



HAL
open science

hiss, a new tool for H i stacking: application to NIBLES spectra

J Healy, S-L Blyth, E Elson, W. Van driel, Z Butcher, S Schneider, M Lehnert, R Minchin

► **To cite this version:**

J Healy, S-L Blyth, E Elson, W. Van driel, Z Butcher, et al.. *hiss*, a new tool for H i stacking: application to NIBLES spectra. *Monthly Notices of the Royal Astronomical Society*, 2019, 487 (4), pp.4901-4938. 10.1093/mnras/stz1555 . hal-03103887

HAL Id: hal-03103887

<https://hal.science/hal-03103887>

Submitted on 8 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HISS, a new tool for H I stacking: application to NIBLES spectra

J. Healy^{1b},^{1,2*} S.-L. Blyth^{1b},¹ E. Elson^{1b},^{1,3} W. van Driel,^{4,5} Z. Butcher,⁶
S. Schneider,⁶ M. D. Lehnert⁷ and R. Minchin^{1b}^{8,9}

¹Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

²Kapteyn Astronomical Institute, University of Groningen, Landlevan 12, NL-9747 AV Groningen, the Netherlands

³Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, Bellville 7535, South Africa

⁴GEPI, Observatoire de Paris, PSL Research University, CNRS, 5 place Jules Janssen, F-92190 Meudon, France

⁵Station de Radioastronomie de Nançay, Observatoire de Paris, CNRS/INSU USR 704, Université d'Orléans OSUC, route de Souesmes, F-18330 Nançay, France

⁶Astronomy Program, University of Massachusetts, 619E LGRT-B, Amherst, MA 01003, USA

⁷CNRS UMR 7095, Institut d'Astrophysique de Paris, Sorbonne Université, 98bis bd Arago, F-75014 Paris, France

⁸Arecibo Observatory, National Astronomy and Ionosphere Center, Arecibo, PR 00612, USA

⁹SOFIA-USRA, NASA Ames Research Center, MS 232-12, Moffett Field, CA 94035, USA

Accepted 2019 June 4. Received 2019 June 4; in original form 2017 December 5

ABSTRACT

H I stacking has proven to be a highly effective tool to statistically analyse average H I properties for samples of galaxies which may or may not be directly detected. With the plethora of H I data expected from the various upcoming H I surveys with the SKA Precursor and Pathfinder telescopes, it will be helpful to standardize the way in which stacking analyses are conducted. In this work we present a new PYTHON-based package, HISS, designed to stack H I (emission and absorption) spectra in a consistent and reliable manner. As an example, we use HISS to study the H I content in various galaxy sub-samples from the NIBLES survey of SDSS galaxies which were selected to represent their entire range in total stellar mass without a prior colour selection. This allowed us to compare the galaxy colour to average H I content in both detected and non-detected galaxies. Our sample, with a stellar mass range of $10^8 < M_{\star} (M_{\odot}) < 10^{12}$, has enabled us to probe the H I-to-stellar mass gas fraction relationship more than half an order of magnitude lower than in previous stacking studies.

Key words: methods: data analysis – methods: statistical – galaxies: evolution – galaxies: general – radio lines: galaxies.

1 INTRODUCTION

How galaxies evolve over cosmic time is currently a key area of research in astrophysics. A great deal is already known about the evolution of the stellar content of galaxies due to recent infrared (*Spitzer*, Werner et al. 2004), optical (Sloan Digital Sky Survey, York et al. 2000), and ultraviolet (Galaxy Evolution Explorer, Bianchi & GALEX Team 2000) surveys. However, comparatively little is known about the evolution of the gas in galaxies. Understanding how the cold gas evolves is important as it is the raw fuel for the formation of stars and thus galaxies.

Neutral atomic hydrogen (H I) forms the most significant reservoir of neutral gas in galaxies. Studies have shown that blue, star-forming galaxies have a higher fraction of H I gas compared to red, quiescent galaxies (e.g. Roberts & Haynes 1994; Gavazzi, Pierini & Boselli 1996; McGaugh & de Blok 1997; Cortese et al. 2011;

Fabello et al. 2011; Brown et al. 2015), which suggests that H I plays an important role in star formation. H I is difficult to detect in most galaxies beyond the local Universe ($z \sim 0.06$) with existing radio telescopes due to the weak nature of the emission. Researchers have had to exploit different techniques in order to measure very weak H I signals from galaxies. One such technique is H I stacking; with this technique, an average H I mass per galaxy can be estimated by co-adding the H I spectra of a sample of galaxies in which H I is not necessarily directly detected.

The idea of co-adding the undetected H I spectra in studies of the gas content in galaxies was first presented by Zwaan et al. (2001) and Chengalur, Braun & Wieringa (2001). Both groups were studying the H I within galaxies located in and around clusters. With low detection counts in their samples, both groups independently co-added the non-detections in an effort to obtain a statistically meaningful averaged detection for their samples.

Stacking analyses have become common place in the last 15+ yr. The technique has been applied in various areas: H I content of galaxies in dense environments and how the gas content relates

* E-mail: juliahealyza@gmail.com

to other observables (Chengalur et al. 2001; Verheijen et al. 2007; Lah et al. 2009; Jaffé et al. 2016); gas content of active galaxies (Geréb et al. 2013, 2015); measurement of Ω at low to intermediate ($z < 0.4$) redshifts (Lah et al. 2007; Delhaize et al. 2013; Rhee et al. 2013, 2016); and using stacking to study the relations between H I and various stellar mass/star formation indicators (Fabello et al. 2011, 2012; Brown et al. 2015, 2017; Geréb et al. 2015).

The gas scaling relations for galaxies are average trends which relate the H I mass (M_{HI}) or H I mass to stellar mass (M_*) ratio (gas fraction, f_{HI}) to various other galaxy properties. Stacking a sample of ~ 5000 galaxies with $M_* > 10^{10} M_\odot$ that had both ultraviolet and optical imaging, Fabello et al. (2011) found using H I stacking that the gas fraction correlates better with $\text{NUV} - r$ colour than stellar mass. This result was later confirmed and extended to a sample with $M_* > 10^9 M_\odot$ by Brown et al. (2015). To date, studies of the H I gas scaling relations have been limited to samples with $M_* > 10^9 M_\odot$, a mass which has been identified as a turning point around which the M_{HI} versus M_* slope changes (Huang et al. 2012; Maddox et al. 2015).

Directly detecting H I with current radio telescopes beyond $z \sim 0.1$ is challenging due to the long observing times required [e.g. HIGHz (Catinella & Cortese 2015) on Arecibo using > 300 h and CHILES (Fernández et al. 2016; Hess et al. 2018) on the JVLA requiring ~ 1000 h of observing time]. The upcoming surveys with the SKA Precursors and Pathfinders (e.g. APERTIF, Verheijen et al. 2008; ASKAP, Johnston et al. 2008; MeerKAT, Booth et al. 2009) will greatly extend the redshift range over which H I in galaxies is studied, either directly or indirectly. Techniques such as H I stacking will have an important role to play in order to study the average H I properties of different galaxy samples, particularly at higher redshifts ($z \gtrsim 0.6$). Deep SKA Precursor surveys such as LADUMA¹ (Holwerda, Blyth & Baker 2011) on MeerKAT have identified stacking as an important tool to probe higher redshifts, and push to lower H I mass limits.

As we rapidly approach the start of the Precursor Surveys, work is underway on developing a data analysis toolkit that will enable consistent and comparable studies of the survey data. Already available is the versatile source finder, SOFIA (Serra et al. 2015). With the important role that stacking will play in the analysis of high-redshift and low-mass samples, it is imperative that a tool capable of reliably and consistently stacking H I spectra is developed.

In this work, we present our new software package, HISS, that has been designed for the astronomy community in response to the need for a stacking software package. HISS can be downloaded from <https://github.com/healytwin1/HISS>. We use HISS to revisit the gas scaling relations with a sample of 1000 galaxies from the Nançay Interstellar Baryon Legacy Extragalactic Survey (NIBLES, van Driel et al. 2016). NIBLES is an SDSS-selected targeted H I survey with the Nançay Radio Telescope (NRT) which aims to study the H I properties of galaxies as a function stellar mass, covering a representable M_* range of the ensemble of galaxies in the nearby Universe.

The outline for this paper is as follows: the first half of the paper (Section 2) describes the design of the H I Stacking Software (HISS). In the second half of the paper (Sections 3 and 4), we give an introduction to the NIBLES Survey and the ancillary data we use in our stacking analysis and describe the classification of H I non-detections. We use the NIBLES sample to explore the well-known

H I mass to stellar mass scaling relations in Section 4, and finally, summarize our findings in Section 5.

2 THE H I STACKING SOFTWARE (HISS)

The stacking method that the software needs to implement can be summarized by the following steps:

- (i) Ingest 1D 21 cm H I spectra (radio data) and list of associated redshifts for the sample of galaxies to be stacked;
- (ii) Spectrally shift and re-scale the H I spectra to align the expected line emission at a common frequency (usually the H I rest frequency: 1420.4058 MHz);
- (iii) Weight the spectra (according to a preferred weighting scheme) and co-add the spectra.

We have designed and created the H I Stacking Software (HISS) package, after discussion with colleagues and with the following guiding principles in mind:

- (i) be freely available in most operating systems
- (ii) to be open source
- (iii) easy to modify (extensible)
- (iv) stack hundreds or thousands of galaxy spectra in an efficient and reliable manner.

The PYTHON² programming language was chosen for development of HISS because it is freely available, can be used on any operating system, and is also one of the most commonly used languages within the astronomy community. Where appropriate, HISS makes use of the publicly available PYTHON modules (such as ASTROPY, NUMPY, SCIPY, etc.) which have been optimized for data input, manipulation, and display.

Fig. 1 shows a flow diagram of how the processes required to create a stacked spectrum are incorporated in the different modules of HISS. In the sections that follow we will illustrate how each of the processes highlighted in Fig. 1 are implemented. Basic instructions on how to run HISS can be found in Appendix A1.

2.1 Simulated data

When developing new data analysis tools, one of the most important steps is to quantify the accuracy and reliability of any generated results. For this purpose, simulated spectra are used instead of real spectra because properties of the input data are then well known. In this work, we use a data set of 1000 simulated H I profiles of galaxies to illustrate and test the capabilities of HISS. The simulated H I spectra are created using the formulation outlined in Obreschkow et al. (2009c, Appendix A), and are based on the evaluated properties extracted from the S³-SAX catalogue (Obreschkow, Rawlings & Road 2009a; Obreschkow et al. 2009b,c) for galaxies in a redshift range of $0.1 < z < 0.2$ with a channel width of 7 km s^{-1} and Gaussian noise of $16 \mu\text{Jy}$ per channel.

2.2 Catalogue of H I spectra

When initiated, HISS requires some information about the sample of H I spectra – see Fig. 2 for an overview. The input module requires two input files: a configuration file and a catalogue file. There are two ways to provide HISS with the required information: the user

¹Looking At the Distant Universe with the MeerKAT Array.

²Developed in PYTHON 2.7, but compatible with PYTHON 3.

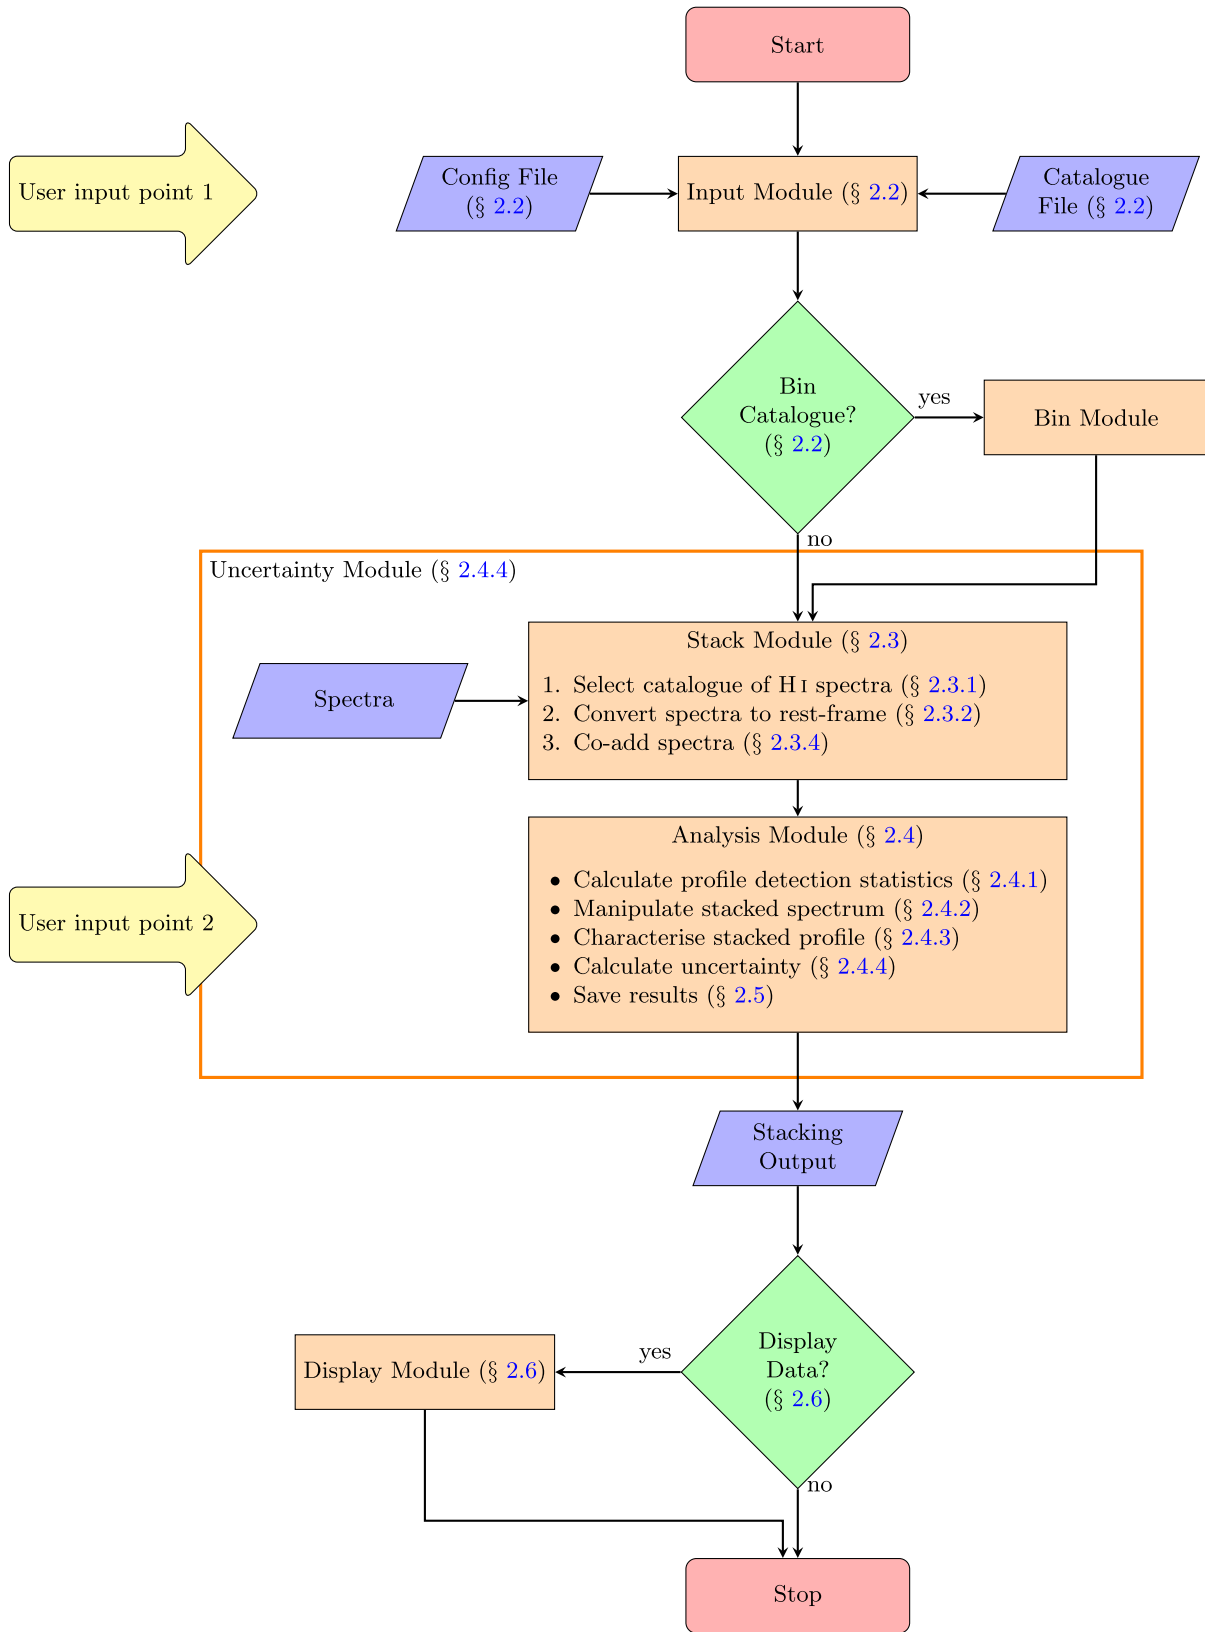


Figure 1. This flow diagram shows how the individual spectra and user information are taken by HISS to produce a stacked spectrum from which average galaxy properties (such as total HI mass, M_{HI} , and the HI-to-stellar mass ratio, f_{HI}) may be extracted. The orange rectangles show the six modules of the package, the blue parallelograms show the points of input or output, and the green diamonds show where the user may choose to incorporate the optional modules.

Figure 2. Screen shot from GUI main window, useful for first-time HISS users. It shows an overview of all initial parameters required by HISS. This graphical interface fills in a configuration file whose parameters are explained in Section A2.

can provide a text file in JSON³ format or use the graphical user interface shown in Fig. 2 to generate a configuration file (also in JSON format). The JSON format is used for the configuration file as it is easy to read/write, and when imported into PYTHON, the information is dumped into an easy-to-access PYTHON dictionary. The configuration file includes options such as different weighting

functions for the stacking procedure, and options to bin the stacking sample based on additional data provided in the catalogue file.

The catalogue file is a user-created text file in CSV⁴ format that contains at least the following columns:

- (i) Object ID
- (ii) Spectrum filename
- (iii) Redshift

³JSON stands for JavaScript Object Notation, a JSON formatted file has the extension `.json`

⁴Comma separated values.

- (iv) Redshift uncertainty
- (v) Other data (optional)

for each spectrum that the user would like to include in the stack. Although the catalogue file may contain any number of columns, the non-optional columns mentioned above are those that are required to run HISS. Additional columns containing numerical information such as stellar mass or optical colour may be used to refine the catalogue into a number of different sub-samples to be stacked separately. This refining is done through the Bin Module (Fig. 1). In this first version of HISS, bins can only be created using one quantity, e.g. stellar mass or colour, and the stacked results are stored for each sub-sample.⁵

During the processing of the configuration file, the code checks that the catalogue file and location of the spectra exist. If these checks are passed, HISS will continue.

2.3 Stacking the spectra

The stacking procedure is the heart of HISS, and is controlled by the Stack Module which is responsible for reading in and preparing each spectrum for stacking, as well as maintaining the stacked spectrum and associated information (e.g. number of objects in the stack, stacked noise). One of the features of this module is the option that allows the user to watch the progress of the stacking in a window such as the one in Fig. A1. The progress window shows each of the spectra in the observed frame as they are read in, as well as the individual spectra as they are converted to the rest frame and then added to the total stacked spectrum.

2.3.1 Reading in the spectra

Regardless of how the HI spectra were created, HISS requires that the spectra be in a text file type format containing a column for the spectral axis (either frequency or velocity) and another for the flux density (in units of μJy , mJy, Jy, or flux density/beam).

Fig. A2 shows a selection of ten of the 1000 simulated HI profiles that will be used to illustrate the different parts of HISS. Each of these have had 16 μJy Gaussian noise added, so as to simulate the spectra that are anticipated from the deep LADUMA survey.

As each spectrum is read into HISS, it is put through a series of quality checks:

- (i) The spectrum's channel width is checked to ensure it is within 5 per cent of the entered channel width. A warning is raised if the spectrum channel width is 5–10 per cent different. If the channel width is more than 10 per cent different the spectrum is discarded.
- (ii) The input object contains the minimum and maximum redshift values for the sample of galaxies. This check makes sure that the redshift associated with the spectrum falls within these bounds.
- (iii) If the user has decided to stack the spectra in units of gas fraction, the catalogue will be checked for an associated stellar mass value.
- (iv) The final check is for spectrum length. A galaxy mask is created based on the expected maximum width galaxy width for the sample. Every spectrum is checked to ensure that it has more channels than the length of mask. An example of a spectrum that

would fail such a check is shown in the panel (j) of Figs A2 and A3.

Due to the discrete nature of the input spectra, the channel width check is important to avoid co-adding arrays where the rest-frame size of the channel (i.e. velocity width in km s^{-1}) changes significantly (> 10 per cent) between different arrays. The spectra that do not pass at least one of the quality checks are excluded from the stack. The object IDs for the excluded spectra and the reason for exclusion are recorded in the log file.

2.3.2 Convert spectra to rest frame

If the spectrum passes all the quality checks, it moves to the next step which is to align the spectrum such that the galaxy signal is centred about the central channel of the spectrum. There are two steps to aligning the spectra. The first step is to convert the observed frequencies to a set of rest-frame frequencies corresponding to the systemic velocity of the galaxy being placed at $z = 0$. The spectral axes are converted to the rest frame by

$$v_{\text{emit}} = v_{\text{obs}} - cz \quad \text{OR} \quad v_{\text{emit}} = v_{\text{obs}}(1 + z), \quad (1)$$

where v_{obs} and v_{emit} are the spectrum's spectral axis in either velocity or frequency, respectively. The galaxy's redshift is z , and c is the speed of light given in kilometres per second.

An input spectrum with a velocity spectral axis is assumed to have rest-frame channel widths; this means that the conversion from the observed frame to the rest frame only requires a shift from the recessional velocity to 0 km s^{-1} as given by the velocity relation in equation (1). For a spectrum with a frequency spectral axis, the rest frame channel width is larger than the observed frame by a factor of $(1 + z)$; thus in order to convert the spectrum to the rest frame, both the spectral axis and flux density must be appropriately scaled such that the integrated line flux density is conserved. The frequency axis is converted to the rest frame using the frequency relation in equation (1), and the flux is scaled according to: $S_{v,\text{emit}} = S_{v,\text{obs}}/(1 + z)$.

2.3.3 Spectrum wrapping

The second step is to shift the galaxy emission to the centre of the spectrum. Since the spectra are now in the rest frame, the target emission is located at either 0 km s^{-1} or 1420 MHz (depending on the units of the spectral axis) and is easily shifted to the centre of the spectrum. In this step, some authors (e.g. Zwaan et al. 2001) will concatenate the spectra such that every channel has the same number of measurements (i.e. all of the spectra are the same length). Other authors (e.g. Lah et al. 2007, 2009; Rhee et al. 2013) do not concatenate the spectra, instead choosing to weight each channel in the stacked spectrum relative to the number of measurements in that channel.

Instead, HISS wraps the flux that would be shifted out of the now centred spectrum to append to the other end of the spectrum so that the original spectrum length is maintained; this method also maintains the spectrum noise properties. This is illustrated in Fig. A2: the green part of the spectrum is the main part that is shifted to the centre and purple highlights the section of the spectrum that is appended to the other side of the spectrum array. The centred spectra are shown in blue in Fig. A3. It is clear from each of the panels of input spectra in Fig. A2 that the spectra need not all be the same length. Each spectrum is thus converted to some appropriate

⁵In the current implementation this process is sequential i.e. each binned sub-sample of spectra is stacked one after the other. This process will be parallelized in future versions.

length specified by the user as an input parameter, (see bottom block of Fig. 2 – *Velocity width of spectrum*⁶).

In order to avoid issues with small number statistics in the spectrum noise calculations, the length of the stacked spectrum should be at least three times the width of the broadest galaxy profile (i.e. the largest W_{50}). The spectra that are longer than necessary are simply truncated at the edges while keeping the galaxy emission at the centre of the spectrum. The shorter spectra are extended by using the so-called ‘noisy’ channels from the outer channels that are not expected to contain any galaxy emission to fill up the new empty channels at either edge of the spectrum.

In order to determine which channels contain only noise, the galaxy emission is masked. The galaxy mask is centred at the spectral location of the galaxy and has some pre-determined width (this is usually a conservative estimate on the maximum possible velocity width of galaxies in the sample). In Fig. A2, for illustrative purposes, the galaxy emission is masked using the pale blue band, all other channels are considered to contain only noise, and so can be used in the spectrum extension process. There are some caveats associated with this process: if there are too few noisy channels in the spectrum, then a ringing phenomenon may appear in the extended spectrum. Fig. A3 shows the original spectra centred in blue, while black indicates the noise that has been used to fill up the channels of the extended spectrum. Fig. A3(j) highlights the ringing phenomenon.

2.3.4 Co-adding the spectra

Before adding each rest-frame spectrum to the stack, the spectra are converted to the units (Jy , M_{\odot} , f_{HI}) that the user specified as the ‘stacking units’ during the configuration process.

To convert the HI flux density to HI mass, we use the following relation from Wieringa, de Bruyn & Katgert (1992):

$$\frac{M_{\text{HI}}}{M_{\odot}} = \frac{2.36 \times 10^5}{1+z} \left(\frac{D_L(z)}{\text{Mpc}} \right)^2 \left(\frac{\int S_v dv}{\text{Jy km s}^{-1}} \right), \quad (2)$$

where z is the galaxy redshift and $D_L(z)$ the associated luminosity distance. $\int S_v dv$ is the galaxy rest frame integrated line flux density for the profile in Jy km s^{-1} . This relation between the flux density and the HI mass assumes a spherical HI cloud that is optically thin with a uniform internal velocity distribution.

The HI gas fraction (or more accurately, the HI to stellar mass ratio) is simply

$$f_{\text{HI}} = \frac{M_{\text{HI}}}{M_{\star}}, \quad (3)$$

where M_{\star} is the total stellar mass and M_{HI} is the total HI mass. The right y-axes of the spectra in Fig. A3 are in units of M_{\odot} per channel; this choice of units allows us to simply add the values in each channel together when measuring the integrated value.

The final step required to produce a stacked spectrum is to co-add the spectra. This is typically done as a weighted average. The benefit of using the average is that Gaussian noise decreases proportional to $1/\sqrt{N}$ (N is the number of profiles included in the stacked spectrum), which is favourable to the creation of a spectrum with a higher signal-to-noise ratio, S/N. Fig. 3 shows how the average noise decreases as the number of profiles (in our simulated




Figure 3. This plot shows how the S/N of the average spectrum increases with increasing number of profiles. The average noise decreases as $1/\sqrt{N}$.

sample) included in the stack increase. The weighted average is defined by

$$S_{\text{stack}} = \frac{\sum_{i=1}^N w_i S_i}{\sum_{i=1}^N w_i}, \quad (4)$$

where i is the number of the spectrum to be included in the stack and N is the total number of spectra in the sample. Each spectrum S_i has an associated weighting factor w_i . By default, the weighting factor is 1, but the user may choose an alternate option such as

- (i) $w_i = \sigma_{\text{rms},i}^{-2}$ (Fabello et al. 2011)
- (ii) $w_i = \sigma_{\text{rms},i}^{-1}$ (Lah et al. 2007)
- (iii) $w_i = (\sigma_{\text{rms},i} D_{L(z),i}^2)^{-2}$ (Delhaize et al. 2013)

where $\sigma_{\text{rms},i}$ is the RMS of the channels in a spectrum that contains no galaxy emission (i.e. the noisy channels), and $D_{L(z),i}$ is the luminosity distance to the galaxy. We use the interquartile range of each spectrum as the noise estimate over the conventional standard deviation in order to deal with outlying high or low values for example from RFI spikes.

2.3.5 Reference/control spectrum

The practice of creating a reference spectrum has been used since the first HI stacking experiments. A reference spectrum is a stacked spectrum created by stacking the spectra in the target sample using the incorrect redshift to align the spectra. The various stacking studies have used two methods to create the reference spectrum. Some authors, who have had access to the radio data cube from which the HI spectra were extracted, choose to create the reference spectrum by extracting spectra from random spatial locations in the data cube and using the target redshifts to align the spectra (e.g. Zwaan et al. 2001). Other authors use the spectra from their sample but randomize the sample redshifts (e.g. Chengalur et al. 2001; Delhaize et al. 2013). By randomizing the sample redshifts, both the reference spectrum and the stacked spectrum have the same systematics arising from the shifting process (Chengalur et al. 2001). If the redshifts are randomly assigned to the sample spectra, the resulting average spectrum should take the appearance of a noise spectrum. If this is the case, it provides further confirmation that a detection in a stacked spectrum is legitimate.

The left-hand panels of Fig. 4 show how when the spectra are converted to the rest frame with the wrong redshift they end up

⁶The default length is 2000 km s^{-1} which is $4 \times$ width of a galaxy rotating at 500 km s^{-1} . A wider spectral axis is not advised due to possible noise variations across the bandwidth.

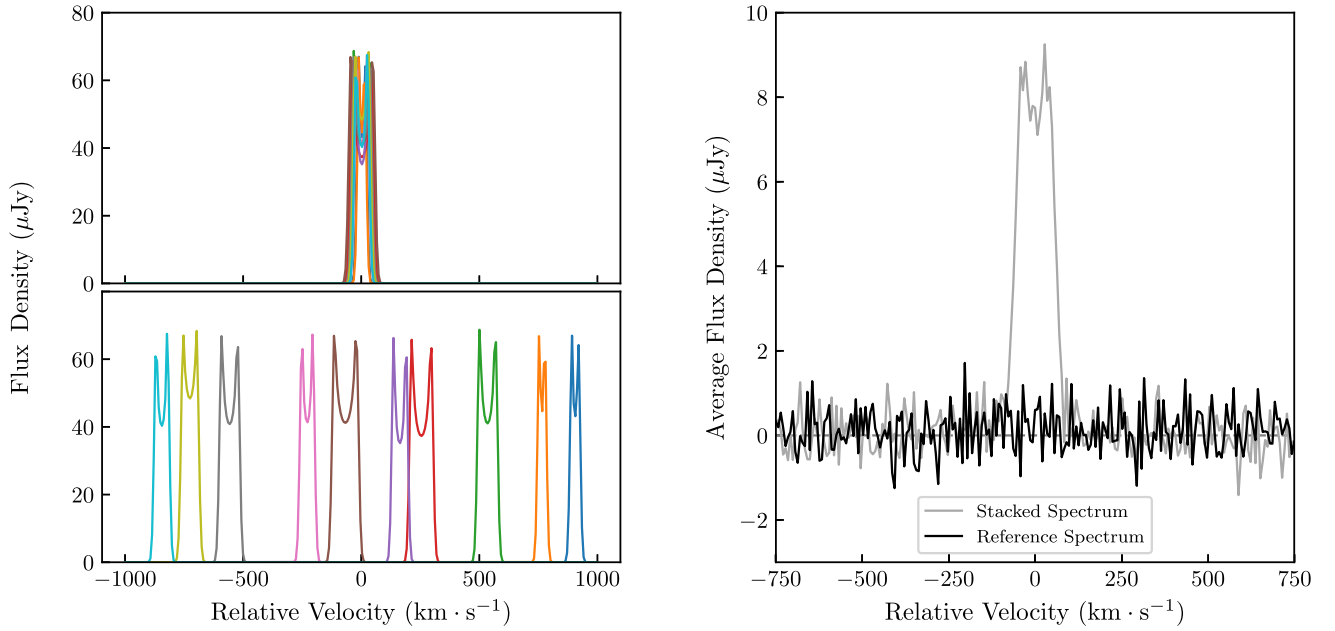


Figure 4. *Top left:* the noise-free versions of the same 10 simulated spectra that have been used previously. These spectra have been aligned using the correct redshifts. *Bottom left:* each of the 10 spectra has been shifted using the redshift from another galaxy. *Right:* the reference spectrum for the 1000 noisy simulated spectra is plotted in black. The stacked spectrum for this sample is plotted in grey for comparison.

completely misaligned. Each of the spectra in the bottom left panel of Fig. 4 has been shifted using the redshift belonging to one of the other galaxies. The result of co-adding the misaligned spectra using equation (4) is shown by the black line in the right-hand panel of Fig. 4. The lack of any significant emission at 0 km s^{-1} supports the integrity of the detection of the corresponding stacked spectrum that has been created using the correct redshifts. For comparison, the thin grey line shows the stacked spectrum of the correctly aligned spectra.

2.4 Analysing the stacked spectrum

The Analysis Module provides the user with a collection of statistics, spectrum manipulation, and profile characterization routines without the need for third-party software. Upon initialization of the Analysis Module, an initial statistical analysis is performed and the results are displayed either to a window or to the command line if running in suppress mode. At this point the first visual of the stacked spectrum is displayed in the form of Diagnostic Plot 1 (see Fig. A4). The user is also asked if the integration window size should be changed from its setting from the input phase.⁷ The next steps are then displayed to the user:

- (i) Smooth the spectrum (Section 2.4.2).
- (ii) Fit a number of other functions to the spectrum (Section 2.4.3).
- (iii) Rebin the spectrum (Section 2.4.2).
- (iv) Do not fit anything to the spectrum, but continue with the analysis.
- (v) Exit HISS without doing anything else. (This option will not save data.)

⁷The original integration window is specified by the user when setting *Max galaxy velocity width of sample* (see bottom block of Fig. 2).

2.4.1 Profile detection statistics

The collection of statistics that are calculated by this module all characterize the significance of the possible detections in the stacked spectrum in some way. The statistics are displayed in the terminal after calculation, and are saved in the log file for perusal after HISS has completed. The statistical analysis of the stacked spectrum includes three different methods of calculating the signal-to-noise ratio (S/N):

- (i) peak flux density (S_{peak}) to noise (σ_{rms}):

$$S/N = S_{\text{peak}}/\sigma_{\text{rms}}, \quad (5)$$

where S_{peak} is the maximum value of the spectrum within the integration window.

- (ii) integrated signal-to-noise:

$$S/N_{\text{int}} = \frac{\sum_i^{N_{\text{chan}}} S_i \, dv}{\sigma_{\text{rms}} \, dv \sqrt{N_{\text{chan}}}}, \quad (6)$$

where S_i is the flux density (in Jy) associated with channel i , σ is the rms noise, dv is the channel width (in km s^{-1}), and N_{chan} is the number of channels integrated over.

- (iii) ALFALFA S/N (from Haynes et al. 2011):

$$S/N_{\text{ALFALFA}} = \left(\left(\frac{F_{\text{HI}}}{W_{50}} \right) \sqrt{\frac{W_{50}}{2R}} \right) / \sigma_{\text{rms}}, \quad (7)$$

where the integrated flux density of the detection is given by F_{HI} (Jy km s^{-1}), W_{50} (km s^{-1}) is the width of the spectrum at 50 per cent the peak value,⁸ and R is the velocity resolution of the spectrum (this is also called the channel width, unless the spectrum has been smoothed along the spectral axis).

⁸The W_{50} is obtained from the single Gaussian fit to the spectrum shown in Fig. A4.

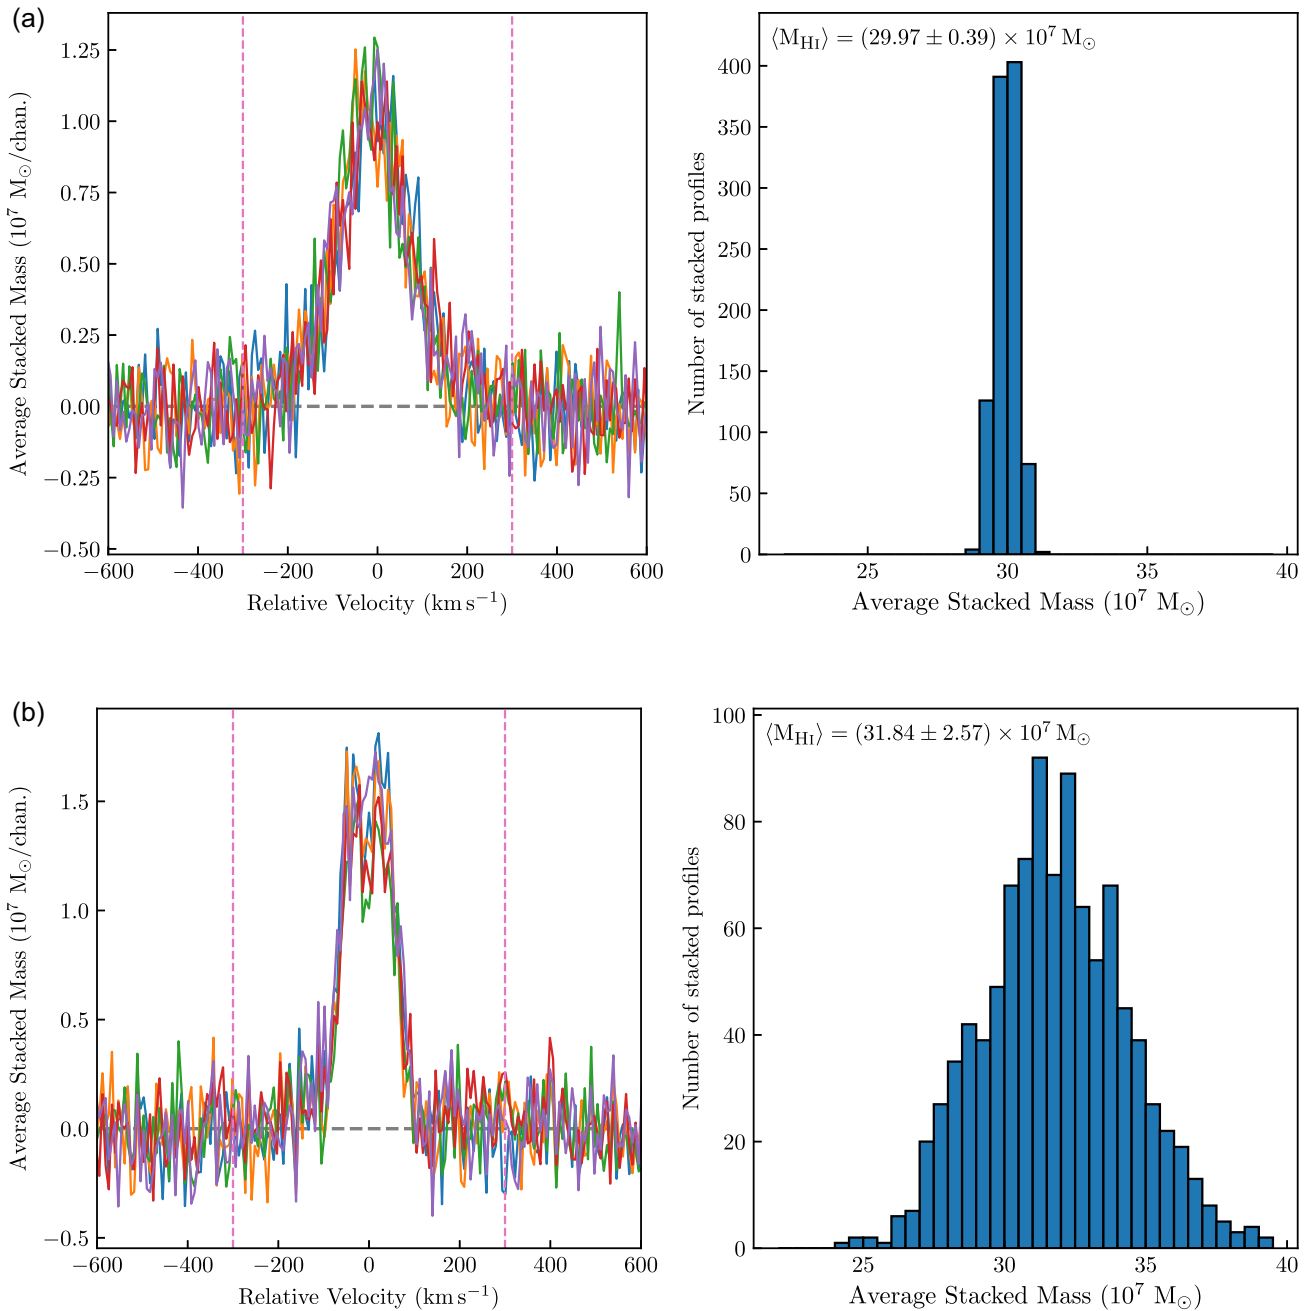


Figure 5. Uncertainty methods. (a) Redshift error analysis: The left-hand panel shows five iterations of the stacked spectrum. Each of the stacked spectra has been created by varying the redshifts of the individual spectra. The error bars for the final stacked spectrum are determined from the range of values in each channel after 1000 iterations. The right-hand panel shows the $\langle M_{\text{HI}} \rangle$ calculated from each iteration of the stacked spectrum by integrating the spectrum between the two vertical dashed lines. (b) Statistical error analysis: The left-hand panel shows five iterations of the stacked spectrum. The error bars for the final the stacked spectrum are determined from the range of values in each channel after 1000 iterations. Each of the stacked spectra has been created using a sub-sample of 75 per cent of the full sample. The right-hand panel shows the $\langle M_{\text{HI}} \rangle$ calculated from each iteration of the stacked spectrum by integrating the spectrum between the two vertical dashed lines.

The next statistic calculated is the probability that the signal in a stacked spectrum arises due to a random fluctuation – this is known as the p -value, which will be as low as possible for a strong detection, from which the significance level of the detected signal can be calculated. This is done by fitting a single Gaussian to the spectrum and calculating the χ^2 value. The single Gaussian

is used in favour of other models as it is the simplest model that can describe the shape of an H I detection to first order; please note that in the case of a very clear detection which cannot be characterized by a single Gaussian, the p -value statistic provides no extra information. The null hypothesis is a horizontal line at zero.

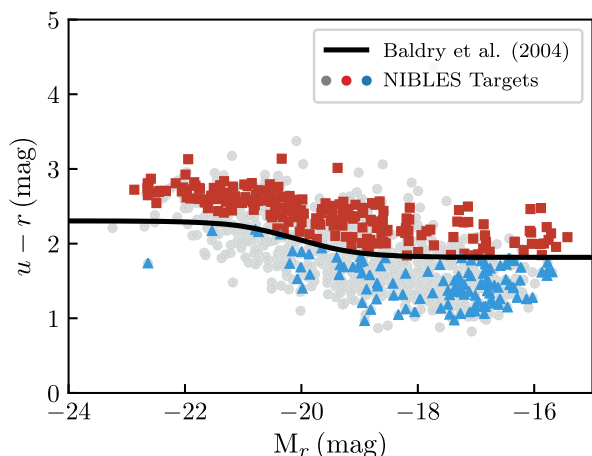


Figure 6. Optical colour–magnitude diagram for the NIBLES stacking sample galaxies: $u - r$ colours derived from SDSS model magnitudes corrected for Galactic extinction as a function of absolute r -band magnitude. The H I non-detected galaxies are plotted in either red (squares) or blue (triangles) based on their position above or below the Baldry et al. (2004) red/blue colour divider which is indicated by the black line. The grey data points represent the H I detected galaxies.

2.4.2 Spectrum manipulation

In order to boost the S/N of a stacked spectrum, the Analysis Module offers the user the choice of either re-binning the data (lowering the data resolution, this method has been used by Lah et al. 2007; Rhee et al. 2016), or smoothing the spectrum using either a boxcar algorithm or a Hanning Window.

All three manipulation routines allow the user to select either the window or the number of old bins to be rebinned into new bins. HISS will not continue to the next step until the user is satisfied with the new version of the stacked spectrum. Diagnostic Plot 2 (which is identical to Diagnostic Plot 1, Fig. A4, but displays the updated spectrum) along with a recalculation of the initial statistics is produced every time a change is made to the spectrum.

2.4.3 Profile characterization

Given that stacked spectra can have very low S/N, with a lot of noise spikes and dips that are unhelpful in determining the integrated flux density and line width, fitting a function to a stacked spectrum can assist in estimating the stacked profile properties. For example, determining the H I line width (W_{50}) of stacked spectra can be useful for studies such as determining the H I Tully–Fisher relation⁹ of non-detected galaxies (Meyer et al. 2016). The analysis module prompts the user to choose from a selection of seven different functions (between one and seven can be chosen simultaneously) that have been used in previous H I studies to characterize H I line profiles:

- (i) Single Gaussian
- (ii) Double Gaussian
- (iii) Lorentzian Distribution
- (iv) Voigt Profile
- (v) Third Order Gauss–Hermite Polynomial

- (vi) Busy Function (Westmeier et al. 2014) – with double-horn
- (vii) Busy Function – without double-horn

Fig. A5 shows the seven functions fitted to the stack of the 1000 simulated profiles. Also shown is a table containing the integrated flux density, a measure of the goodness of fit, and the full width at half-maximum of the fitted function. The user may select any number of the functions, and HISS checks that the user is satisfied with the selection before continuing. A version of Fig. A5 is saved in PNG¹⁰ format in the output location as Diagnostic Plot 3. The diagnostic plot is saved to file if HISS is run in suppress mode (an option the user may select when initializing HISS which suppresses all displayed output) so that the user may inspect the function fit before continuing with the analysis.

2.4.4 Uncertainty calculation

The Uncertainty Module allows for two types of uncertainty calculations: a statistical error analysis and a redshift error analysis. The user specifies the type of uncertainty calculation (the two methods cannot be run simultaneously). The Uncertainty Module facilitates the storing of the final analysis options chosen by the user and uses those options to repeat the stacking process 1000 times, each time with slightly different parameters. Fig. 1 indicates the Uncertainty Module as a wrapper around the Stack and Analysis modules.

(1) Redshift error analysis: if the user chooses to engage the redshift uncertainty calculations, another loop is activated around the Stack and Analysis Modules which repeats the two processes one thousand times. In every iteration, the redshift associated with each spectrum in the input catalogue is changed by an amount that is selected from a normal distribution with a standard deviation equivalent to the particular spectrum’s redshift uncertainty. Thus, the redshift for each spectrum becomes $z = z + dz$, where dz has been sampled from a normal distribution defined by $N(\mu = z; \sigma = u(z))$.

The result of repeating the stack and analysis process means that there are 1000 slightly different versions of the stacked spectrum and their corresponding integrated flux values. Each of the 1000 versions of the analysis object (which each contain a selection of integrated fluxes and the stacked spectrum data) are stored by the Uncertainty Class. Upon completion, the Uncertainty Class processes the stored data: the error bars on the spectrum data are calculated from the minimum and maximum values per channel of the 1000 stored spectra, the data points are given by the mean values. The error bars on the integrated fluxes are determined from the difference between the mean value (which serves as the quoted value) and the 25th and 75th percentile.

An example of the stacked spectrum using this error calculation method and an average $dz = 60 \text{ km s}^{-1}$ is shown in Fig. 5(a).

(2) Statistical error analysis: the statistical error analysis implemented by this module is the *Delete-A-Group Jackknife* (Kott 2001) error method. In the same way as mentioned above, another loop is activated around the Stack and Analysis Modules repeating the process a thousand times, however each iteration discards a percentage (as chosen by the user) of the total catalogue without replacement. The 1000 analysed stacked spectra are stored by the Uncertainty Module. The flux values for the final stacked spectrum

⁹An empirical relation which relates the luminosity of galaxy to its rotation velocity (Tully & Fisher 1977).

¹⁰Portable Network Graphics.




Figure 7. The stellar mass distribution for the full sample (*left*) and the non-detections (*right*). The grey bars indicate the total number of galaxies per bin, while the blue and red lines indicate the number of blue and red galaxies per mass bin. Each bin is 0.5 dex wide. The dashed green line highlights the ‘gas-richness’ threshold (Kannappan et al. 2013), where gas-rich galaxies are to the left of the line (lower M_*) and gas-poor galaxies are to the right of the line (higher M_*).

are taken as the mean value of the 1000 versions in each channel. The error bars on the flux in each channel are calculated according to

$$\sigma(s) = \sqrt{\frac{R-1}{R} \sum_{n=0}^R (s_n - \bar{s})^2}, \quad (8)$$

where R is the number of repeated estimates (which would be 4 if only 75 per cent of the population was used for each iteration), s_n is the flux for the particular channel of the n th subset, and \bar{s} is the mean flux value for the particular channel (Kott 2001; Brown et al. 2015). This is illustrated in Fig. 5(b).

The shape of the redshift error analysis profile is smeared out into a more Gaussian shape than the statistical error analysis stacked profile due to each spectrum being allocated its correct redshift plus an offset value. This results in the individual line centres being smeared around the zero-point. This was also observed by Maddox et al. (2013).

The choice of preferred error analysis method is left to the user to decide based on their particular sample selection and optical catalogue redshift precision.

2.5 Save results

There are nine possible files that can be saved to the output location – five data files and four possible diagnostic plots (these inform user interactions at point 2):

(i) **Stacked catalogue:** this text file contains a list of all the spectra included in the stack. Along with the data provided by the user from the input catalogue, this table also contains the integrated flux density for each spectrum.

(ii) **Output data:** a FITS table file containing the spectrum data (stacked spectrum, spectral axis, and reference spectrum), the fitted parameters of any fitted functions, the stacked noise, and if the

spectrum has been smoothed then the original version of the spectrum data is also saved.

(iii) **Stacked spectrum plot:** a PDF file of the stacked spectrum plotted with the fitted functions and the reference spectrum (a version of this is shown in the top left panel of Fig. A7).

(iv) **Stacked noise plot:** a PDF file of the stacked noise which has the expected σ/\sqrt{N} line overplotted (top right panel of Fig. A7).

(v) **Integrated flux data file:** a table of the calculated integrated flux from the different functions as well as the flux integrated within the galaxy window. This file is saved in Encapsulated Comma Separated Value (ecsv) format using the ASTROPY ASCII module, which allows the units of the flux to be saved to file. Other columns in this file include the goodness-of-fit values from the function fits. A version of this table is shown in the bottom panel of Fig. A7.

(vi) **Diagnostic Plot 1:** this is the first plot of the stacked spectrum that is displayed and saved upon initialization of the Analysis Module (Fig. A4).

(vii) **Diagnostic Plot 2:** a PNG file that is only produced if the user decides to rebin or smooth the spectrum. This plot is almost identical to Fig. A4.

(viii) **Diagnostic Plot 3:** a PNG file that is created when the user selects a function to fit to the stacked spectrum (Fig. A5).

(ix) **Diagnostic Plot 4:** this plot is only produced if the uncertainties have been calculated. The plot contains a series of histograms showing the spread of the integrated fluxes. Diagnostic Plot 4 is saved to file in PNG format (Fig. A6).

2.6 Displaying the results

Upon conclusion of the analysis of the stacked spectrum, all calculated quantities and plots are saved to the user-chosen output location. Amongst the saved files is a plot of the stacked spectrum which makes this final module an optional one. The display module




Figure 8. The three columns in the figure above show the stacked profiles for the full sample, blue sample, and red sample respectively. In the top row are the stacked profiles for all spectra (blue triangles), only detections (purple squares), and only non-detections (black circles). The bottom row shows the stacked profiles for the non-detections, the reference spectrum is shown in green, and the grey band represents the uncertainty of the stacked spectrum. The vertical yellow regions in each of the panels indicate the region over which the spectra are integrated to obtain the average H I mass, $\langle M_{\text{HI}} \rangle$.

reads in the files saved to disc at the end of the analysis. The results are then displayed in an easy to read manner – the stacked spectrum (along with the reference spectrum) is shown on the top left, to the right of the stacked spectrum is the stacked noise as a function the number of profiles. Below the two plots are the statistics associated with stacked spectrum, and a table containing the calculated integrated fluxes. A screenshot of the display window is shown in Fig. A7.

3 APPLICATION OF HISS TO THE NIBLES SURVEY

In this second half of the paper, we will demonstrate some of the capabilities of HISS by stacking different sub-samples of galaxies observed as part of the NIBLES (van Driel et al. 2016), and revisiting the well-known H I scaling relations (e.g. Catinella et al. 2010, 2012, 2013; Brown et al. 2015).

NIBLES is a targeted H I survey of galaxies selected by M_z , a proxy for stellar mass, from the Sloan Digital Sky Survey (SDSS; York et al. 2000). The H I data were obtained using the 100 m NRT in the centre of France. NIBLES provides a unique opportunity to study the baryonic content in galaxies with a wide range of stellar mass [$10^6 < M_*(M_{\odot}) < 10^{12}$].

3.1 H I data

The NIBLES sample of 2600 galaxies was selected from SDSS Data Release 5 (DR5; Adelman-McCarthy et al. 2007). The selected galaxies were chosen to uniformly sample the stellar mass range of the ensemble of galaxies in the local Universe. The particular selection criteria (van Driel et al. 2016) were as follows:

- (i) The target must have both SDSS DR5 spectroscopy and photometry.
- (ii) NIBLES targets were limited to the Local Volume: $900 < cz < 12000 \text{ km s}^{-1}$.
- (iii) In order to study the H I in galaxies as a function of stellar mass, ~ 150 galaxies were selected per 0.5 mag bin for $-16.5 \text{ mag} \leq M_z \leq -24 \text{ mag}$ ($H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$), if available. M_z was used as a proxy for stellar mass.¹¹
- (iv) In order to maintain a morphologically diverse sample, no colour selection was made.

In addition to these criteria, van Driel et al. (2016), when selecting the target galaxies tried to avoid objects in and around the Virgo

¹¹ SDSS stellar mass catalogues were not yet publicly available at the time of the original NIBLES sample selection which was based on DR5 photometry.

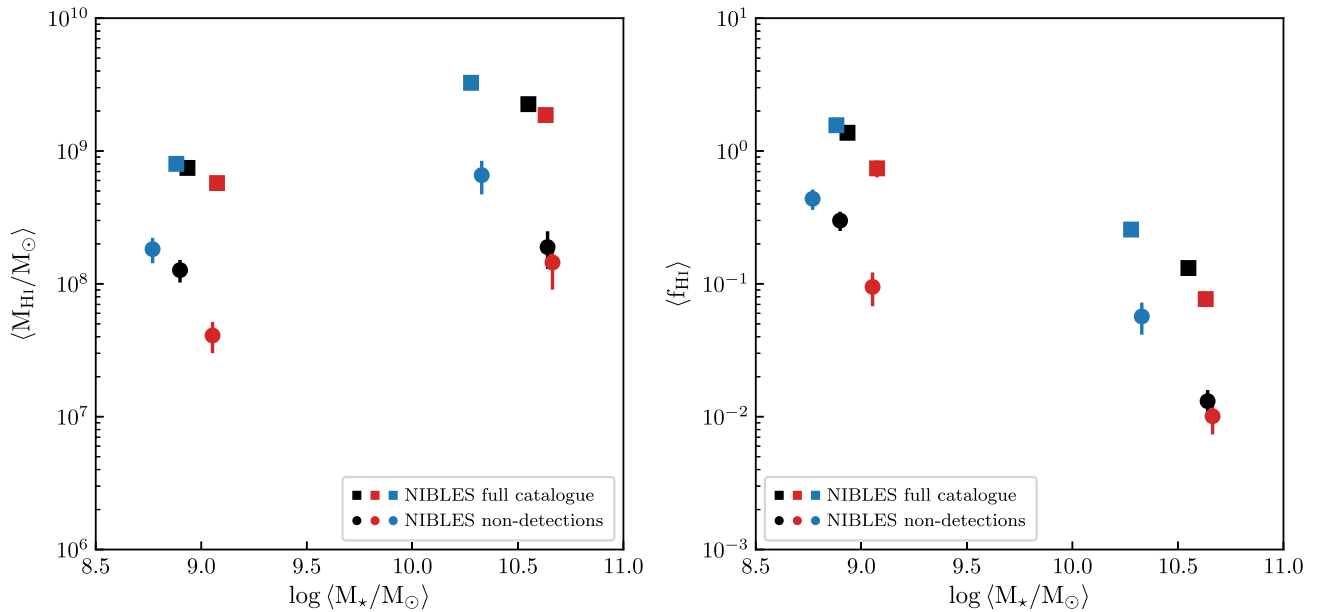


Figure 9. Average $M_{\text{HI}}/M_{\text{HI}}$ (left), and the $(M_{\text{HI}}/M_{\star})$ mass ratio, the average gas fraction $\langle f_{\text{HI}} \rangle$ (right), as a function of M_{\star} for the full NIBLES sample as well as the sub-sample of H I undetected galaxies. Each sample has been split into two stellar mass bins: $10^8 < M_{\star} (M_{\odot}) < 10^{9.5}$ and $10^{9.5} < M_{\star} (M_{\odot}) < 10^{12}$. Shown here are the mean M_{\star} of the galaxies in each of the two stellar mass bins, for each of the various samples. The round data points represent the non-detections, and the square points represent the full sample. The blue, red, and black data points represent respectively the blue, red, and total sub-samples. The error bars represent the statistical uncertainties, including on the stacks of non-detections. The stacked profiles which each data point represents, are presented in Appendix D2 (the full NIBLES sample in Figs D1 and D2, and the non-detected sub-sample in Figs D3 and D4).

cluster as well as the ALFALFA $\alpha.40$ footprint (Haynes et al. 2011). The dense environment in the Virgo cluster is known to have noticeable effects on the H I properties of the galaxies in the region. The distances to galaxies in the Virgo region also have large uncertainties.

The NRT (see van Driel et al. for further details) is a 100 m single dish radio telescope located in France. The telescope is a transit instrument with a collecting area of 6900 m² and consists of two large reflectors. At 21 cm, the NRT beam spans 3.6 arcmin in right ascension and varies in declination from 22 to 33 arcmin depending on the declination of the source; the exact declination size of the beam can be determined using equation (B1).

The spectra were collected from 2007 January to 2010 December, for a total of 3500 h of telescope time. Each target was initially observed for roughly 40 min using successive 40 s ON source and 40 s OFF source pointings, which was repeated as required and as observing time permitted. The OFF source pointing was chosen to be 20' E of each target. The data reduction process is discussed in detail by van Driel et al. (2016). Each published spectrum has a velocity resolution of 18 km s⁻¹ and a typical noise level of 2.5 mJy.

3.1.1 Higher-sensitivity Arecibo follow-up H I observations

An indication that stacking Nançay H I spectra may lead to the detection of a significant signal can be found in the higher sensitivity follow-up observations that were obtained with the 305m Arecibo radio telescope, which resulted in a high detection rate of NRT non-detections (Butcher et al. 2016). Of the Nançay targets, 71 non-detections and 19 marginal detections were observed at Arecibo with a four times higher sensitivity (mean rms 0.57 mJy). This led to the detection of 85 percent and 89 percent of the NRT non-detections and marginal detections, respectively. It should be noted that most (59 or 66 percent) of the re-observed NRT non-

detections are blue, low-mass galaxies ($\sim 10^{6.5} < M_{\star} (M_{\odot}) < 10^{8.5}$) which are expected to be relatively gas-rich (see Section 4). Within the same stellar mass range, the $\langle f_{\text{HI}} \rangle$ of the new Arecibo detections is on average three times lower than for the Nançay detections. Alternatively, using the HISS approach to reach the four times lower Arecibo rms level we could also have stacked 16 NRT spectra. The analysis of similar Arecibo data of a much larger sample of NIBLES Nançay non-detections of all colours has recently been published by Butcher et al. (2018). The Arecibo data mentioned above were not used for the stacking study presented here.

3.1.2 Comparison between NIBLES and other H I surveys

Upon comparison with other single-dish H I surveys [e.g. ALFALFA, HIPASS (Barnes et al. 2001); see table A.4 in van Driel et al. 2016], the ALFALFA $\alpha.40$ fluxes were found to be higher than the NIBLES fluxes by a factor $\alpha.40/\text{NRT} = 1.45 \pm 0.17$. The comparisons between NRT fluxes and the ALFALFA fluxes, as well as possible reasons for the difference in flux scales, are discussed in detail in section 4.6 in van Driel et al. (2016). In this work, the NIBLES fluxes are quoted. In the event of comparison between NIBLES measurements and ALFALFA measurements, the ALFALFA measurements are scaled to the NIBLES flux values using the ALFALFA/NRT flux ratio mentioned above.

3.2 Optical data

The optical photometric data used in this work is obtained from the SDSS Data Release 12 where the photometric quality flag was marked clean with no deblended ‘child’ objects. In the DR12 database, the spectroscopic sources (from which the NIBLES sample is derived) are matched to the ‘value-added’ JHU/MPA catalogues (Brinchmann et al. 2004) which contain the stellar

masses used in this work. The catalogue provides a number of different measures of the stellar mass, however for the purposes of the NIBLES analyses, the median values are used (van Driel et al. 2016; Butcher et al. 2018)

3.3 NIBLES stacking sample

3.3.1 Accounting for known contamination sources

Studies have shown that galaxies which are nearby to the target galaxies in terms of both distance on the sky and systemic velocity can significantly contribute to the total emission in a target H I galaxy spectrum (Elson, Blyth & Baker 2016; Jones et al. 2016). For example, Elson et al. (2016) showed that for a stacking experiment using simulated Parkes data at 15 arcmin resolution at $0.04 < z < 0.13$, the average contaminant H I mass per galaxy spectrum was $\sim 1.4 \times 10^{10} M_{\odot}$. Thus, before finalizing the sample of spectra for stacking, it is necessary to remove spectra that are known to have nearby galaxies that could contribute significant contaminating emission to the spectrum. In other works (e.g. Fabello et al. 2012), a secondary source (irrespective of morphology) is considered to contaminate the target spectrum if it lies spatially within the half-power beam width and $\pm 300 \text{ km s}^{-1}$ of the target. In this work, galaxies near the target source are considered as possible contaminators if

- (i) the secondary galaxy is located spatially within $1.2 \times \text{N-S}$ half-power beam width and $2 \times \text{E-W}$ half-power beam width, and
- (ii) the redshift of the secondary galaxy is within $\pm 300 \text{ km s}^{-1}$ of the target redshift.

The above criteria were used to search for sources in both the SDSS Spectroscopic Database and the NASA/IPAC Extragalactic Database. A total of 761 spectra were classified as potentially contaminated using these criteria. For more details see Appendix B.

3.3.2 Defining H I detections versus non-detections

NIBLES galaxies were originally classified as ‘detected’, ‘marginal’, or ‘not detected’ as explained in van Driel et al. (2016) by three independent adjudicators, who used statistical parameters describing the profile S/N and also inspected all spectra by eye. For this work we wanted a binary classification of ‘detected’ or ‘not detected’ for our stacking experiments and we wanted to use a method which did not rely on human intervention. Various methods exist for finding signals and classifying detections in H I spectra (see for example Saintonge 2007; Verheijen et al. 2007; Haynes et al. 2011; Ramatsoku et al. 2016). Using the original classifications in van Driel et al. (2016) as a guide, we used classification criteria, adapted from Ramatsoku et al. (2016) to reproduce the sample selection of van Driel et al. (2016), while at the same time self-consistently and objectively classifying the marginal detections as either detections or non-detections. This method produced very similar detection classifications to van Driel et al. (2016) with over 96 per cent overlap in the classifications between the two schemes for both detections and non-detections, and resulted in the NIBLES marginal detections (52) being allocated as detections versus non-detections in a ratio of one-third to two-thirds in our sample.

Following Ramatsoku et al. (2016) we classify galaxies as detected if there are a number of consecutive channels above a threshold value. The threshold values are set based on the noise level of the individual spectra which we estimate using the sub-interquartile range (IQR – half the difference between the 75th and

25th flux percentiles). The IQR is less sensitive to outliers in a distribution than the standard deviation and therefore eliminates the need for manual masking of the target source and/or artefacts in the spectrum when estimating the noise (see Fig. C1). For classification of NIBLES spectra, we considered a window of 600 km s^{-1} centred on the target redshift. If any emission, satisfying the criteria below, was found in the window, the spectrum was classified as a detection:

- (i) 1 or more channels with flux density above $5 \times \text{IQR}$ (within a 400 km s^{-1} window)
- (ii) 2 or more consecutive channels above $4 \times \text{IQR}$
- (iii) 4 or more consecutive channels above $3 \times \text{IQR}$
- (iv) 5 or more consecutive channels above $2 \times \text{IQR}$

Visually inspecting the spectra after classification showed that some sources had an H I detection in the off-source pointing which presents in the NIBLES spectra as an absorption-like feature, while other spectra had clear emission features from sources not previously identified in our possible contaminant source search (see Section 3.3.1). Examples of spectra with these features are presented in Fig. C2. These features can bias the stacked spectrum and therefore in order to flag these spectra and eliminate them from our sample, we ran the classification algorithm a second time over a wider velocity range ($\pm 600 \text{ km s}^{-1}$ centred on the target redshift and indicated by the vertical green dashed lines in Fig. C2), to identify possible positive and negative emission not related to the target galaxy. We rejected galaxy spectra that had either negative emission satisfying our classification criteria in any part of the larger search window ($c z \pm 600 \text{ km s}^{-1}$) or positive emission satisfying the criteria and falling outside a central range of $c z \pm 200 \text{ km s}^{-1}$ but within $c z \pm 600 \text{ km s}^{-1}$. In total 24 spectra were removed from the stacking sample in this way.

3.3.3 Final stacking sample

We excluded from our sample galaxies with $M_{\star} < 10^8 M_{\odot}$ due to unreliable SDSS photometry and therefore unreliable stellar masses. Our final catalogue consists of 1000 H I spectra deemed free of nearby contaminant emission. Of the 1000 sources, 323 were classified as non-detections and 677 as detections (see Fig. C3 for examples). The catalogue for the non-detections can be found in Table E1 and Fig. E1 in Appendix E.

4 GAS SCALING RELATIONS IN NIBLES

As has been previously stated, the NIBLES sample is morphologically diverse, and it ranges from gas-rich late-type galaxies to gas-poor early-type galaxies. In order to explore the contributions from the different morphological types to the stacked profiles, we separate the stacking sample into red and blue sub-samples using the Baldry et al. (2004, equation 9) optimal colour divider as shown in Fig. 6.

Having separated the full sample into blue and red sub-samples, there are 549 blue galaxies (436 detections and 113 non-detections) and 451 red galaxies (241 detections and 210 non-detections). The stellar mass distributions for the full sample and non-detections alone, and each broken down by colour, are shown in Fig. 7.

It is clear in both the full sample and the non-detected sample that the blue galaxies are predominantly in the low- M_{\star} regime, while the red sample contains mainly high- M_{\star} galaxies. The transition from predominantly blue galaxies to predominantly red galaxies occurs in our full sample around $M_{\star} = 10^{9.5} M_{\odot}$ which also corresponds

to the so-called ‘gas-richness’ threshold (Kannappan et al. 2013) which signifies the transition from HI gas-rich galaxies (at lower M_*) to HI gas-poor galaxies (at higher M_*).

4.1 The average HI properties

We separated the full sample into three samples: all spectra (full sample), detections only, and non-detections only. The spectra are stacked in units of M_{HI} and the resulting spectra are shown in the left column of Fig. 8. The top row of Fig. 8 shows the stacked profiles for the full sample, detections, and non-detections on the same set of axes; the bottom row shows only the stacked profiles for the non-detections, enlarged for clarity. The top left panel shows the comparison between the stacked spectra for the full sample (blue triangles), and the spectra for the sample of only detections (purple squares) and only non-detections (black filled circles).

The middle and right-hand panels of Fig. 8 show the stacked spectra for the blue and red sub-samples, respectively. As indicated by the vertical yellow bands, we integrate over a narrower velocity range for the blue sub-samples compared to the red sub-samples. The blue stacked spectra are expected to be narrower since the average stellar mass of the blue sample is lower than the red sample.

The average HI mass and M_{HI} to M_* ratio (f_{HI}) measured from the stacked profiles for each of the nine sub-samples (all galaxies, all blue galaxies, all red galaxies; all HI detections, blue HI detections, red HI detections; all non-detections, blue non-detections, and red non-detections) are shown in Table D1.

HI stacking is a powerful tool to probe the average HI properties of various galaxy samples, however one should have a clear idea of the morphological properties of the galaxies that make up the sample. We have shown how the stacked HI profiles can change by only looking at blue or red galaxies – we find that the stacked profiles for the red samples are wider than the stacked profiles for the blue samples. This is to be expected based on the bimodal distribution of local galaxies where the red galaxies on average have higher stellar masses, and in line with the Tully–Fisher relation, higher rotation velocities (see also Fabello et al. 2011; Meyer et al. 2016).

We divided our sample into two bins of low and high stellar mass ($M_* < 10^{9.5} M_\odot$ and $M_* > 10^{9.5} M_\odot$) to explore the average HI properties of blue and red galaxies in these two samples. As outlined earlier in this paper, HISS is capable of directly determining the average HI mass or HI gas fraction from stacked spectra. The HI-to-stellar mass fraction (gas fraction or f_{HI}) is a useful tool to probe how gas rich a galaxy or sample of galaxies is. Our results are summarized in Fig. 9 and the stacked profiles are shown in Appendix D2, with the measured quantities presented in Tables D2 and D3. Note that all the stacked profiles have $S/N_{\text{ALFALFA}} > 6.5$ and are also classified as detections according to the criteria in Section 3.3.2.

These results are consistent with work by Kannappan (2004) which shows that while low-mass galaxies tend to have lower HI masses than their high-mass counterparts, the low stellar mass galaxies tend to have higher HI gas fractions. The right-hand panel of Fig. 9 highlights what was found by Brown et al. (2015) that the steep slope of the average $\langle f_{\text{HI}} \rangle$ versus M_* trend (indicated by the black squares and filled circles) from low to high stellar mass is driven by the predominance of high gas fraction blue galaxies at low stellar mass and the low gas fraction red galaxies at high stellar mass. This trend is seen for both the full stacked sample and when stacking non-detections alone.

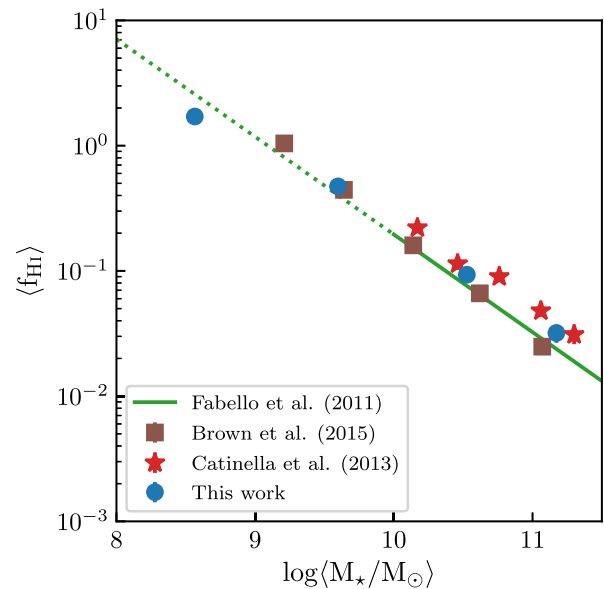


Figure 10. HI gas fraction ($\langle f_{\text{HI}} \rangle$) as a function of M_* for the full NIBLES stacking sample, including both detections and non-detections. Also shown, are $\langle f_{\text{HI}} \rangle$ values from Brown et al. (2015) who stacked sub-samples of SDSS-selected galaxies with HI data from ALFALFA, and Catinella et al. (2013) who calculated weighted averages for sub-samples of GASS galaxies. The green line is taken from Fabello et al. (2011) who fitted the slope of the gas scaling relation for a sub-sample of ALFALFA galaxies; the solid line represents the stellar mass range of their sample while the dotted part is a linear extrapolation. The error bars on the NIBLES data points are the statistical uncertainties. The plotted values are given in Table 1, and the stacked profiles are shown in Fig. D5.

4.2 Gas fractions in NIBLES galaxies

Gas scaling relations have been used in conjunction with stacking (e.g. Fabello et al. 2011, hereafter F11; Brown et al. 2015, hereafter B15) to study various galaxy properties (star formation, stellar mass, etc.) and their influence on the HI content of the galaxies. In Fig. 10 we compare NIBLES gas fractions calculated using HISS to gas fractions from Catinella et al. (2013, hereafter C13), F11, and B15. Adopting the method used by B15, we separated the full NIBLES sample (detections plus non-detections) into bins of $\log M_*$ with widths of 1 dex. The sample stellar mass distribution is presented in Fig. 7. For the NIBLES sample in Fig. 10 (indicated by the round blue circles), the M_* values represent the mean value in each bin. The NIBLES stacked profiles are presented in Fig. D5 and the measured quantities are listed in Table 1.

The C13 sample was selected from the GALEX Arcibo SDSS Survey (GASS, Catinella et al. 2010, 2013), a survey with the Arcibo telescope which targeted ~ 1000 massive galaxies randomly selected from the overlap between the SDSS DR6 spectroscopic survey and GALEX Medium Imaging Survey. The F11 and B15 samples are selected from the overlap between ALFALFA, SDSS, and GALEX. The C13 results show weighted average values (the HI non-detections are set to upper limits) while the F11 and B15 results are obtained from HI stacking. To take into account the known flux offset between ALFALFA HI spectra and NIBLES (see Section 3.1.2), the B15 and F11 gas fractions have been scaled by the mean $\alpha.40/\text{NRT}$ flux ratio (1.45). The solid green line is the fit from F11 to their gas scaling relation and the dotted green line represents a linear extrapolation of this trend to lower stellar mass.

Table 1. Average H I gas fractions ($\langle f_{\text{HI}} \rangle$) for the scaling relation in Fig. 10. The average M_{HI}/M_* ratio ($\langle f_{\text{HI}} \rangle$) values for each data point are given along with the statistical uncertainties. The total number of galaxies per bin are given in the column titled N , with the number of non-detections per bin indicated by the number in brackets. The second-last column gives the Gaussian significance of the stacked profile obtained from HISS (see Section 2.4.1 for details). The final column gives the S/N of the stacked spectrum calculated by HISS using equation (7).

$\log \langle M_*/M_\odot \rangle$	$\langle f_{\text{HI}} \rangle$	N	Significance	S/N _{ALFALFA}
8.56	1.709 ± 0.089	347 (109)	$>8.2\sigma$	174.7
9.60	0.473 ± 0.024	335 (96)	$>8.2\sigma$	231.8
10.53	0.093 ± 0.005	272 (93)	$>8.2\sigma$	216.6
11.17	0.032 ± 0.005	46 (25)	$>8.2\sigma$	49.3

Due to the NIBLES sample selection, we are able to go down an order of magnitude lower in stellar mass than previous results by B15. For $M_* > 10^9 M_\odot$ the NIBLES gas fractions agree well with the results by C13 and the adjusted gas fractions from B15 and F11. Below $M_* < 10^9 M_\odot$, the NIBLES $\langle f_{\text{HI}} \rangle$ is lower than the trend set by the higher stellar mass galaxies.

Huang et al. (2012) and Maddox et al. (2015) have previously noted the change in slope of the M_{HI} versus M_* relation between $10^8 M_\odot < M_* < 10^9 M_\odot$ in the H I-selected population of ALFALFA galaxies where the relation is steeper at lower mass and flatter towards the high-mass end. A change in slope is also seen in the ALFALFA $f_{\text{HI}}-M_*$ scaling relation by Huang et al. (2012) around a similar M_* where the slope is flatter at lower M_* and drops more steeply at higher M_* . Huang et al. (2012) note that their average gas fraction relation has an overall offset to higher values as expected compared to the optically selected samples by Catinella et al. (2010) and Cortese et al. (2011) due to the ALFALFA sample bias towards more gas-rich galaxies per stellar mass bin. Across the mass range investigated, per bin in stellar mass, they also observe a sharp cut-off in number density at the higher end of the gas fraction distribution, but a large dispersion towards lower gas fractions and note that for surveys with longer integration times, a lower average f_{HI} per stellar mass bin would be observed. A similar change in slope towards lower M_* would also therefore be expected for stellar mass selected samples as for the ALFALFA sample. The divergence of the lowest stellar mass NIBLES gas fraction data point from the higher stellar mass trends from the extrapolated F11 fit is therefore consistent with expectations based on the H I-selected ALFALFA sample Huang et al. (2012) gas fraction versus M_* relation.

5 SUMMARY

This work has detailed the development of a new 21 cm H I spectral stacking package in PYTHON (H I Stacking Software: HISS); while we have exclusively used H I emission spectra to test HISS, it is capable of stacking emission or absorption spectra of any characteristic line. HISS takes the spectra for a sample of galaxies along with the accompanying redshifts recorded in the galaxy catalogue to produce an average spectrum from which a number of average properties (e.g. H I mass $\langle M_{\text{HI}} \rangle$, H I-to-stellar mass ratio $\langle f_{\text{HI}} \rangle$) for the sample may be extracted. HISS also offers the user a choice of two built-in error analysis methods, and the ability to characterize the shape of the stacked spectrum through fitting a variety of functions.

We have applied HISS to stacking of 1000 H I spectra from the NIBLES, which is a stellar mass-selected (with no colour-selection) H I survey of SDSS galaxies in the nearby Universe. Due to the wide stellar mass range spanned by the NIBLES dataset ($10^8 < M_* (M_\odot) < 10^{12}$), extending an order of magnitude lower than Brown et al. (2015), we were able to extend the previously

studied gas fraction versus stellar mass gas scaling relation (Fabello et al. 2011; Catinella et al. 2013; Brown et al. 2015) down to lower average stellar mass ($\langle M_* \rangle = 10^{8.6} M_\odot$). With our stellar mass selected sample, we find good agreement with previous results at high stellar mass. At low stellar mass ($M_* < 10^9 M_\odot$) we observe a deviation from the extrapolated high-mass relation which indicates a flattening of the slope in the scaling relation at low stellar mass qualitatively consistent with the trend seen in the H I-blind, gas-rich ALFALFA sample (Huang et al. 2012).

ACKNOWLEDGEMENTS

JH acknowledges the bursary provided by South African Radio Astronomy Observatory; SLB, EE, and JH were supported by the South African National Research Foundation. We thank Barbara Catinella and Kelley Hess for their input and advice on the design of HISS. JH wishes to thank S. Makhathini and L. Lindroos for helpful conversations during a trip to the Onsala Space Observatory, supported by MIDPREP. The Nançay Radio Astronomy Facility is operated as part of the Paris Observatory, in association with the Centre National de la Recherche Scientifique (CNRS) and partially supported by the Région Centre in France. This research has made use of the NASA/IPAC Database (NED) which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the Nation Aeronautics and Space Administration, and images and data from Sloan Digital Sky Survey (SDSS-III). Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>. SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration.

REFERENCES

- Adelman-McCarthy J. K. et al., 2007, *ApJS*, 172, 634
 Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Z., Lupton R. H., Nichol R. C., Szalay A. S., 2004, *ApJ*, 600, 681
 Barnes D. G. et al., 2001, *MNRAS*, 322, 486
 Bianchi L., GALEX Team, 2000, *Mem. Soc. Astron. Ital.*, 71, 1117
 Booth R. S., de Blok W. J. G., Jonas J. L., Fanaroff B., 2009, preprint ([astro-ph/0910.2935](https://arxiv.org/abs/astro-ph/0910.2935))
 Brinchmann J., Charlot S., Heckman T. M., Kauffmann G., Tremonti C., White S. D. M., 2004, preprint ([astro-ph/0406220](https://arxiv.org/abs/astro-ph/0406220))
 Brown T., Catinella B., Cortese L., Kilborn V., Haynes M. P., Giovanelli R., 2015, *MNRAS*, 452, 2479 (B15)
 Brown T. et al., 2017, *MNRAS*, 466, 1275
 Butcher Z., Schneider S., van Driel W., Lehnert M. D., Minchin R., 2016, *A&A*, 596, 20
 Butcher Z., Schneider S., van Driel W., Lehnert M. D., 2018, *A&A*, 619, 14
 Catinella B. et al., 2010, *MNRAS*, 403, 683

- Catinella B. et al., 2012, *A&A*, 544, 30
 Catinella B. et al., 2013, *MNRAS*, 436, 34 (C13)
 Catinella B., Cortese L., 2015, *MNRAS*, 446, 3526
 Chengalur J. N., Braun R., Wieringa M., 2001, *A&A*, 372, 768
 Cortese L., Catinella B., Boissier S., Boselli A., Heinis S., 2011, *MNRAS*, 415, 1797
 Delhaize J., Meyer M. J., Staveley-Smith L., Boyle B. J., 2013, *MNRAS*, 433, 1398
 Elson E. C., Blyth S. L., Baker A. J., 2016, *MNRAS*, 460, 4366
 Fabello S., Catinella B., Giovanelli R., Kauffmann G., Haynes M. P., Heckman T. M., Schiminovich D., 2011, *MNRAS*, 411, 993 (F11)
 Fabello S., Kauffmann G., Catinella B., Li C., Giovanelli R., Haynes M. P., 2012, *MNRAS*, 427, 2841
 Fernández X. et al., 2016, *ApJ*, 824, L7
 Gavazzi G., Pierini D., Boselli A., 1996, *A&A*, 312, 397
 Geréb K., Morganti R., Oosterloo T. A., Guglielmino G., Prandoni I., 2013, *A&A*, 558, 1
 Geréb K., Morganti R., Oosterloo T. A., Hoppmann L., Staveley-Smith L., 2015, *A&A*, 580, A43
 Haynes M. P. et al., 2011, *AJ*, 142, 170
 Hess K., Luber N., Fernandez X., Gim H., Momjian E., van Gorkom J., CHILES, 2018, *MNRAS*, 484, 2234
 Holwerda B. W., Blyth S.-L., Baker a. J., 2011, in Tuffs R. J., Popescu C. C., eds, Proc. IAU Symp. 7, The Spectral Energy Distribution of Galaxies. Preston, UK, p. 496
 Huang S., Haynes M. P., Giovanelli R., Brinchmann J., 2012, *ApJ*, 756, 113
 Jaffé Y. L. et al., 2016, *MNRAS*, 461, 1202
 Johnston S. et al., 2008, *Exp. Astron.*, 22, 151
 Jones M. G., Haynes M. P., Giovanelli R., Papastergis E., 2016, *MNRAS*, 455, 1574
 Kannappan S. J., 2004, *ApJ*, 611, L89
 Kannappan S. J. et al., 2013, *ApJ*, 777, 42
 Kott P. S., 2001, *J. Off. Stat.*, 17, 521
 Lah P. et al., 2007, *MNRAS*, 376, 1357
 Lah P. et al., 2009, *MNRAS*, 399, 1447
 Maddox N., Hess K. M., Blyth S.-L., Jarvis M. J., 2013, *MNRAS*, 433, 2613
 Maddox N., Hess K. M., Obreschkow D., Jarvis M. J., Blyth S. L., 2015, *MNRAS*, 447, 1610
 Matthews L. D., van Driel W., 2000, *A&AS*, 143, 421
 McGaugh S. S., de Blok W. J. G., 1997, *ApJ*, 481, 689
 Meyer S. A., Meyer M., Obreschkow D., Staveley-Smith L., 2016, *MNRAS*, 455, 3136
 Obreschkow D., Rawlings S., Road K., 2009a, *MNRAS*, 394, 1857
 Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009b, *ApJ*, 698, 1467
 Obreschkow D., Klockner H.-R., Heywood I., Levrier F., Rawlings S., 2009c, *ApJ*, 703, 1890
 Ramatsoku M. et al., 2016, *MNRAS*, 460, 923
 Rhee J., Zwaan M. A., Briggs F. H., Chengalur J. N., Lah P., Oosterloo T., van der Hulst T., 2013, *MNRAS*, 435, 2693
 Rhee J., Lah P., Chengalur J. N., Briggs F. H., Colless M., 2016, *MNRAS*, 460, 2675
 Roberts M. S., Haynes M. P., 1994, *ARA&A*, 32, 115
 Saintonge A., 2007, *AJ*, 133, 2087
 Serra P. et al., 2015, *MNRAS*, 448, 1922
 Stoughton C. et al., 2002, *AJ*, 123, 485
 Tully R. B., Fisher J. R., 1977, *A&A*, 54, 661
 van Driel W. et al., 2016, *A&A*, 595, 71
 Verheijen M., van Gorkom J. H., Szomoru A., Dwarakanath K. S., Poggianti B. M., Schiminovich D., 2007, *ApJ*, 668, L9
 Verheijen M. A. W., Oosterloo T. A., Van Cappellen W. A., Bakker L., Ivashina M. V., Van Der Hulst J. M., 2008, in Robert Minchin, Emmanuel Momjian, eds, Evolution of galaxies through the neutral hydrogen window, American Institute of Physics, Melville, p. 265
 Werner M. W. et al., 2004, *ApJS*, 154, 1
 Westmeier T., Jurek R., Obreschkow D., Koribalski B. S., Staveley-Smith L., 2014, *MNRAS*, 438, 1176
 Wieringa M. H., de Bruyn A. G., Katgert P., 1992, *A&A*, 256, 331
 York D. G. et al., 2000, *AJ*, 120, 1579
 Zwaan M., van Dokkum P. G., Verheijen M. A. W., Briggs F. H., 2001, in John E. Hibbard, Michael Rupen, Jacqueline H. van Gorkom, eds, Gas and Galaxy Evolution, ASP Conference Proceedings, Vol. 240. Astronomical Society of the Pacific, San Francisco, p. 640

SUPPORTING INFORMATION

Supplementary data are available at [MNRAS](https://academic.oup.com/mnras/article/487/4/4901/5511788) online.

Table E1. Data for the H I non-detections in the stacking sample.
Figure E1. Optical images and H I spectra for all 323 galaxies classified as H I non-detections from the final NIBLES stacking sample (see Section 3.3.2).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: DETAILS OF HISS**A1 Running HISS**

HISS can be downloaded from <https://github.com/healytwin1/HISS>. Below is a summary of the README file which details the software requirements as well as how to install and run HISS from the command line.

An example of HISS run-time: to stack the 1000 NIBLES spectra with 1000 jackknife iterations on a standard laptop with 8 GB of RAM, took ~4–5 h.

FIRST TIME USERS: It is recommended to use the graphical interface to populate a config file. The graphical interface is called by the ‘hiss’ executable. To use HISS from the command-line, enter the following command:

```
python pipeline.py [-h] [-f <filepath + filename>] [-d] [-p] [-s] [-c] [-l]
```

There are a number of different options that can be used to run this package:

optional arguments:

-h, --help	show this help message and exit.
-d, --display	Option to display progress window during the stacking process.
-p, --saveprogress	Option to save progress window during the stacking process.
-s, --suppress	Use this flag to suppress all output windows. Note that [suppress] and [progress] cannot be used simultaneously.
-c, --clean	Use[clean] for testing purposes and for stacking noiseless spectra as this option will bypass any noise-related functions and actions.
-l, --latex	This option enables to the use of latex formatting in the plots.

A2 HISS configuration file

Below is the configuration file that is used by HISS. This file is filled by the interface shown in Fig. 2.

Config parameter	GUI parameter/section	Description
CatalogueFilename	Catalogue file	Path to catalogue file – this should be the full path
CatalogueColumnNumbers	Catalogue information	Column numbers for [Object ID, Filename, Redshift, Redshift Uncertainty, Stellar Mass, Other Data]. ‘Other Data’ refers to data used to bin the sample - if one of the columns is not needed leave the column entry as an empty string.
z_max	Redshift range	Maximum redshift value of sample.
z_min		Minimum redshift value of sample.
SpectrumFluxUnit	Input spectra block	Specify flux density units of input spectra. Values to use in config file: 1-Jy, 2-mJy, 3-uJy, 4-Jy/beam, 5-mJy/beam, 6- μ Jy/beam
SpectralAxisUnit		Values to use in config file: 1-Hz, 2-kHz, 3-MHz, 4-m s ⁻¹ , 5-km s ⁻¹
VelocityType		The choices here are optical/radio. The default is optical, and should be changed with care. This option will only appear in the GUI if a velocity spectral type is chosen.
ChannelWidth		This should be entered in the same units as the spectral axis. This parameter is used as a diagnostic to flag spectra which have vastly different channel widths than the rest of the sample.
SpectrumLocation	Spectra location	Full path to location of spectra.
FirstRowInSpectrum	Input spectra file info	First row of data in spectrum file (numbering starts at 1). This option allows for header information in the spectrum file to be skipped.
RowDelimiter	Data delimiter	Spectrum file data delimiter – if the delimiter is whitespace, use ‘none’.
SpectrumColumns	Spectrum data columns (axis/flux density)	Column numbers for [freq/vel, flux density].
StackFluxUnit	Stacked spectrum information	Choose unit in which to stack. Choices: 1-Jy, 2-M _☉ , 3-M _{H_I} /M _*
H0	Cosmology	Choice of H ₀ – default is 70 km s ⁻¹ Mpc ⁻¹
Om0		Choice of $\Omega_{m,0}$ – default is 0.3
OutputLocation	Output location	Full path to where results should be saved.
GalaxyWidth	Max galaxy velocity width of sample	This is the maximum velocity width in the sample. This value is used to mask the spectrum where emission is expected as well as uses to define the integration window. The value can be changed during the analysis process.
StackedSpectrumLength	Velocity width of spectrum	This defines the length of the spectrum. There is a hard limit that this value cannot be less than three times the maximum galaxy velocity width. Default is 2000 km s ⁻¹ .
WeightOption	Select weighting function	1-w = 1[default], 2-w = $\frac{1}{rms}$, 3-w = $\frac{1}{rms^2}$, 4-w = $\frac{1}{(D_L^2 rms)^2}$
NumberCatalogueObjects	Stack entire catalogue	If ‘y’ is chosen, need to entire the number of catalogue objects to stack. This will be set to the length of the catalogue if the chosen number is larger.
UncertaintyYN	Calculate uncertainties	This is a [y]es/[n]o option.
UncertaintyMethod	Uncertainty method	redshift – uses redshift uncertainties to determine uncertainty of stacked profile. dagjk – Delete-a-group-jackknife, determines the statistical uncertainty of stacked profile.
BinYN	Stack in different bins?	This is a [y]es/[n]o option.
BinInfo	Information on bins	Column name of data to bin, Bin width, Start of Bin range, End of Bin range

A3 Plots created by HISS

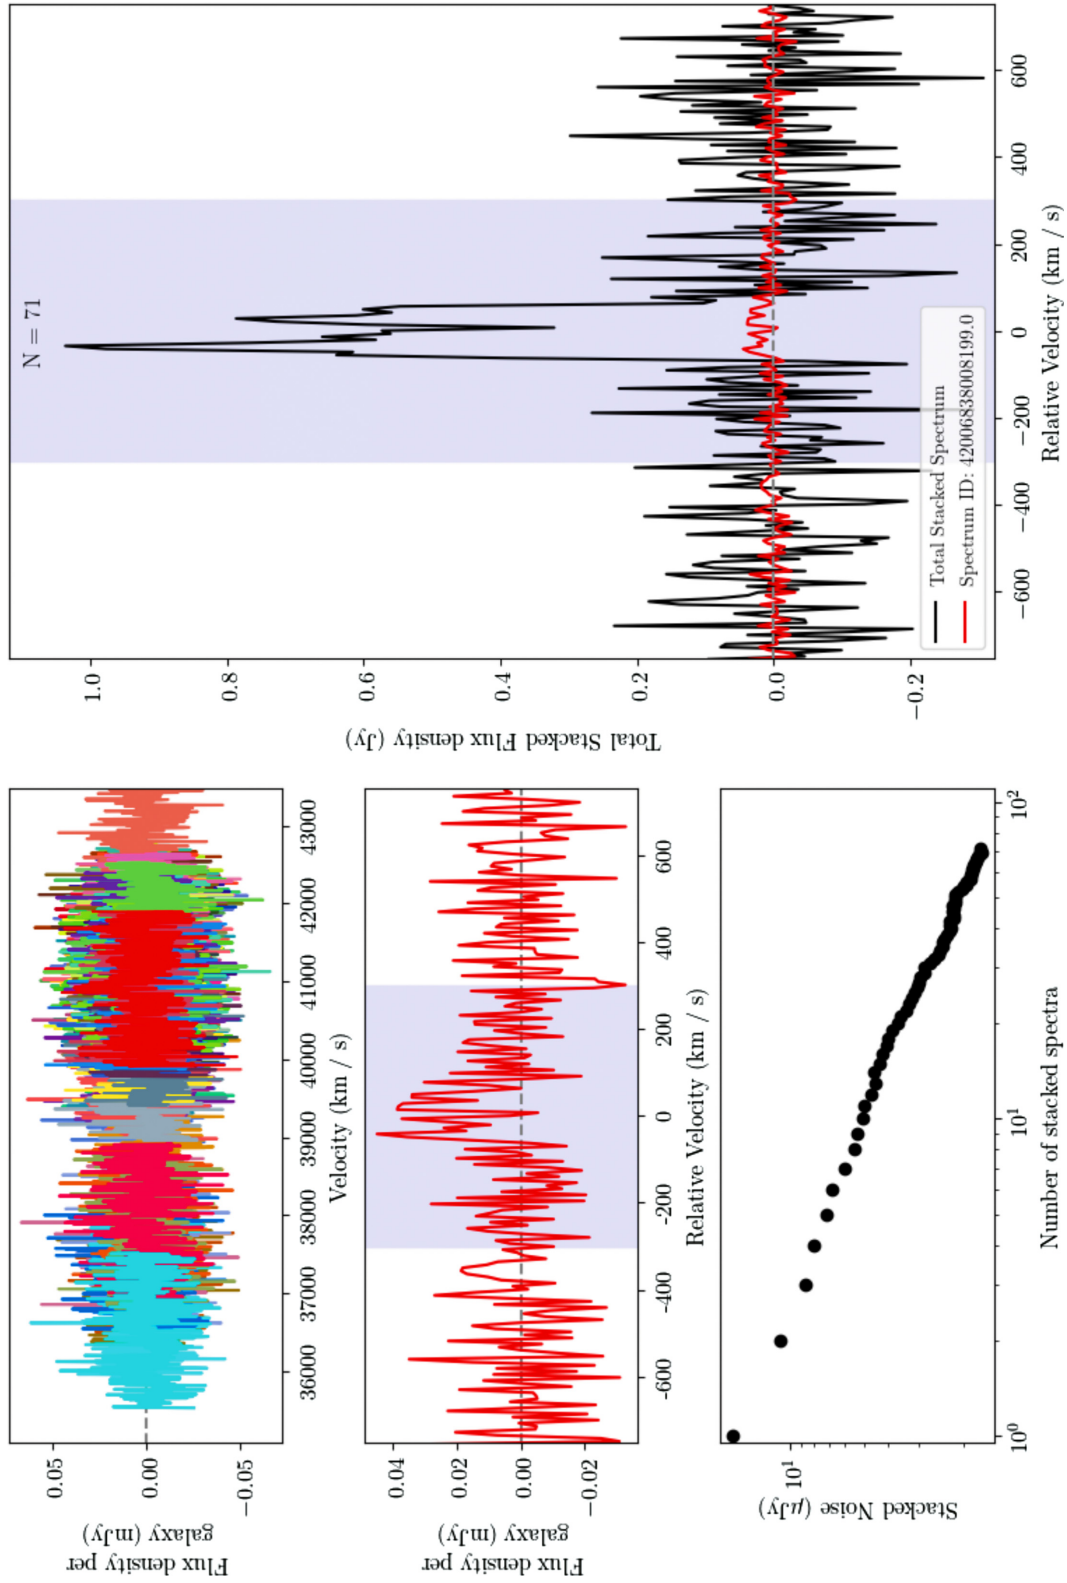


Figure A1. The top left panel shows where the input spectra lie in the observed frame. The middle left panel shows the now aligned spectrum in the velocity and flux density units in which it will be stacked. The bottom left panel shows how the average noise changes with the addition of each new spectrum. The large panel on the right shows the total stacked spectrum, in black, together with the incoming spectrum in blue.

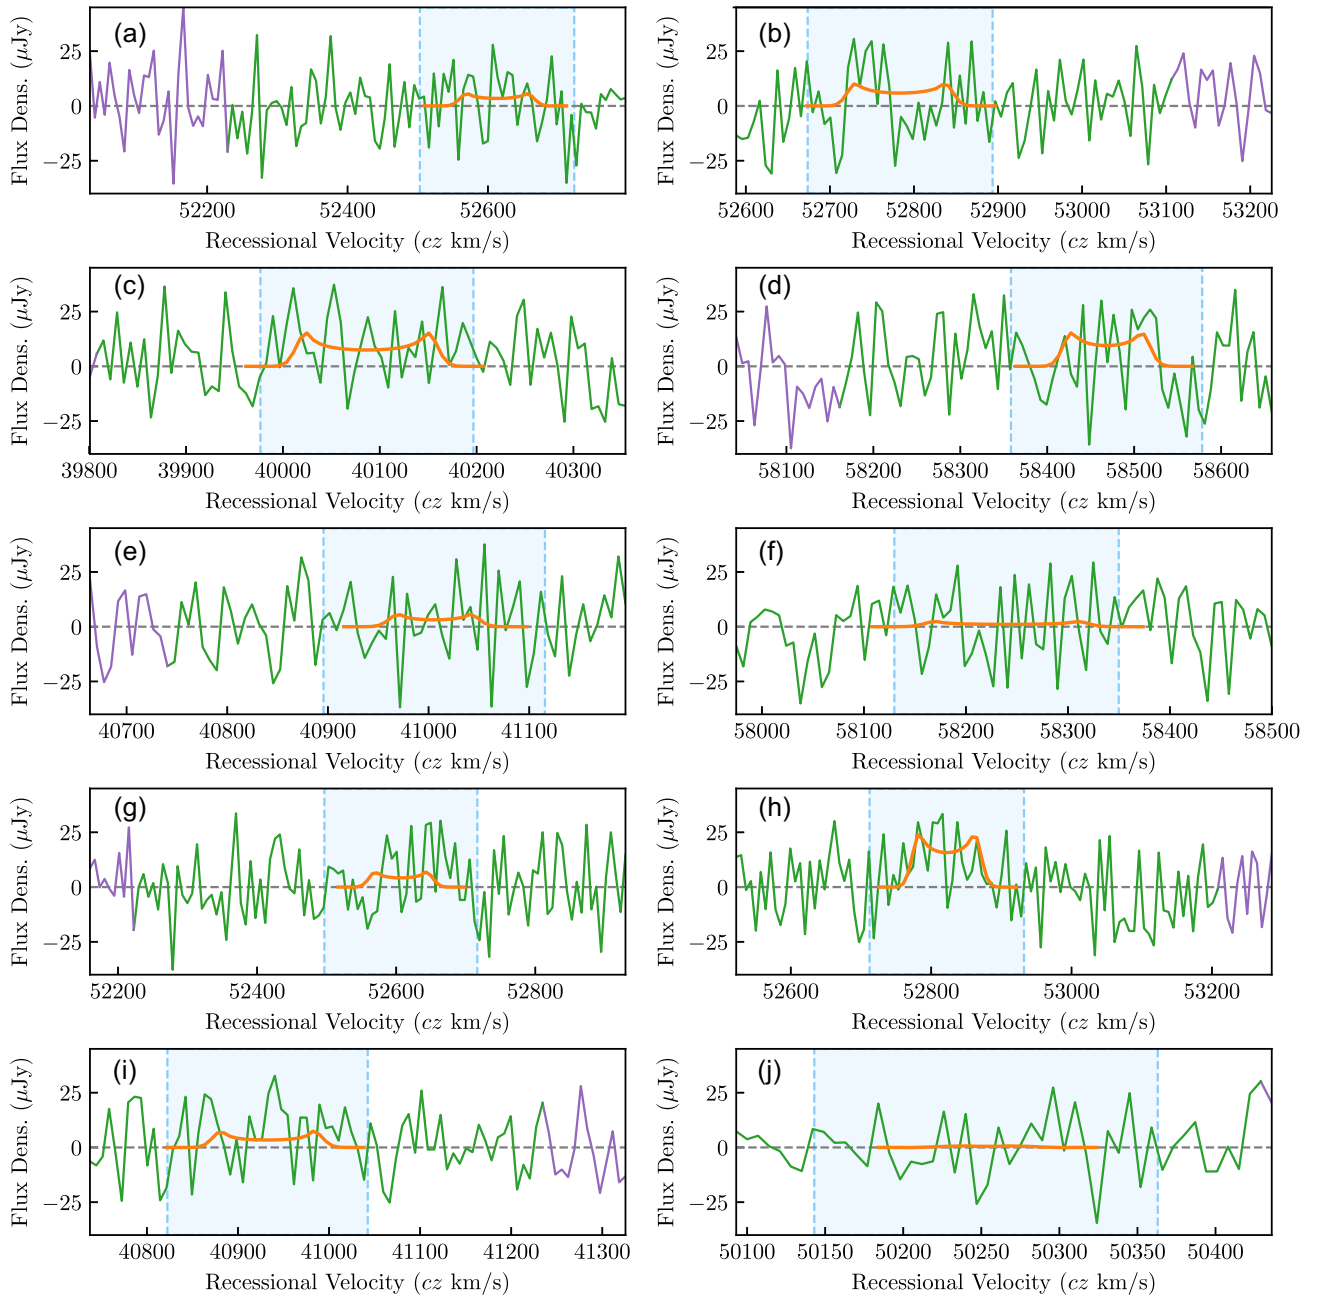


Figure A2. These 10 spectra are a selection of the 1000 simulated spectra used in the stacking example throughout the paper. The spectra are plotted using their native flux density and spectral axes. The simulated, noiseless spectra are shown in orange. Shown in green and purple are the spectra after adding simulated Gaussian noise. The green section of the spectrum shows the part that will be shifted to the centre of the array while the purple section is the part that will be wrapped around and appended to the other end of the spectrum in the shifting process. The light blue region shows the section of the profile where the galaxy emission is expected to be located. The width of this region is defined by the user and forms part of the galaxy mask. Panel (j) is an example of a spectrum that will be excluded during the stacking process as there are not enough channels outside of the galaxy window.

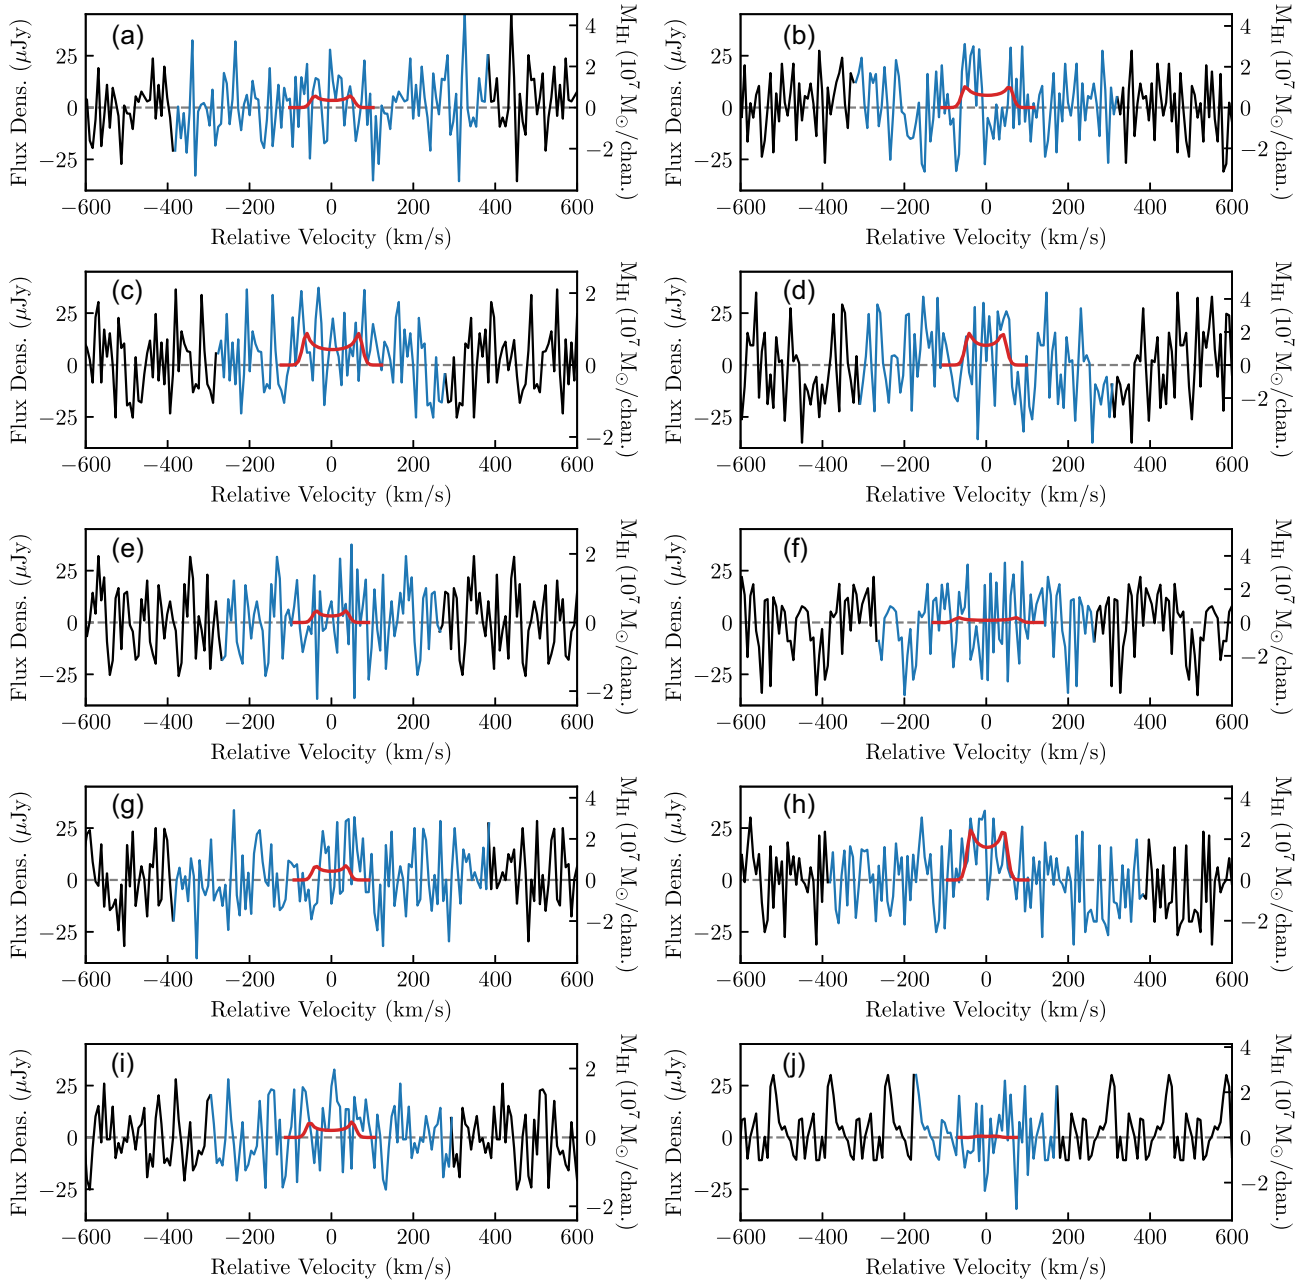


Figure A3. These are the same 10 simulated spectra as those plotted in Fig. A2. Shown here in blue are the spectra that have been converted to their rest frame and centred. The original noise-free versions of each of the spectra are now shown in red. Shown in black are the parts of the extended spectrum that have been filled using the flux from outside the galaxy window. The right y-axis shows that each of these spectra have also been converted to mass spectra.

Diagnostic Plot 1

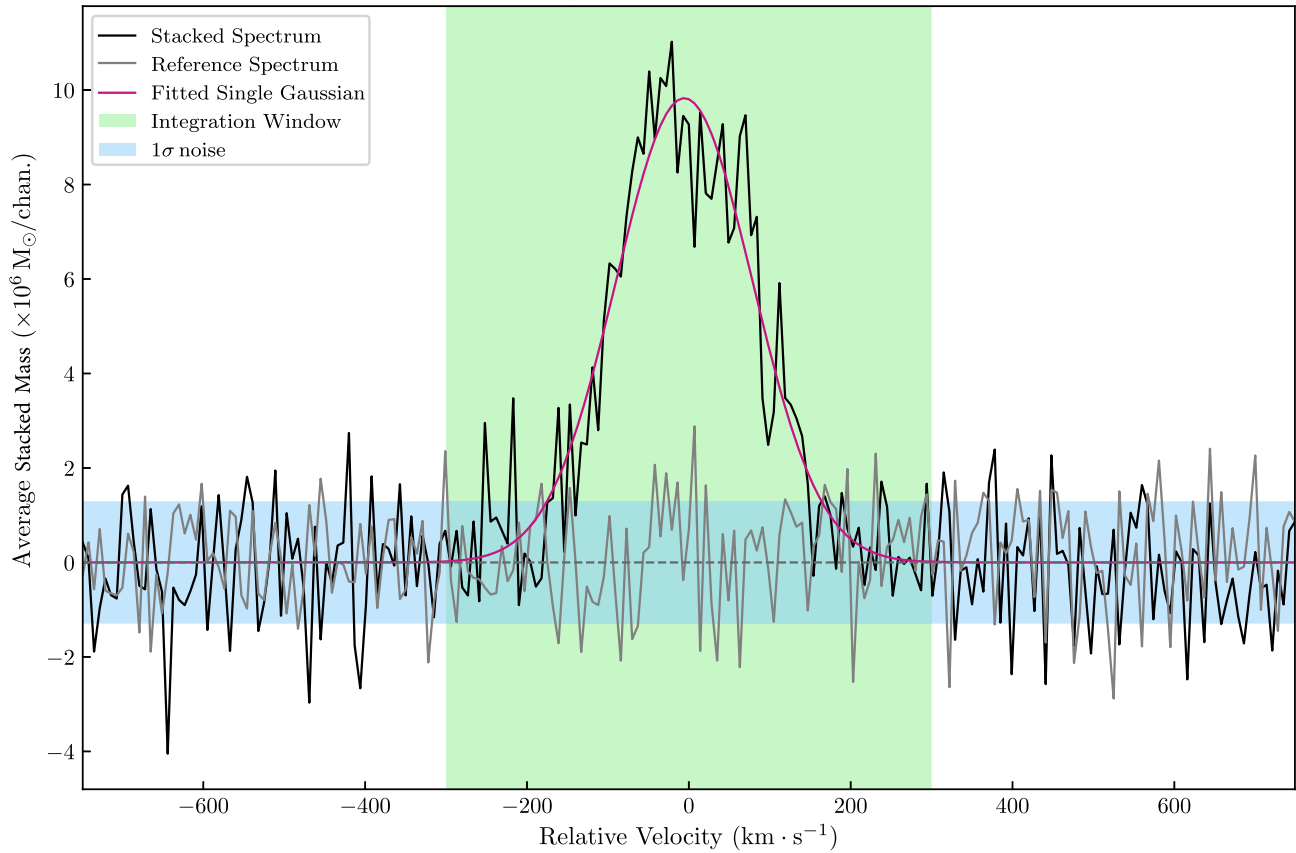


Figure A4. Diagnostic Plot 1: this is the first plot of the stacked spectrum displayed to the user. The stacked spectrum is shown in black in the units the user chose as the stacking units. This plot is shown to the user before any spectrum manipulation routines have been applied to the spectrum. The magenta curve shows the Single Gaussian fit to the spectrum to determine the significance of the detection. The grey spectrum is the reference spectrum. The green vertical window indicates the part of the spectrum over which the average quantity will be measured. The 1σ noise level is shown by the horizontal blue band.

Diagnostic Plot 3

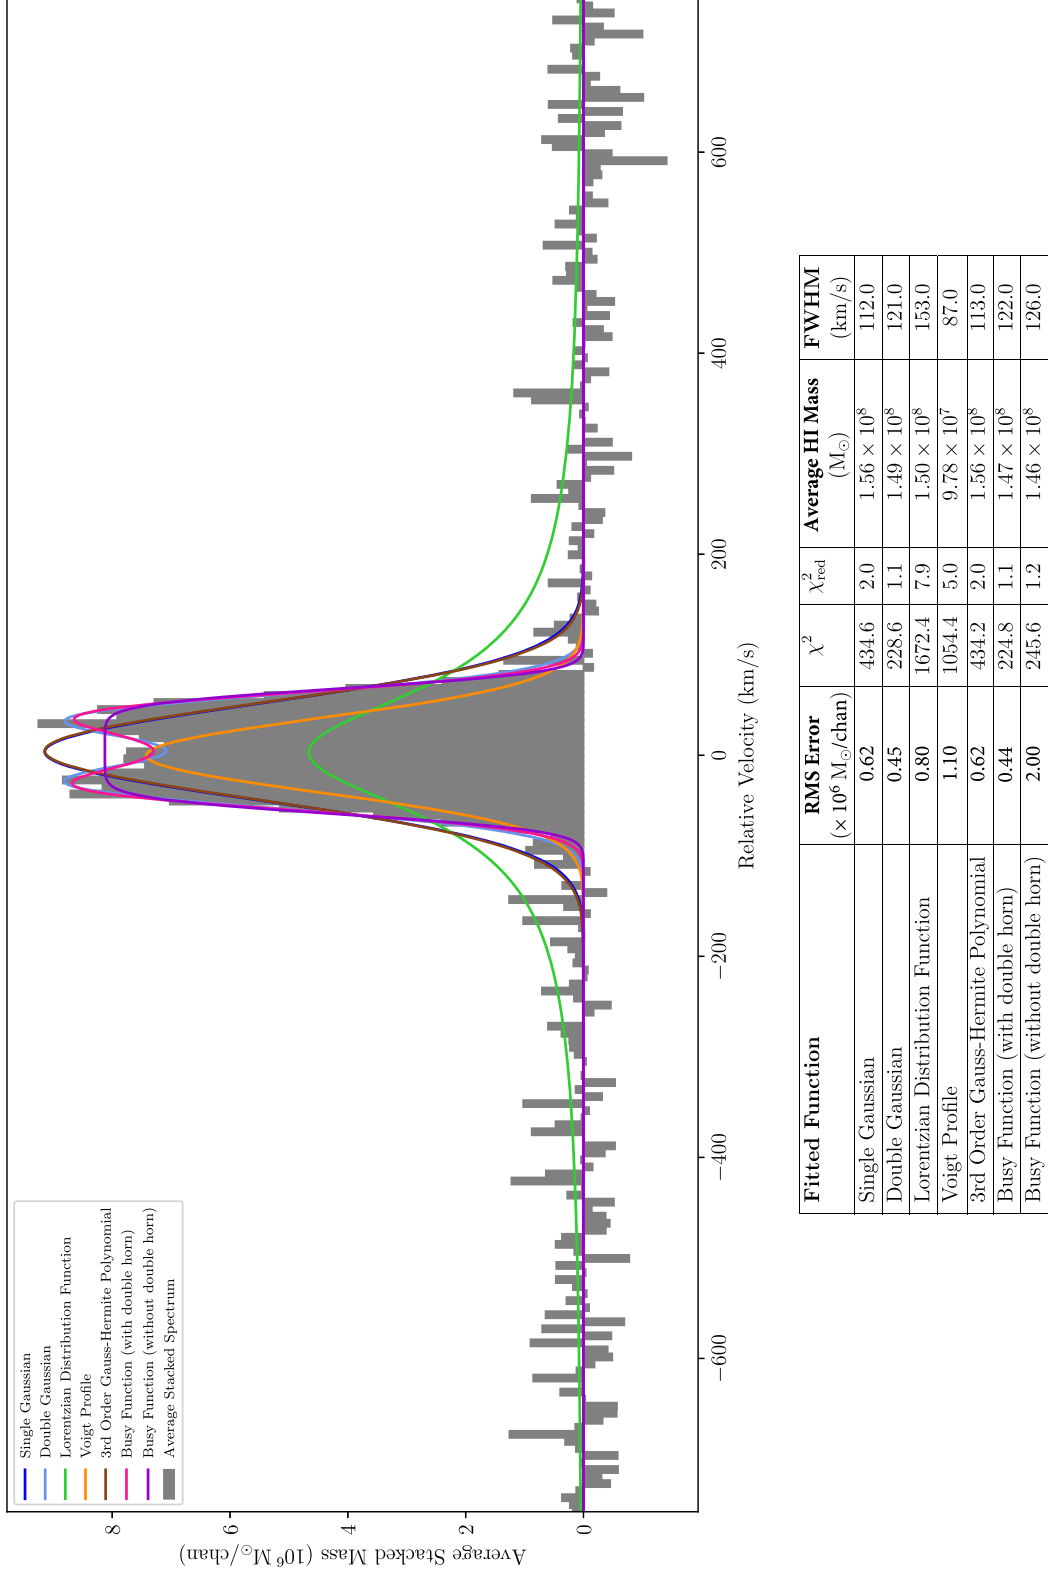


Figure A5. Diagnostic Plot 3: the top panel shows the seven different functions (in various colours) fitted to the stacked spectrum which is represented by the grey bars. The table in the bottom panel shows the measured average HI mass as well as three measures of the goodness of fit – here we use the standard χ^2 goodness of fit parameter and the $\chi^2_{\text{red}} = \chi^2/\text{dof}$, where dof is the number of degrees of freedom, and the last measure is the RMS error. This plot enables the user to inspect the quality of the fit to the data before committing to a selection of functions.

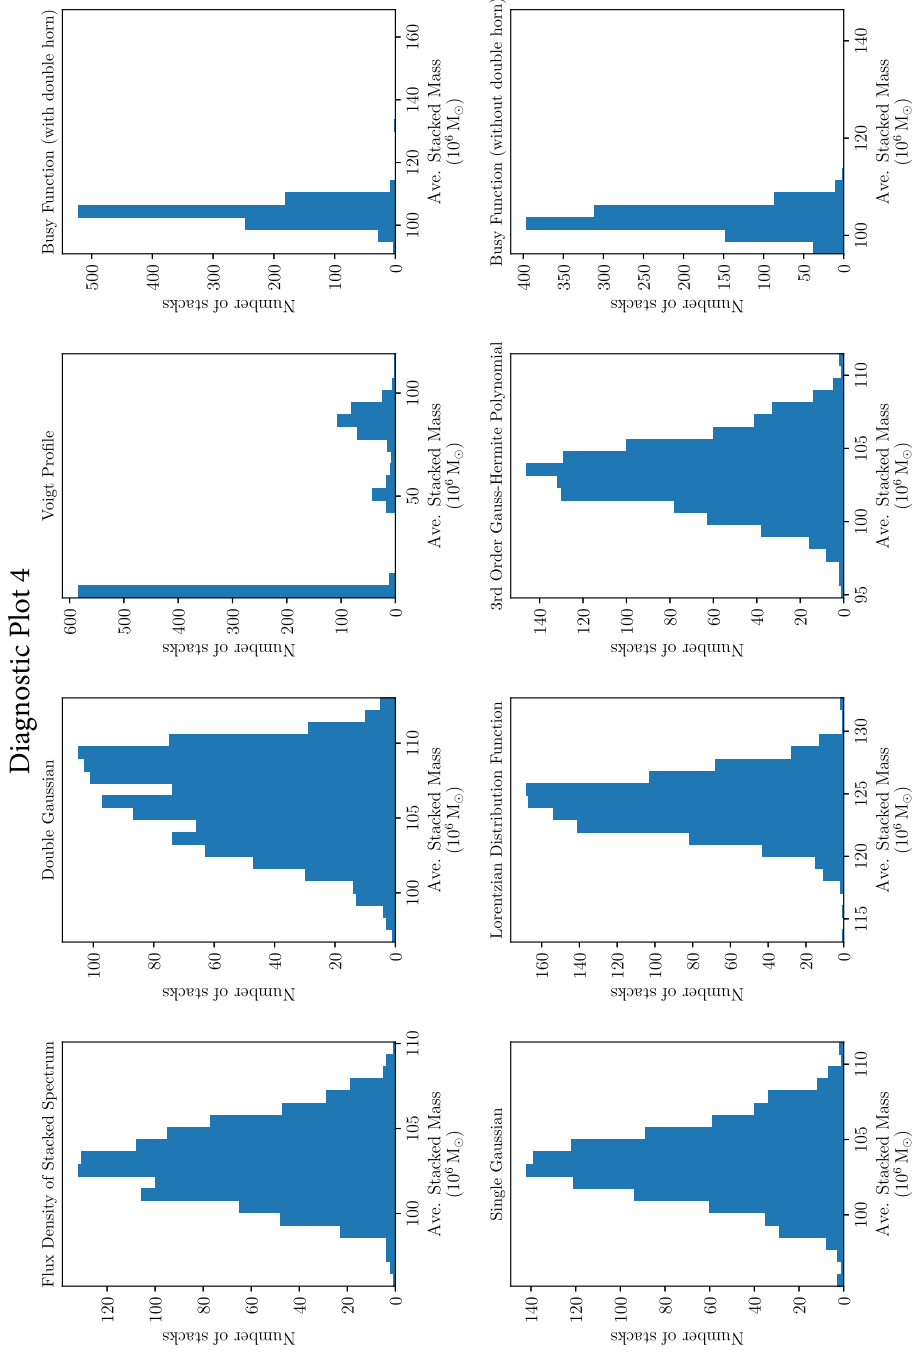



Figure A6. Diagnostic Plot 4; each panel shows a histogram of the average stacked mass measured from each profile fit for each of the 1000 stacking iterations (i.e. stacks). This plot provides a picture of how well each function was fit. It is from these histograms that the best estimate of the average quantity and associated error are calculated.



Number of galaxy profiles included in Stacked Spectrum: 100
 The peak signal-to-noise ratio: 20
 The significance of the peak: $> 8.2\sigma$ (p-value = 1.45×10^{-139})
 The ALFA signal-to-noise ratio: 70
 Integrated signal-to-noise ratio: 54

Fitted Function	Average HI Mass (M_{\odot})	Average HI Mass Uncertainty (M_{\odot})	RMS Error ($\times 10^6 M_{\odot}/\text{chan}$)	Fit ChiSquare	FWHM (km/s)
Single Gaussian	9.53×10^7	2.65×10^6	0.14	12	175.0
Average Mass of Stacked Spectrum	9.22×10^7	2.60×10^6	0	0	175.0

Figure A7. Screen shot of the output in full display mode. The top left panel shows the final average stacked spectrum with error bars in black, and the chosen fitted Single Gaussian overlaid in blue. The top right panel shows the stacked noise as a function of the number of profiles in grey; overlaid in green is the theoretical response to stacked noise: σ/\sqrt{N} . The bottom half of the figure contains the table of measured quantities (average HI mass and associated uncertainty), as well as the χ^2 and RMS Error calculated from the function fit to the stacked spectrum. The second entry in the table (Average Mass of Stacked Spectrum) shows the quantities measured directly from the stacked spectrum over the range specified by the user. The S/N statistics, as discussed in Section 2.4.1, for the stacked profile are given in the text above the table.

APPENDIX B: ACCOUNTING FOR KNOWN NEARBY CONTAMINATION

The Nançay telescope has a large beam which extends between 22 and 33 arcmin in the N–S direction depending on the declination of the source. The change in the N–S beam width can be approximated by the polynomial

$$y = (4.02 \times 10^{-9})x^5 - (6.90 \times 10^{-8})x^4 + (3.63 \times 10^{-6})x^2 + (5.45 \times 10^{-3})x + 21.95, \quad (\text{B1})$$

where y is the beam width in arcminutes and x is the declination in degrees (fig. 1; Matthews & van Driel 2000).

In this work, galaxies near the target source are considered as possible contaminants if

- (i) the secondary galaxy is located spatially within $1.2 \times$ N–S half-power beam width and $2 \times$ E–W half-power beam width, and
- (ii) the redshift of the secondary galaxy is within $\pm 300 \text{ km s}^{-1}$ of the target redshift.

A number of the target galaxies were found to have possible contaminants with very small angular separation distances. Upon inspection of the optical images for these sources, it was discovered that the possible contaminants were, in fact, part of the target galaxy. The SDSS pipeline has a known issue of de-blending large sources, particularly those with knots of star formation. In order not to overclassify the contaminant sources, a new criterion was therefore introduced to the search: the contaminant source must lie outside the edge of the target source, as defined as twice the r -band Petrosian radius (r_p), which according to Stoughton et al. (2002) is a good estimate of the size of the galaxy. Fig. B1 is an illustration of the method used to find any secondary contaminating sources.

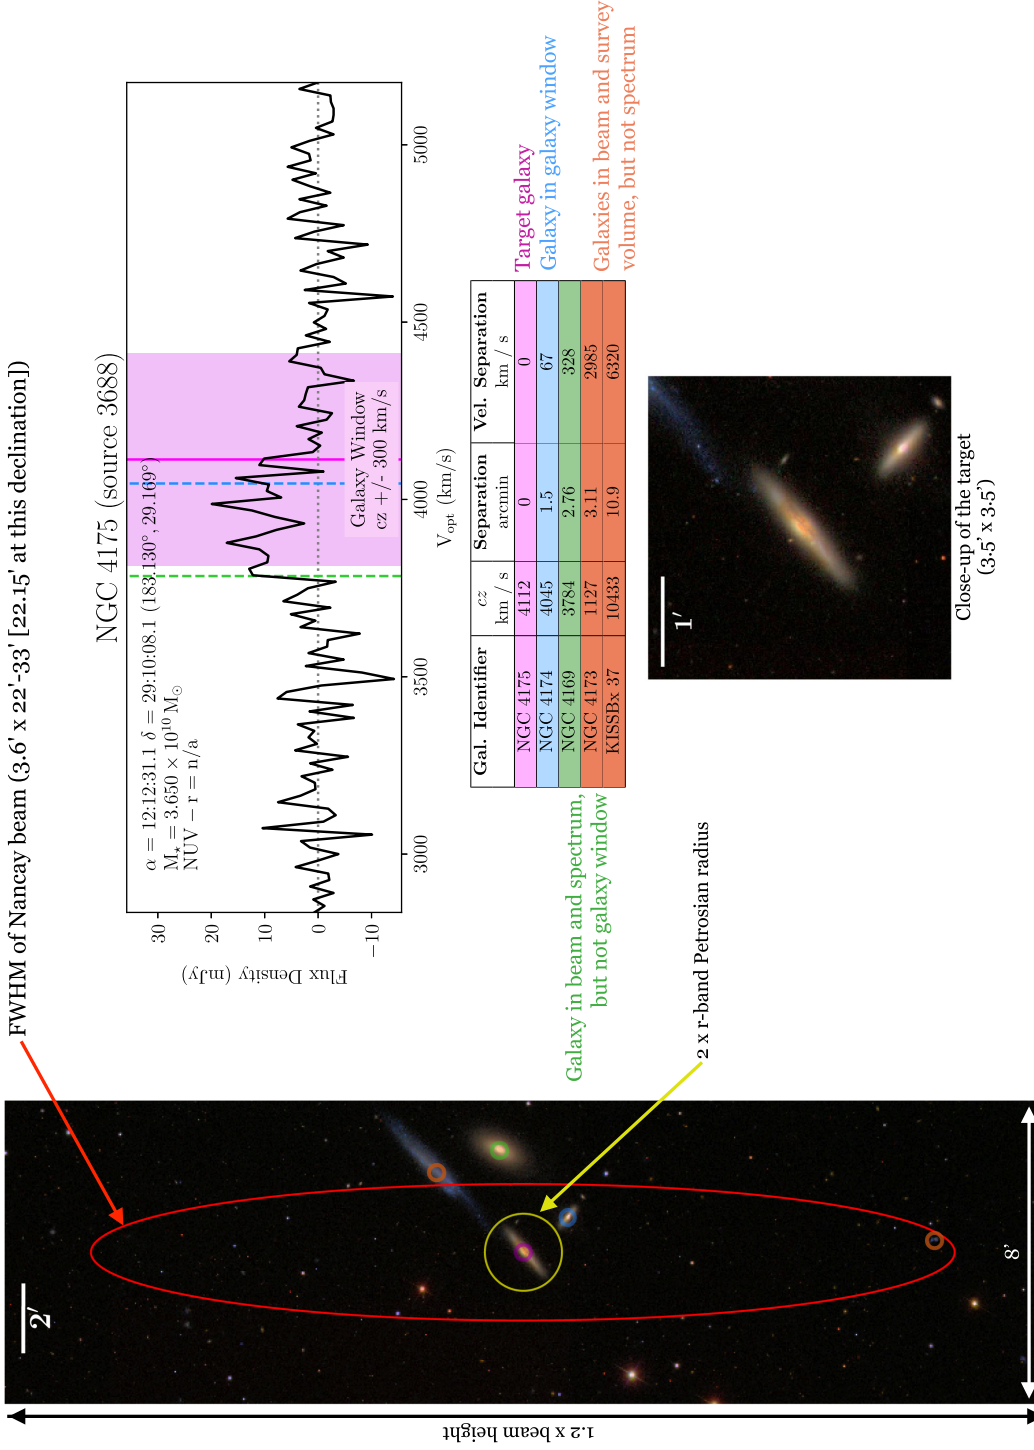


Figure B1. Illustration of the method used to find any secondary contaminating sources. On the left of the figure is an SDSS optical image that is 8 arcmin wide and $1.2 \times$ beam in declination centred on the coordinates of the target galaxy. Overlaid in red is the half-power size of the beam; the yellow circle has a radius of $2r_p$ which indicates the estimated size of the galaxy. The magenta open circular marker indicates the target source while the blue, green, and orange open circular markers indicate secondary targets in the entire survey velocity range ($900 \text{ km s}^{-1} < cz < 12000 \text{ km s}^{-1}$). Each of the targets highlighted in the large optical image are listed in the table to the right. The galaxies that are in the survey volume, but whose recessional velocities are not in the spectrum velocity range, are indicated by the orange circles. The green circles highlight galaxies that are within the velocity range covered by the spectrum – these sources are also marked on the spectrum by the green dashed vertical lines. The blue circled sources are the ones that are considered contamination sources. These sources have a redshift (highlighted by the dashed blue line on the spectrum) that is within $\pm 300 \text{ km s}^{-1}$ of the target redshift (indicated by the magenta line on the spectrum). The small optical image below the table is a $3.5 \text{ arcmin} \times 3.5 \text{ arcmin}$ close-up centred on the target source.

APPENDIX C: DEFINING A SAMPLE OF H I NON-DETECTIONS

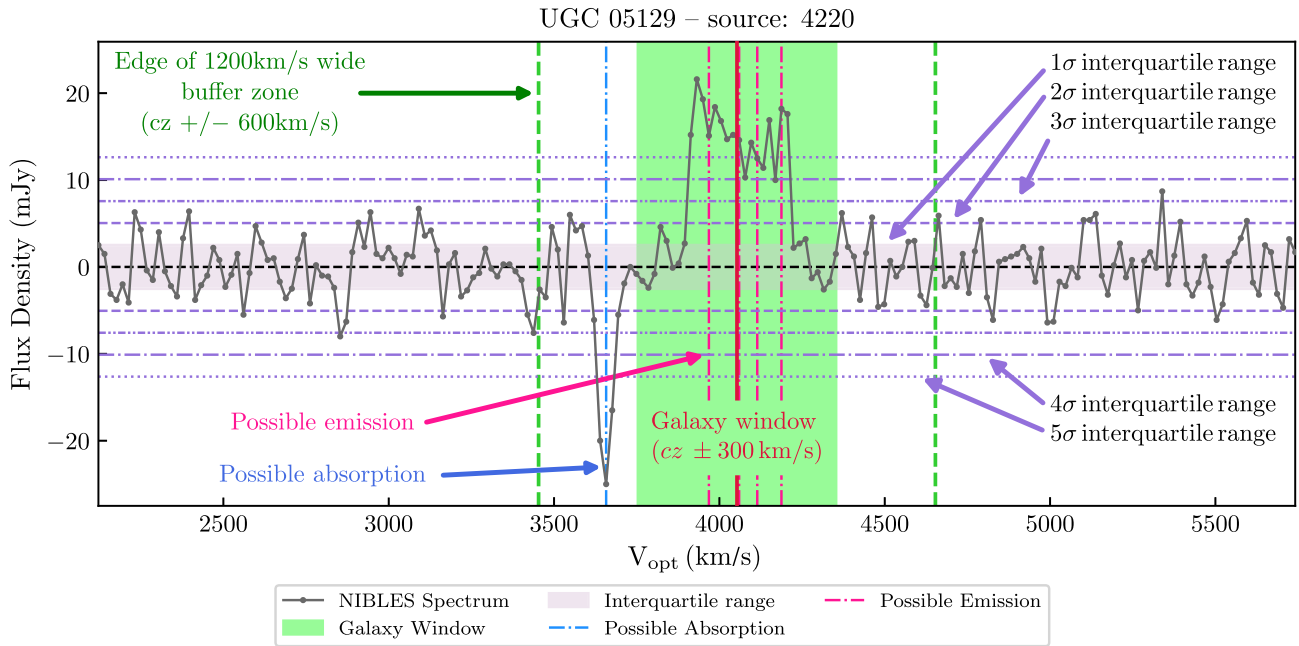


Figure C1. Graphic example of how the detection algorithm flags different features in a spectrum. The spectrum used in this example is classified as a detection. The criteria we used for defining H I detections is outlined in Section 3.3.2. Using these criteria we reproduced with 96 per cent overlap, the detection and non-detection classifications published in van Driel et al. (2016). The marginal detections in van Driel et al. (2016) were classified in an approximately one-third/two-thirds ratio as detections versus non-detections, respectively.

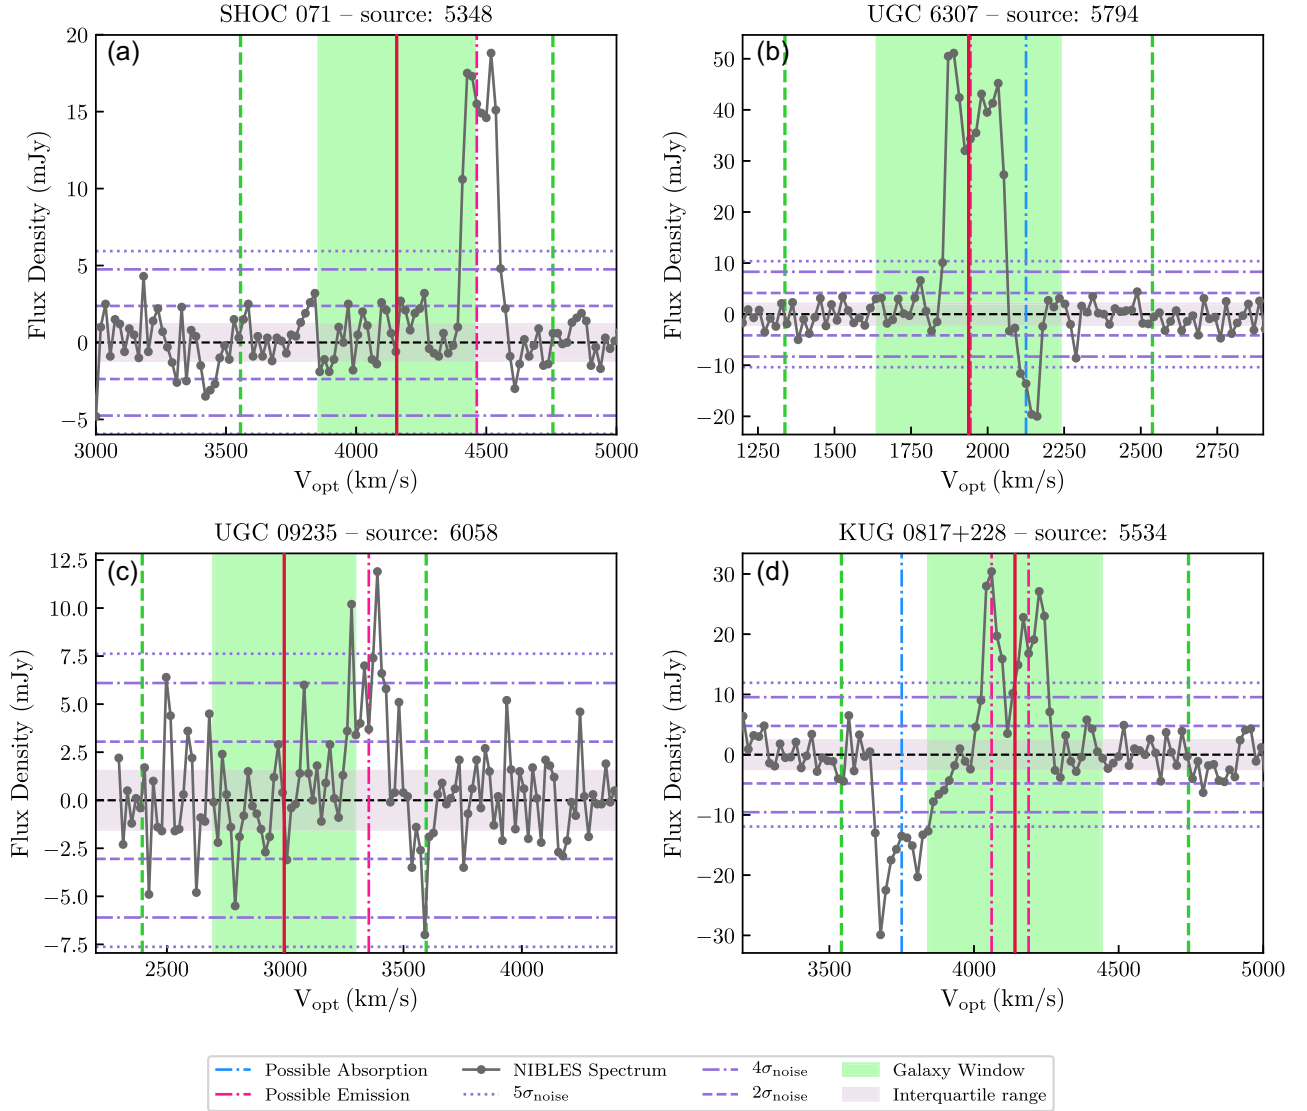


Figure C2. Sources flagged by the second run of the detection algorithm as containing emission/absorption not related to the target galaxy. All four sources were removed from the sample. Panels (b) and (d) show evidence of an H I detection in the off-source pointing. Panels (a) and (c) show clear H I detections that are not associated with sources previously identified in our contaminant source search. The vertical red line in each panel indicates the optical redshift of the target source.

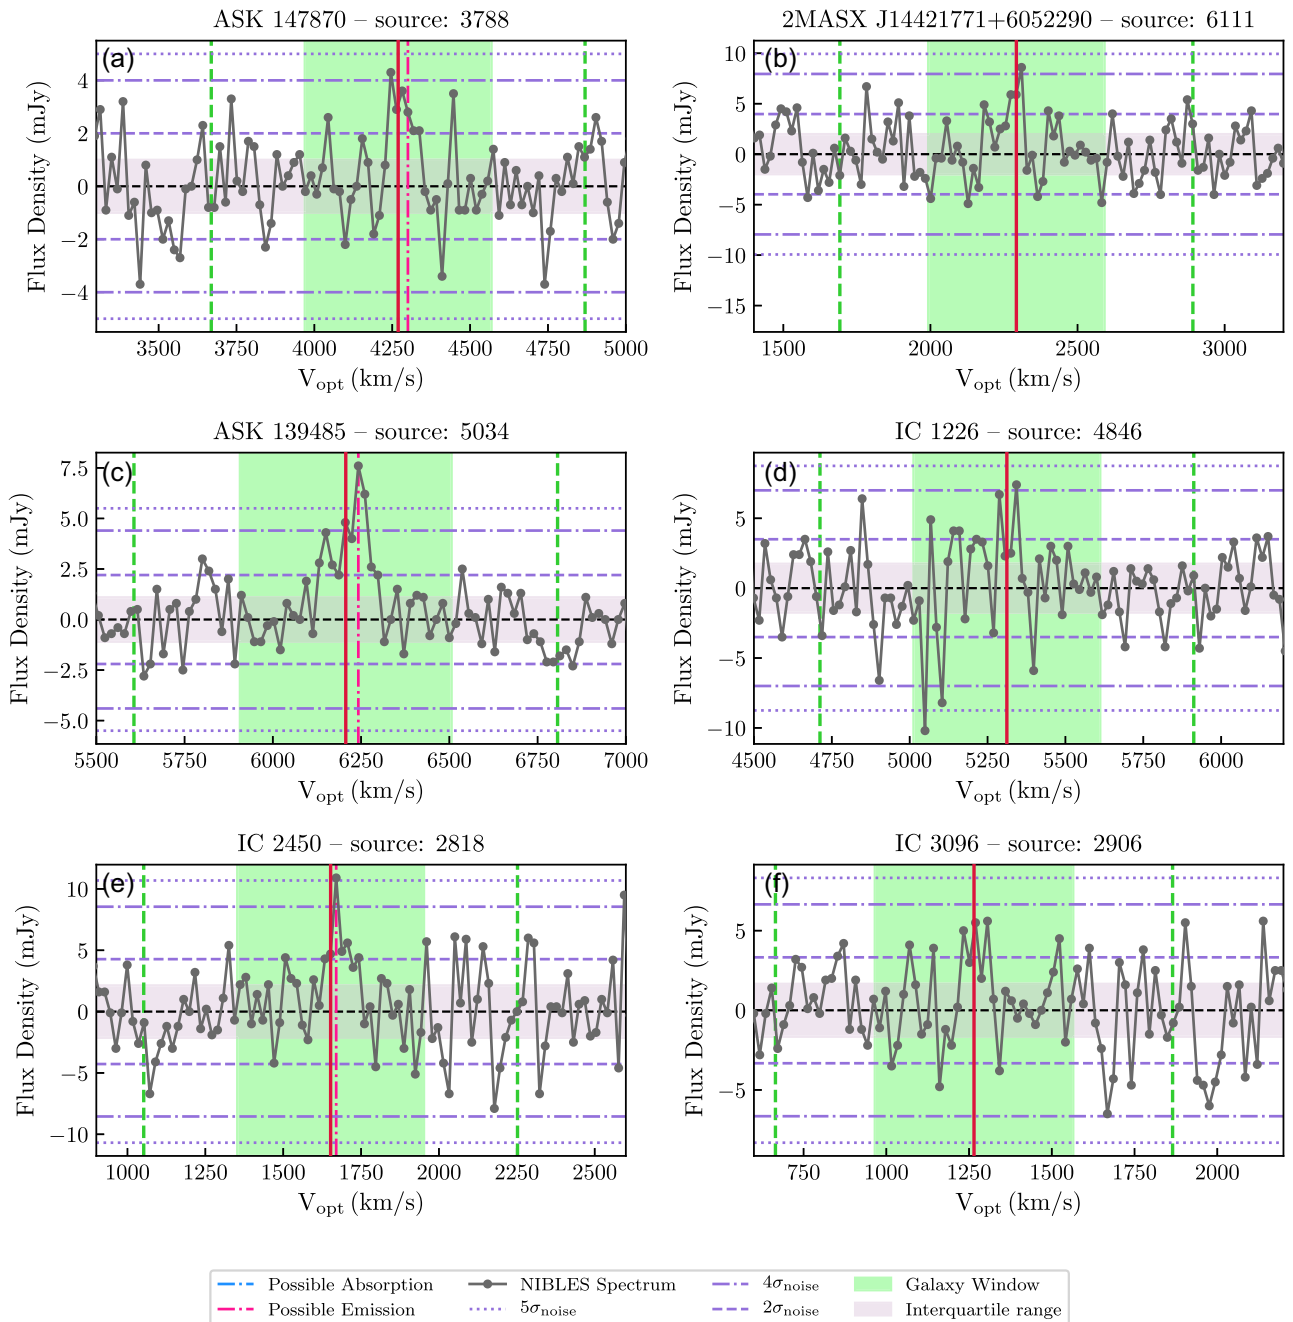


Figure C3. Examples of NIBLES spectra classified as H I detections (left column) and not classified as detections (right column) – i.e. non-detections. Panel (a) Classified as a detection with ≥ 5 consecutive resolution elements with flux density $> 2 \times \text{IQR}$. Panel (c) Classified as a detection with ≥ 1 consecutive resolution elements with flux density $> 5 \times \text{IQR}$. Panel (e) Classified as a detection with one resolution element with flux density $> 5 \times \text{IQR}$. The vertical red line in each panel indicates the optical redshift of the target source. Panels (a), (c), (e) are also identified by van Driel et al. (2016) as detections, however only panel (c) would be identified as a detection using the S/N criterion used by ALFALFA.

APPENDIX D: STACKING NIBLES GALAXIES

D1 Quantities measured from the stacked spectra

Table D1. Average quantities obtained from stacking the NIBLES galaxies in colour sub-samples and the full sample. The average H I mass values are calculated from the stacked spectra shown in Fig. 8. Gaussian significance values (i.e. $\sigma_{M_{\text{H I}}}/\sigma_{f_{\text{H I}}}$) of $>8.2\sigma$ have corresponding p -values that are too small to be numerically resolved by a computer, thus have been assigned the significance value of the smallest resolvable p -value.

	Full sample	Blue sample	Red sample	
All	1000	549	451	N
	$>8.2/>8.2$	$>8.2/>8.2$	$>8.2/>8.2$	$\sigma_{M_{\text{H I}}}/\sigma_{f_{\text{H I}}}$
	63.1/204.5	289.3/182.8	26.9/112.7	$S/N_{\text{ALFALFA}, M_{\text{H I}}}/S/N_{\text{ALFALFA}, f_{\text{H I}}}$
	$1.4939^{\pm 0.0024(\text{sys})}_{\pm 0.0572(\text{stat})} \times 10^9$	$1.47156^{\pm 0.0028(\text{sys})}_{\pm 0.0715(\text{stat})} \times 10^9$	$1.51545^{\pm 0.0042(\text{sys})}_{\pm 0.0946(\text{stat})} \times 10^9$	$\langle M_{\text{H I}} \rangle (M_{\odot})$
	1.795×10^{10}	5.641×10^9	3.280×10^{10}	$\langle M_{\star} \rangle (M_{\odot})$
$0.7776^{\pm 0.0020(\text{sys})}_{\pm 0.0358(\text{stat})}$	$1.2100^{\pm 0.0038(\text{sys})}_{\pm 0.0585(\text{stat})}$	$0.2533^{\pm 0.0011(\text{sys})}_{\pm 0.0321(\text{stat})}$	$\langle f_{\text{H I}} \rangle$	
Detections	677	436	241	N
	$>8.2/7$	$>8.2/>8.2$	$>8.2/>8.2$	$\sigma_{M_{\text{H I}}}/\sigma_{f_{\text{H I}}}$
	61.2/339.4	194.9/242.6	25.9/122.3	$S/N_{\text{ALFALFA}, M_{\text{H I}}}/S/N_{\text{ALFALFA}, f_{\text{H I}}}$
	$2.1231^{\pm 0.0030(\text{sys})}_{\pm 0.0825(\text{stat})} \times 10^9$	$1.5926^{\pm 0.0026(\text{sys})}_{\pm 0.0864(\text{stat})} \times 10^9$	$2.5806^{\pm 0.0080(\text{sys})}_{\pm 0.1728(\text{stat})} \times 10^9$	$\langle M_{\text{H I}} \rangle (M_{\odot})$
	1.519×10^{10}	5.837×10^9	3.226×10^{10}	$\langle M_{\star} \rangle (M_{\odot})$
$1.0802^{\pm 0.0021(\text{sys})}_{\pm 0.0487(\text{stat})}$	$1.3222^{\pm 0.0038(\text{sys})}_{\pm 0.0739(\text{stat})}$	$0.3409^{\pm 0.0017(\text{sys})}_{\pm 0.0402(\text{stat})}$	$\langle f_{\text{H I}} \rangle$	
Non-detections	323	113	210	N
	$>8.2/>8.2$	$>8.2/>8.2$	$>8.2/7.7$	$\sigma_{M_{\text{H I}}}/\sigma_{f_{\text{H I}}}$
	21.1/32.5	18.7/27.6	8.6/8.8	$S/N_{\text{ALFALFA}, M_{\text{H I}}}/S/N_{\text{ALFALFA}, f_{\text{H I}}}$
	$1.5183^{\pm 0.0409(\text{sys})}_{\pm 0.3251(\text{stat})} \times 10^8$	$2.4889^{\pm 0.0589(\text{sys})}_{\pm 0.4674(\text{stat})} \times 10^8$	$1.1879^{\pm 0.0577(\text{sys})}_{\pm 0.4310(\text{stat})} \times 10^8$	$\langle M_{\text{H I}} \rangle (M_{\odot})$
	2.536×10^{10}	4.428×10^9	3.351×10^{10}	$\langle M_{\star} \rangle (M_{\odot})$
$0.1408^{\pm 0.0039(\text{sys})}_{\pm 0.0277(\text{stat})}$	$0.3780^{\pm 0.0098(\text{sys})}_{\pm 0.0710(\text{stat})}$	$0.0349^{\pm 0.0013(\text{sys})}_{\pm 0.0078(\text{stat})}$	$\langle f_{\text{H I}} \rangle$	

Table D2. Average quantities and statistical errors obtained from HISS for Fig. 9 for the low stellar mass sample ($10^8 < M_{\star} (M_{\odot}) < 10^{9.5}$). The results with Gaussian significance values of >8.2 have corresponding p -values that are too small to be numerically resolved by a computer, thus have been assigned the significance value of the smallest resolvable p -value.

	N	$\langle M_{\text{H I}} \rangle (M_{\odot})$	$\sigma_{M_{\text{H I}}}$	$S/N_{\text{ALFALFA}, M_{\text{H I}}}$	$\langle f_{\text{H I}} \rangle$	$\sigma_{f_{\text{H I}}}$	$S/N_{\text{ALFALFA}, f_{\text{H I}}}$
Full catalogue (all)	521	$7.46 \pm 0.33 \times 10^8$	>8.2	277.7	1.372 ± 0.069	>8.2	205.7
Full catalogue (blue)	401	$8.00 \pm 0.39 \times 10^8$	>8.2	232.8	1.563 ± 0.074	>8.2	190.5
Full catalogue (red)	120	$5.74 \pm 0.63 \times 10^8$	>8.2	196.0	0.740 ± 0.107	>8.2	127.3
Non-detections (all)	156	$12.69 \pm 2.47 \times 10^7$	>8.2	35.1	0.299 ± 0.050	>8.2	33.7
Non-detections (blue)	97	$18.25 \pm 3.95 \times 10^7$	>8.2	29.5	0.437 ± 0.077	>8.2	28.1
Non-detections (red)	59	$4.09 \pm 1.08 \times 10^7$	6	11.1	0.095 ± 0.027	5.9	7.5

Table D3. Same as Table D2 but for the high stellar mass sample ($10^{9.5} < M_{\star} (M_{\odot}) < 10^{12}$).

	N	$\langle M_{\text{H I}} \rangle (M_{\odot})$	$\sigma_{M_{\text{H I}}}$	$S/N_{\text{ALFALFA}, M_{\text{H I}}}$	$\langle f_{\text{H I}} \rangle$	$\sigma_{f_{\text{H I}}}$	$S/N_{\text{ALFALFA}, f_{\text{H I}}}$
Full catalogue (all)	479	$2.25 \pm 0.11 \times 10^9$	>8.2	54.2	0.132 ± 0.007	>8.2	172.6
Full catalogue (blue)	148	$3.26 \pm 0.22 \times 10^9$	>8.2	171.6	0.256 ± 0.014	>8.2	203.0
Full catalogue (red)	331	$1.87 \pm 0.13 \times 10^9$	>8.2	33.3	0.077 ± 0.005	>8.2	93.9
Non-detections (all)	174	$1.89 \pm 0.60 \times 10^8$	>8.2	16.6	0.013 ± 0.003	>8.2	19.0
Non-detections (blue)	17	$6.58 \pm 1.86 \times 10^8$	8	12.1	0.057 ± 0.015	4.6	10.8
Non-detections (red)	157	$1.45 \pm 0.54 \times 10^8$	7.9	8.0	0.010 ± 0.003	>8.2	16.0

D2 Stacked profiles for Figs 9 and 10

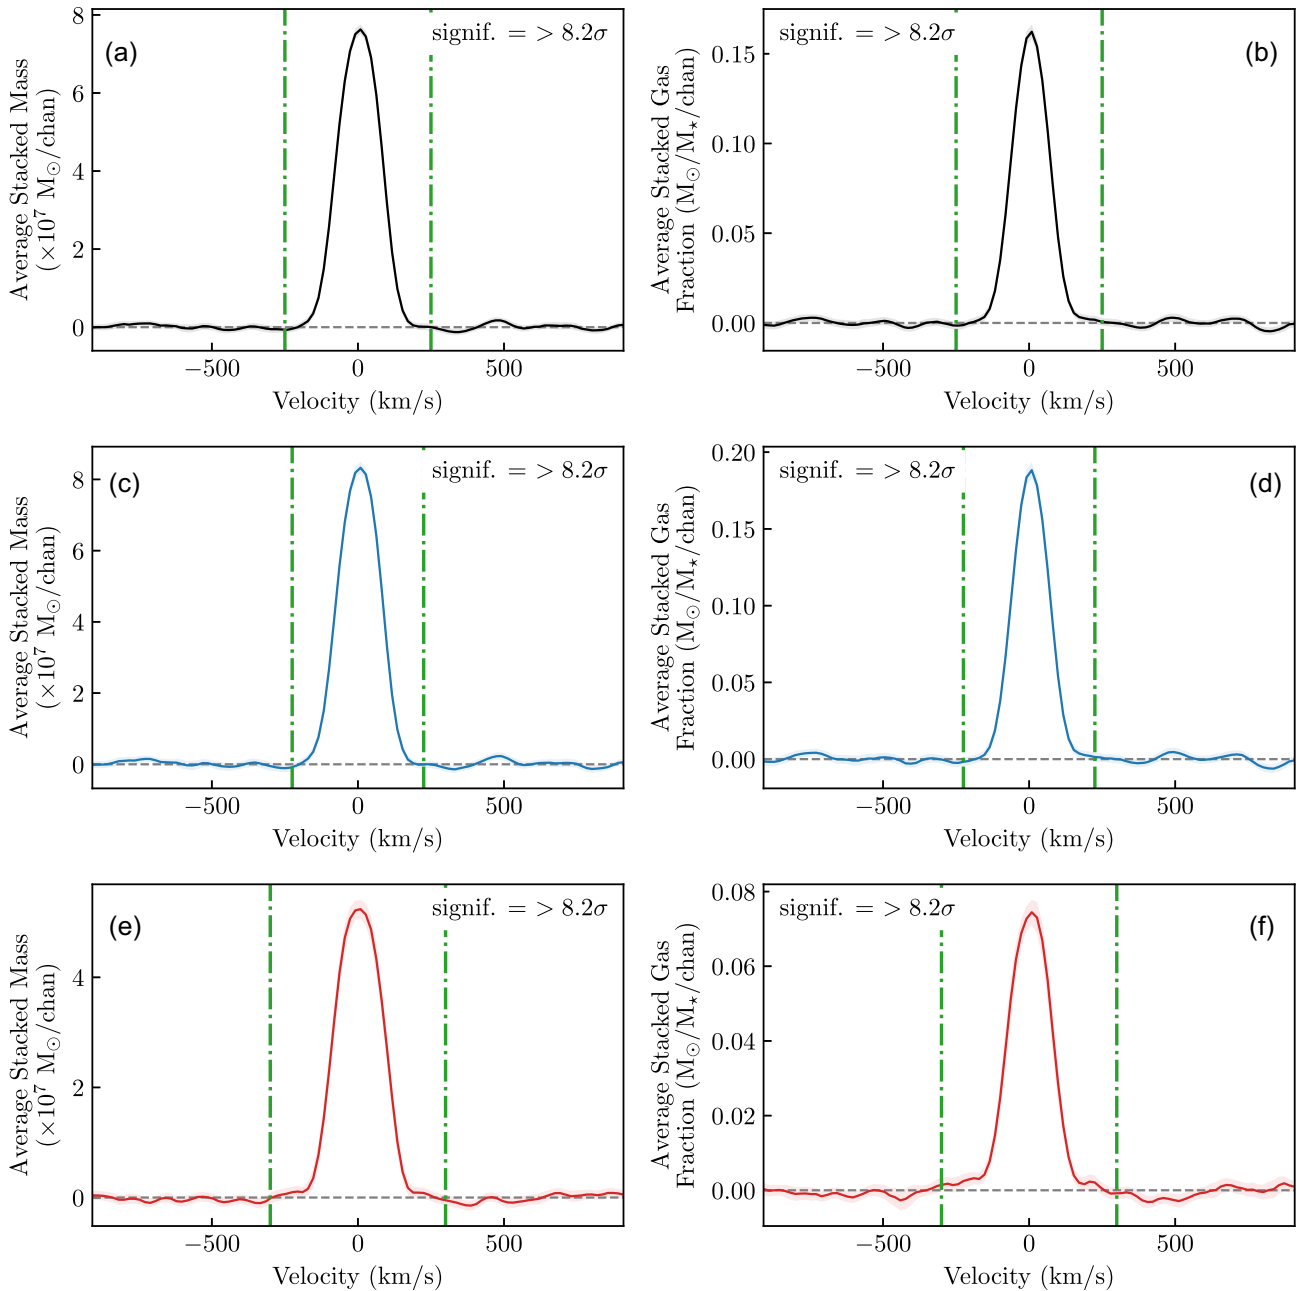


Figure D1. Stacked profiles for the full sample of H I detections and non-detections with $10^8 < M_{*}(M_{\odot}) < 10^{9.5}$. The left column (panels a, c, and e) shows the stacked profiles from which the average H I mass ($\langle M_{\text{H I}} \rangle$) is measured and the right column (panels b, d, and f) shows the stacked profiles from which the average gas fraction ($\langle f_{\text{H I}} \rangle$) is measured. The stacked profiles for all sources in this sample are shown in the top row. The middle row shows the red sub-sample, and the bottom row shows the blue sub-sample. The shading around each spectrum indicates the uncertainty in each channel calculated based on the redshift errors. The green dot-dashed lines in every panel mark the section of the stacked profiles used to measure the $\langle M_{\text{H I}} \rangle$ and $\langle f_{\text{H I}} \rangle$. These spectra correspond to the measurements of $\langle M_{\text{H I}} \rangle$ and $\langle f_{\text{H I}} \rangle$ presented in Fig. 9. The measurements made from the above spectra can be found in Table D2.

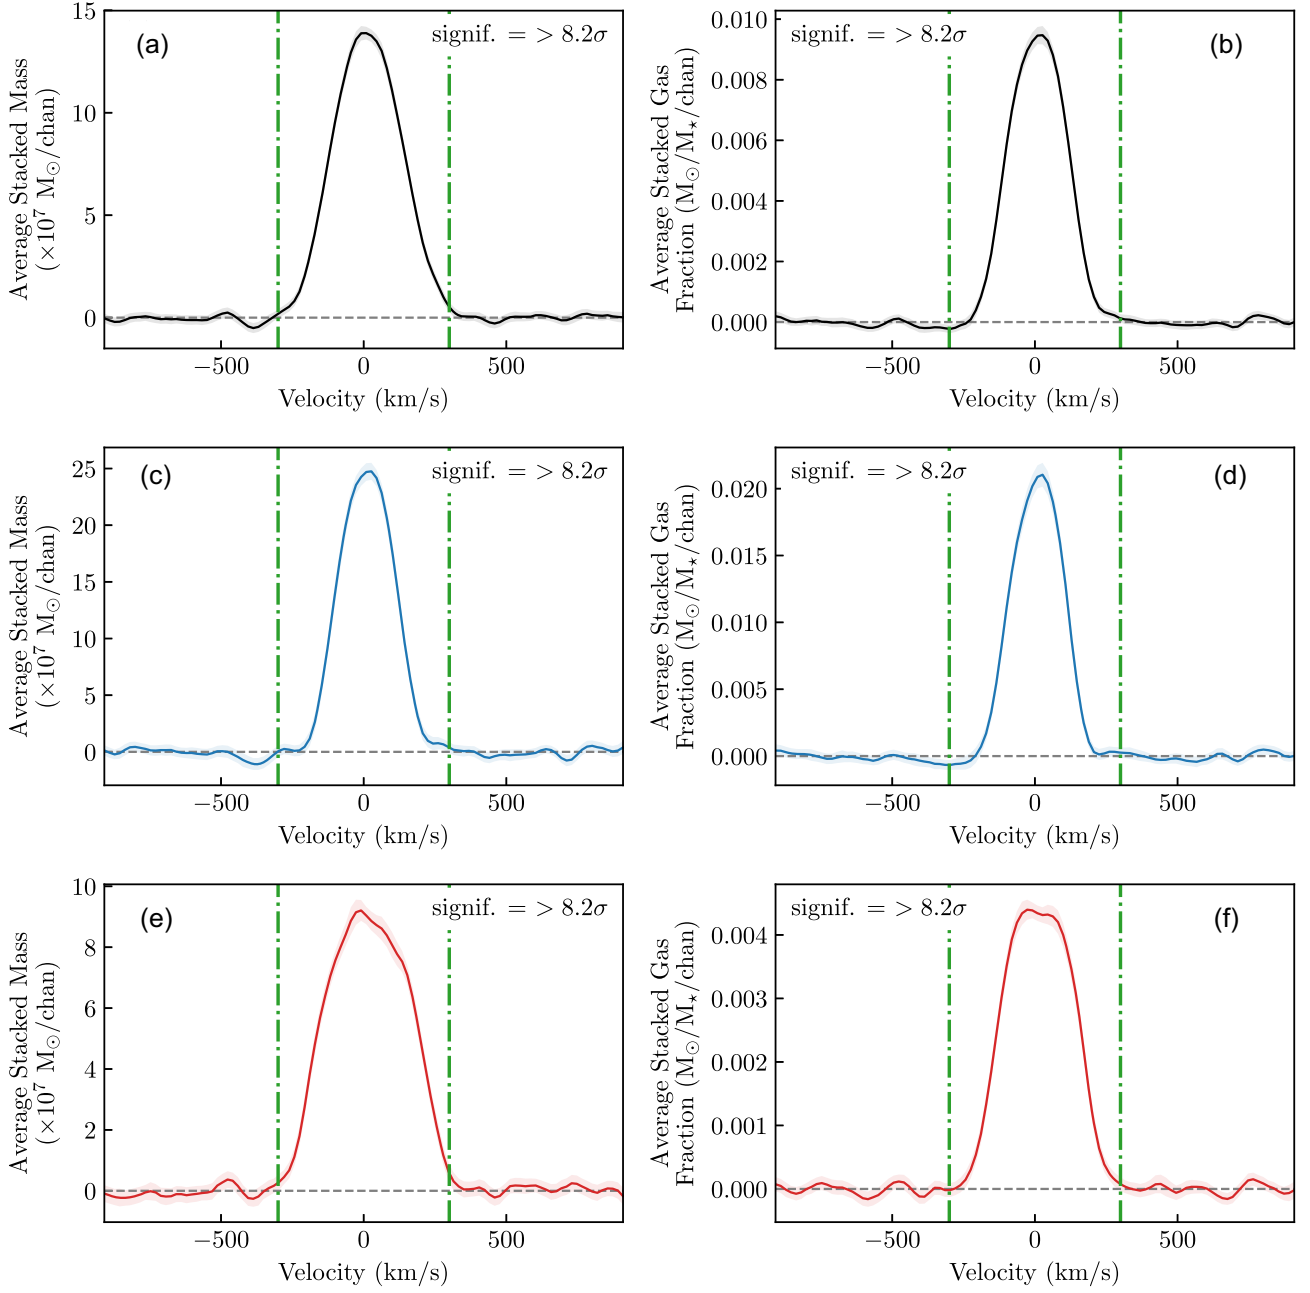


Figure D2. Same as Fig. D1, but for the sample with $10^{9.5} < M_{*}(M_{\odot}) < 10^{12}$. The measurements made from the above spectra can be found in Table D3.

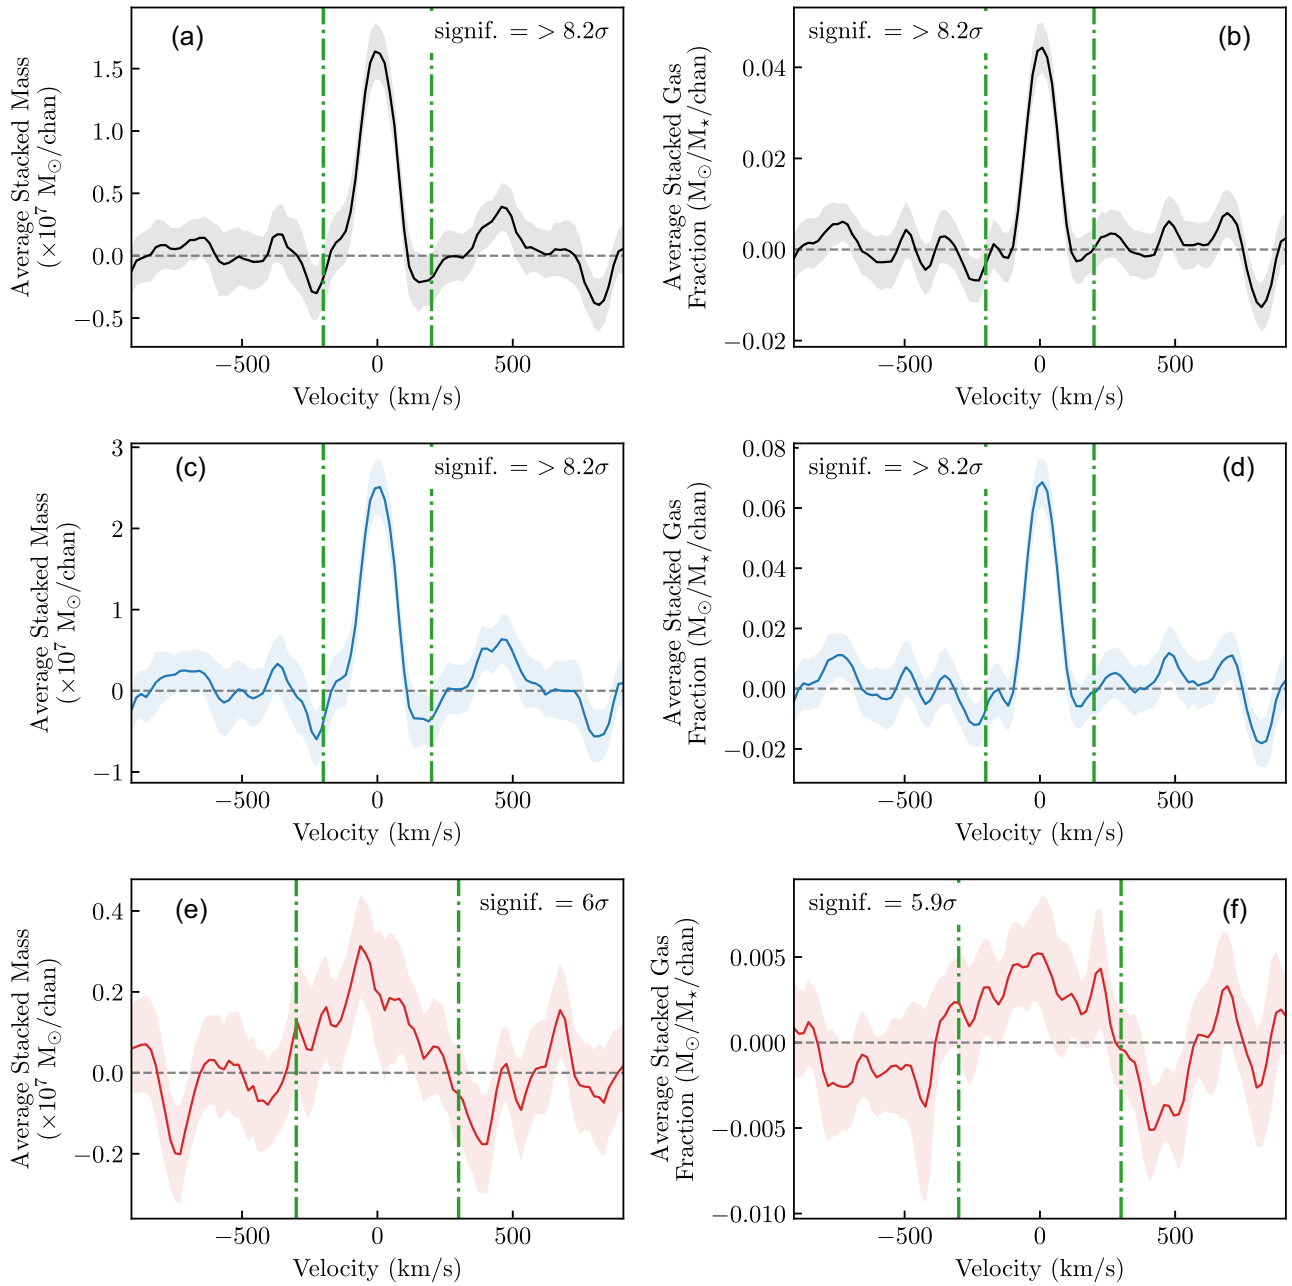


Figure D3. Same as Fig. D1, but for the sample of only H I non-detections with $10^8 < M_{*}(M_{\odot}) < 10^{9.5}$. Table D2 contains the measurements made from the above spectra.

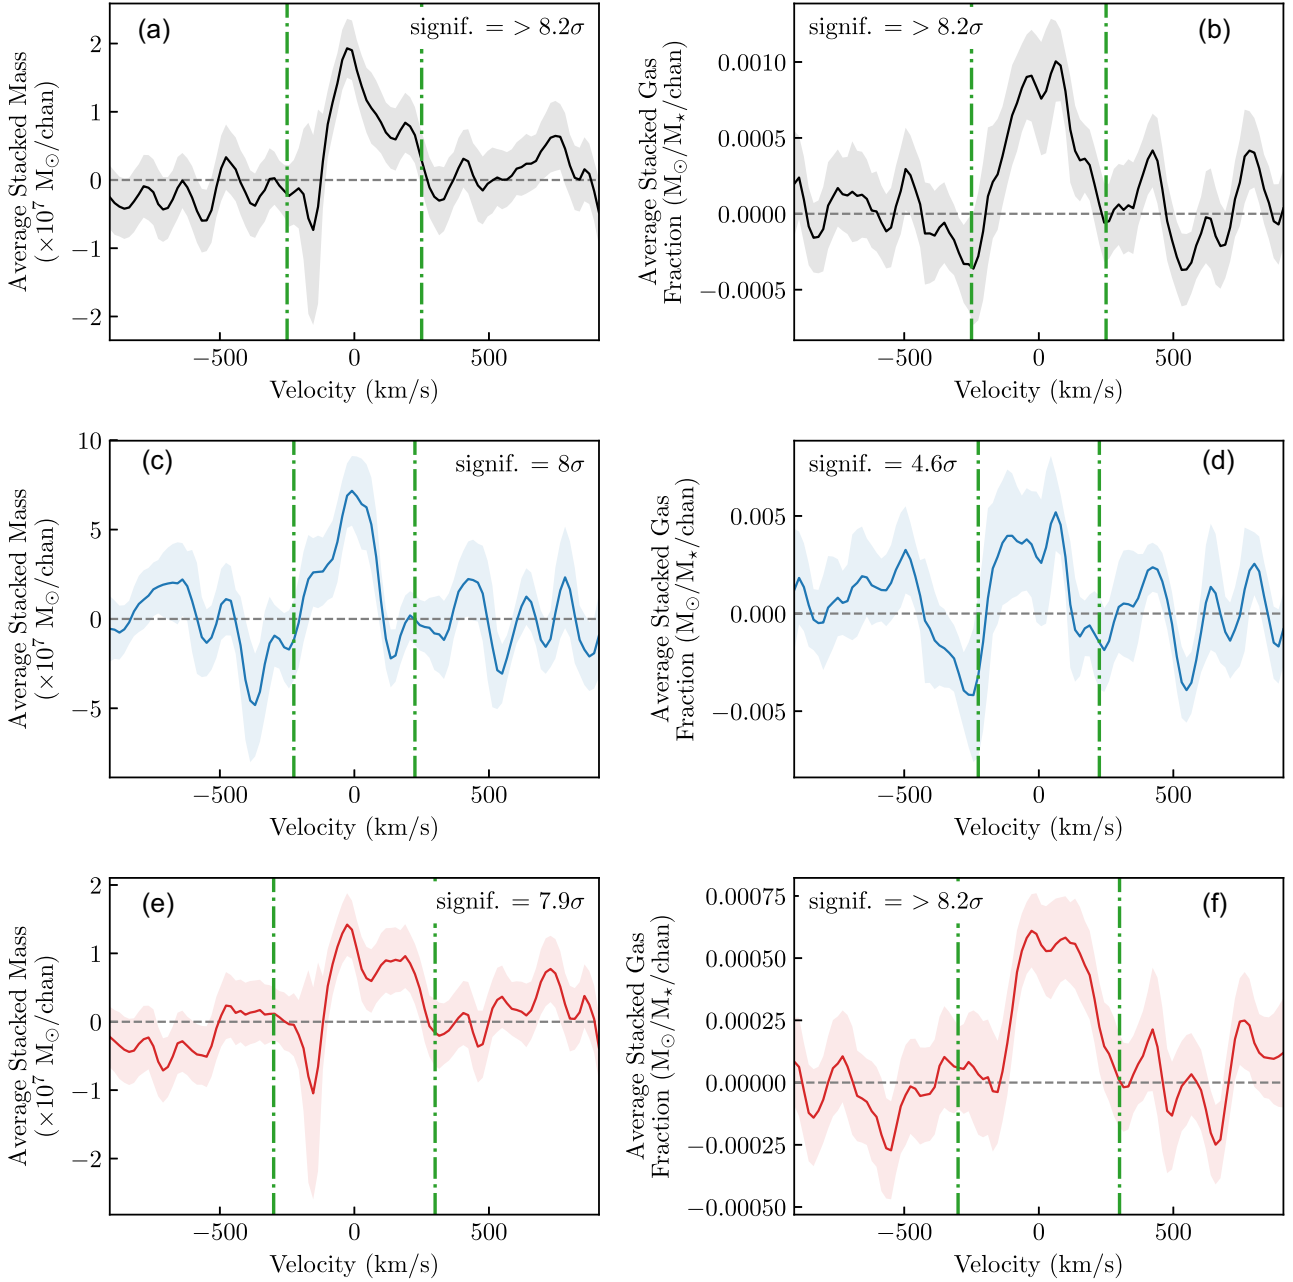


Figure D4. Same as Fig. D1, but for the sample of H I non-detections with $10^{9.5} < M_{*}(M_{\odot}) < 10^{12}$. Table D3 contains the measurements made from the above spectra.




Figure D5. All detections and non-detections stacked in bins of stellar mass. The bins increase in stellar mass from panel (a) to (d). The green dot-dashed lines in every panel mark the section of the stacked profiles used to measure the $\langle f_{\text{H I}} \rangle$. The grey band around the spectrum represents the uncertainty of the stacked spectrum. These spectra correspond to the $\langle f_{\text{H I}} \rangle$ measurements in Fig. 10. See Table 1 for measured quantities.

APPENDIX E: CATALOGUE OF H I NON-DETECTIONS

Table E1. Data for the H I non-detections in the stacking sample. The full table is available online.

Source	RA	Dec	Name	V_{opt}	$u - r$	M_r	M_{\star}
	(J2000)			(km s^{-1})	(mag)	(mag)	[log] (M_{\odot})
1926	11 23 18.84	03 57 18.78	CGCG 039-161	1582 ± 19	1.91	-16.18	8.23
1954	11 49 12.79	-03 39 34.64	PGC 1068442	1572 ± 3	1.57	-15.78	8.04
2012	12 19 21.12	04 46 24.49	ASK 173829	2065 ± 2	1.61	-15.87	8.09
2244	14 04 15.84	04 06 43.85	UGC 08986	1239 ± 3	2.19	-17.79	9.08
2272	14 25 08.99	-01 06 48.52	CGCG 019-019	2643 ± 3	2.11	-17.94	9.21
2442	15 17 29.45	03 35 07.94	PGC 3124577	1882 ± 2	1.62	-15.69	8.03
2525	09 35 44.04	31 42 19.68	NGC 2918	6756 ± 2	2.65	-22.64	11.28
2533	10 06 07.44	47 15 45.53	NGC 3111	7437 ± 2	2.68	-22.52	11.19
2535	14 30 25.58	11 55 40.77	NGC 5644	7649 ± 2	2.72	-22.87	11.33
2541	13 53 38.43	36 08 02.57	NGC 5352	7956 ± 3	2.81	-22.64	11.26
2565	00 35 59.83	-10 07 18.07	NGC 0163	5917 ± 4	2.70	-22.05	11.01
2574	11 00 35.40	12 09 41.65	NGC 3491	6351 ± 2	2.72	-21.98	11.06
2589	12 19 32.09	49 48 56.71	UGC 07367	4093 ± 2	2.60	-21.45	10.78
2597	12 19 32.59	56 44 11.69	NGC 4271	4756 ± 2	2.82	-21.84	10.93
2606	14 41 32.96	38 51 04.95	UGC 09473	4669 ± 2	2.17	-21.54	10.80
2614	12 12 17.27	13 12 18.72	NGC 4168	2273 ± 2	2.44	-21.43	10.63
2615	12 32 47.65	63 56 21.18	NGC 4512	2516 ± 1	2.72	-20.75	10.60
2661	10 36 38.44	14 10 15.97	NGC 3300	3017 ± 1	2.54	-21.02	10.53
2675	12 08 09.62	10 22 43.96	NGC 4119	1651 ± 2	2.28	-19.74	10.05
2682	12 27 13.34	12 44 05.22	NGC 4425	1891 ± 1	2.39	-20.42	10.22
2688	12 24 05.01	11 13 05.04	NGC 4352	2083 ± 1	2.36	-20.12	10.08
2710	12 27 02.54	15 27 41.34	NGC 4421	1551 ± 1	2.58	-20.15	10.04
2715	14 05 12.42	55 44 30.67	NGC 5475	1647 ± 1	2.54	-19.58	9.98
2731	12 37 48.38	05 22 06.69	NGC 4580	1035 ± 2	2.37	-19.23	9.75
2742	12 15 59.87	66 13 51.00	NGC 4221	1314 ± 1	2.51	-19.48	9.83
2743	10 12 41.25	03 07 45.79	NGC 3156	1266 ± 29	1.99	-19.14	9.42
2744	11 06 32.10	11 23 07.51	NGC 3524	1357 ± 1	2.36	-19.08	9.71
2756	11 08 40.38	57 13 48.71	NGC 3530	1876 ± 1	2.42	-18.95	9.69
2769	12 25 18.78	64 56 00.52	NGC 4391	1320 ± 1	2.52	-18.73	9.58
2770	13 00 10.57	12 28 59.92	NGC 4880	1362 ± 3	2.24	-19.29	9.51
2775	12 22 04.11	12 47 14.95	NGC 4306	1520 ± 3	2.21	-18.70	9.35
2789	11 21 24.99	03 00 50.23	NGC 3643	1742 ± 2	2.28	-18.67	9.53
2790	12 38 17.87	13 06 35.64	NGC 4584	1715 ± 3	2.01	-18.96	9.46
2795	12 41 59.35	12 56 34.27	NGC 4620	1125 ± 3	2.09	-18.63	9.24

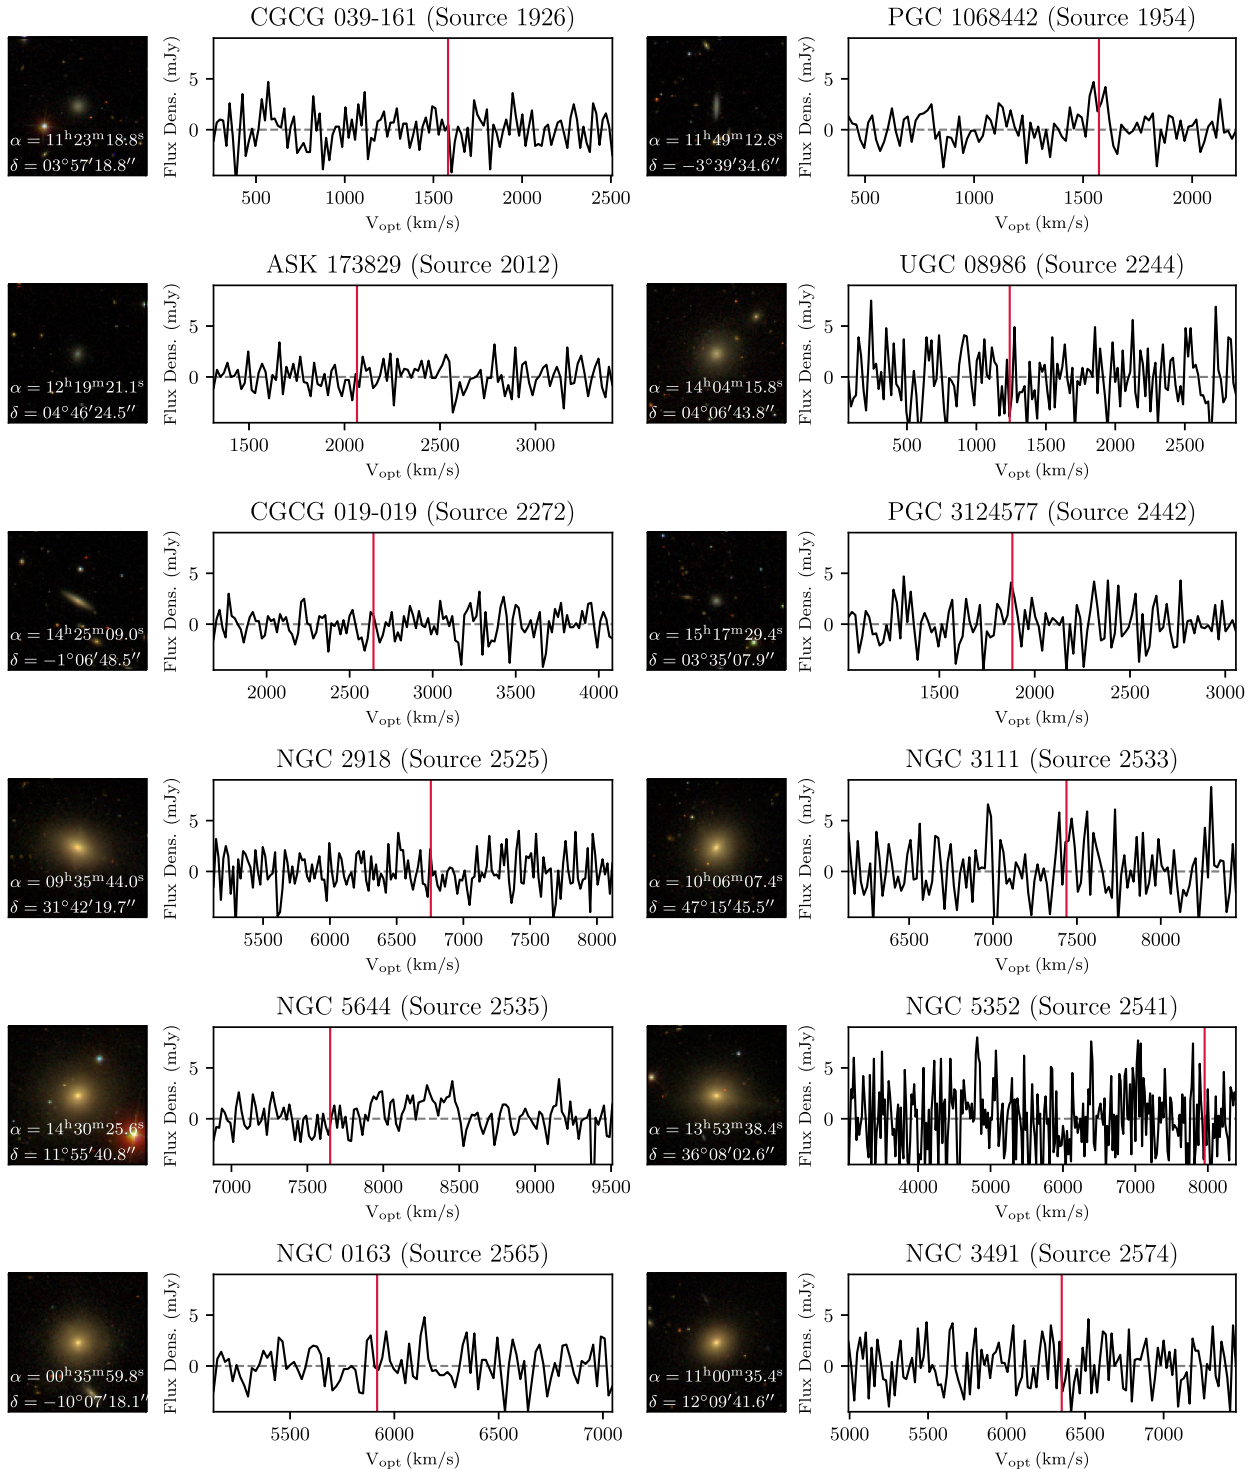


Figure E1. Optical images and H I spectra for all 323 galaxies classified as H I non-detections from the final NIBLES stacking sample (see Section 3.3.2). The SDSS redshift of each galaxy is represented by the red line in each spectrum. The rest of the images can be found in the online supplementary material.

This paper has been typeset from a \LaTeX file prepared by the author.