fernández, X., Gim, H. B., van Gorkom, J. H., et al. 2016, *ApJL*, **824**, L1
- Gogate, A. R., Verheijen, M. A. W., Deshev, B. Z., et al. 2020, *MNRAS*, **496**, 3531
- Gim, H. B., Yun, M. S., & Luber, N. M. 2025, arXiv:2504.20200
- Golap, K., & Momjian, E. 2016, VLA Expansion Project Memoranda #198
- Greisen, E. W. 2003, in *Astrophysics and Space Science Library*, Vol. 285, Information Handling in Astronomy - Historical Vistas, ed. A. Heck (Dordrecht: Kluwer), **109**
- Gunn, J. E., & Gott, J. R. I. 1972, *ApJ*, **176**, 1
- Hales, C. A., & Stephenson, N. 2019, VLA Expansion Project Memoranda
- Haynes, M. P., & Giovanelli, R. 1986, *ApJ*, **306**, 466
- Hess, K. M., Luber, N. M., Fernández, X., et al. 2019, *MNRAS*, **484**, 2234
- Hess, K. M., & Wilcots, E. M. 2013, *AJ*, **146**, 124
- Heywood, I., Ponomareva, A. A., Maddox, N., et al. 2024, *MNRAS*, **534**, 76
- Högbom, J. A. 1974, *A&AS*, **15**, 417
- Hopkins, A. M., & Beacom, J. F. 2006, *ApJ*, **651**, 142
- Huchtmeier, W. K. 1994, *A&A*, **286**, 389
- Jagannathan, P., Bhatnagar, S., Briske, W., & Taylor, A. R. 2018, *AJ*, **155**, 3
- Jones, M. G., Haynes, M. P., Giovanelli, R., & Papastergis, E. 2016, *MNRAS*, **455**, 1574
- Kenney, J. D. P., van Gorkom, J. H., & Vollmer, B. 2004, *AJ*, **127**, 3361
- Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009, *MNRAS*, **395**, 160
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *MNRAS*, **363**, 2
- Kilborn, V. A., Forbes, D. A., Barnes, D. G., et al. 2009, *MNRAS*, **400**, 1962
- Kreckel, K., Platen, E., Aragón-Calvo, M. A., et al. 2012, *AJ*, **144**, 16
- Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, **224**, 24
- Luber, N., van Gorkom, J. H., Hess, K. M., et al. 2019, *AJ*, **157**, 254
- Luber, N. M., Hasan, F., van Gorkom, J. H., et al. 2025, *ApJ*, **985**, 214
- Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
- Maddox, N., Hess, K. M., Blyth, S. L., & Jarvis, M. J. 2013, *MNRAS*, **433**, 2613
- Maddox, N., Frank, B. S., Ponomareva, A. A., et al. 2021, *A&A*, **646**, A35
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in *ASP Conf. Ser. 376, CASA Architecture and Applications*, ed. R. A. Shaw, F. Hill, & D. J. Bell (San Francisco, CA: ASP), **127**
- Morganti, R., de Zeeuw, P. T., Oosterloo, T. A., et al. 2006, *MNRAS*, **371**, 157
- Murphy, E. J., Condon, J. J., Schinnerer, E., et al. 2011, *ApJ*, **737**, 67
- Pearson, S., Besla, G., Putman, M. E., et al. 2016, *MNRAS*, **459**, 1827
- Perley, R. A., Chandler, C. J., Butler, B. J., & Wrobel, J. M. 2011, *ApJL*, **739**, L1
- Poggianti, B. M., Moretti, A., Gullieuszik, M., et al. 2017, *ApJ*, **844**, 48
- Ramatsoku, M., Serra, P., Poggianti, B. M., et al. 2020, *A&A*, **640**, A22
- Renzini, A., & Peng, Y.-j. 2015, *ApJL*, **801**, L29
- Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, **172**, 1
- Scoville, N., Lee, N., Vanden Bout, P., et al. 2017, *ApJ*, **837**, 150
- Sekhar, S., & Athreya, R. 2018, *AJ*, **156**, 9
- Sinigaglia, F., Rodighiero, G., Elson, E., et al. 2022, *ApJL*, **935**, L13
- Solanes, J. M., Manrique, A., García-Gómez, C., et al. 2001, *ApJ*, **548**, 97
- Sousbie, T. 2011, *MNRAS*, **414**, 350
- Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, **214**, 15
- Swaters, R. A., Sancisi, R., van Albada, T. S., & van der Hulst, J. M. 2009, *A&A*, **493**, 871
- Swaters, R. A., van Albada, T. S., van der Hulst, J. M., & Sancisi, R. 2002, *A&A*, **390**, 829
- Verheijen, M., van Gorkom, J. H., Szomoru, A., et al. 2007, *ApJL*, **668**, L9
- White, R. L., Helfand, D. J., Becker, R. H., Glikman, E., & de Vries, W. 2007, *ApJ*, **654**, 99
- Wu, C., Tobar, R., Vinsen, K., et al. 2017, *A&C*, **20**, 1
- Yun, M. S., & Carilli, C. L. 2002, *ApJ*, **568**, 88
- Yun, M. S., Reddy, N. A., & Condon, J. J. 2001, *ApJ*, **554**, 803
- Zwaan, M. A., van Dokkum, P. G., & Verheijen, M. A. W. 2001, *Sci*, **293**, 1800