



Figuring Out Gas & Galaxies in Enzo (FOGGIE). I. Resolving Simulated Circumgalactic Absorption at $2 \leq z \leq 2.5$

Molly S. Peeples^{1,2} , Lauren Corlies^{2,3} , Jason Tumlinson^{1,2} , Brian W. O’Shea⁴ , Nicolas Lehner⁵ , John M. O’Meara^{6,7} , J. Christopher Howk⁵ , Nicholas Earl¹ , Britton D. Smith⁸ , John H. Wise⁹ , and Cameron B. Hummels¹⁰

¹Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA; molly@stsci.edu

²Department of Physics & Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA

³Large Synoptic Survey Telescope, Tucson, AZ 85719, USA

⁴Department of Computational Mathematics, Science and Engineering, Department of Physics and Astronomy, National Superconducting Cyclotron Laboratory, Michigan State University, East Lansing, MI 48823, USA

⁵Department of Physics, University of Notre Dame, Notre Dame, IN 46556, USA

⁶Department of Physics, St. Michael’s College, Colchester, VT 05439, USA

⁷W. M. Keck Observatory, Waimea, HI 96743, USA

⁸San Diego Supercomputer Center, University of California, San Diego, 10100 Hopkins Drive, La Jolla, CA 92093, USA

⁹Center for Relativistic Astrophysics, School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA

¹⁰Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

Received 2018 October 11; revised 2019 January 23; accepted 2019 January 26; published 2019 March 13

Abstract

We present simulations from the new “Figuring Out Gas & Galaxies in Enzo” (FOGGIE) project. In contrast to most extant simulations of galaxy formation, which concentrate computational resources on galactic disks and spheroids with fluid and particle elements of fixed mass, the FOGGIE simulations focus on extreme spatial and mass resolution in the circumgalactic medium (CGM) surrounding galaxies. Using the Enzo code and a new refinement scheme, FOGGIE reaches spatial resolutions of 381 comoving h^{-1} pc and resolves extremely low masses ($\lesssim 1\text{--}100 M_{\odot}$) out to 100 comoving h^{-1} kpc from the central halo. At these resolutions, cloud and filament-like structures giving rise to simulated absorption are smaller, and better resolved, than the same structures simulated with standard density-dependent refinement. Most of the simulated absorption arises in identifiable and well-resolved structures with masses $\lesssim 10^4 M_{\odot}$, well below the mass resolution of typical zoom simulations. However, integrated quantities such as mass surface density and ionic covering fractions change at only the $\lesssim 30\%$ level as resolution is varied. These relatively small changes in projected quantities—even when the sizes and distribution of absorbing clouds change dramatically—indicate that commonly used observables provide only weak constraints on the physical structure of the underlying gas. Comparing the simulated absorption features to the KODIAQ (Keck Observatory Database of Ionized Absorption toward Quasars) survey of $z \sim 2\text{--}3.5$ Lyman limit systems, we show that high-resolution FOGGIE runs better resolve the internal kinematic structure of detected absorption and better match the observed distribution of absorber properties. These results indicate that circumgalactic medium resolution is key in properly testing simulations of galaxy evolution with circumgalactic observations.

Key words: galaxies: evolution – hydrodynamics – intergalactic medium – quasars: absorption lines

Supporting material: machine-readable table

1. Introduction

In the galactic ecosystem, the majority of baryons and heavy elements are in the diffuse gas outside of galaxies rather than in the relatively dense gas and stars composing the disks and central spheroids (Pagel 2008; Stocke et al. 2013; Peeples et al. 2014; Werk et al. 2014; Prochaska et al. 2017). Hydrodynamic simulations of galaxy evolution have met with mixed success in reproducing the observed column densities of low- or high-ionization circumgalactic gas, despite creating realistic-looking galaxies (Hummels et al. 2013; Ford et al. 2016; Oppenheimer et al. 2016; Suresh et al. 2017; see also Tumlinson et al. 2017). The focus of these theoretical studies has usually been the improvement of stellar feedback models: how the energy and momentum from star formation and supermassive black holes are coupled with the interstellar medium (ISM) and propagated into the circumgalactic medium (CGM) and beyond (e.g., Oppenheimer & Davé 2008; Trujillo-Gomez et al. 2015; Christensen et al. 2016; Salem et al. 2016; Christensen et al. 2018). While feedback is certainly important for understanding how galaxies and the circumgalactic medium coevolve, we posit here that numerical

resolution in the circumgalactic medium is a critically important factor in the comparison between circumgalactic medium data and cosmological simulations that influences results for any subgrid physical models. We show here that improving the resolution of circumgalactic gas strongly affects the simulated circumgalactic medium structures that give rise to detected absorption; in Paper II of this series (Corlies et al. 2018), we show that a resolved circumgalactic medium is crucial for predicting and subsequently understanding the sources of circumgalactic emission.

Observational evidence has existed for decades that the circumgalactic medium is structured on scales smaller than the $\gtrsim 1\text{--}10$ kpc scales that are typically resolved in cosmological simulations. Absorbers seen along lines of sight toward multiply lensed quasars have structure on the scale of $\lesssim 1\text{--}10$ kpc, with lower-ionization species (e.g., Mg II or Si II, thought to trace cooler, higher-density gas) showing structure on smaller scales than more highly ionized species (e.g., C IV or O VI; Rauch et al. 2001a, 2001b; Ellison et al. 2004; Rubin et al. 2018b). Milky Way high-velocity clouds (HVCs) with well-constrained distances show HI masses of $\sim 10^5\text{--}5 \times 10^6 M_{\odot}$ (e.g., Wakker et al. 2008;

Putman et al. 2012) and sizes in the few to ~ 15 kpc range. Given how relatively close to the disk most HVCs are (e.g., Wakker et al. 2007, 2008; Thom et al. 2008; Lehner & Howk 2011; Lehner et al. 2012), these clouds may or may not be representative of the Milky Way circumgalactic medium as a whole. Recent evidence from absorption toward background galaxies—which have a larger effective beam size than quasars and thus wash out small-scale structure—indicates that the “coherence scale” of circumgalactic absorption is on the order of tens of kiloparsecs (Lopez et al. 2018; Péroux et al. 2018; Rubin et al. 2018a). Therefore, to fully resolve the *observed* structure of the circumgalactic medium, simulations must reach spatial scales of $\lesssim 1$ kpc in order to have at least tens of resolution elements composing such structures.

There are several physical processes that can give rise to multiphase structure at subkiloparsec length scales. The cooling length (the sound speed times the cooling time; see, e.g., Figure 8 of Smith et al. 2017) has scales from tens of parsecs to hundreds of kiloparsecs for typical circumgalactic densities, temperatures, and metallicities, with photoionizational cooling from the metagalactic UV background.¹¹ At low densities, length scales on the smaller end of this range translate to very small mass scale—tens to hundreds of solar masses—which can be challenging to achieve in simulations evolved with fixed mass resolution (e.g., those evolved with particle-based or moving mesh codes). Resolving the cooling length in the density and temperature regimes of interest to the inner circumgalactic medium, where $l_c \simeq 0.1\text{--}10$ kpc, is critical to resolving the thermal instability in the circumgalactic medium. This instability is potentially responsible for the formation of multiphase gas in the circumgalactic medium (Voit et al. 2015). If the simulation’s physical resolution is larger than the local cooling length in the region of interest, the formation of a multiphase medium will be either dampened or suppressed entirely.

Other physical scales that may affect the properties of the multiphase circumgalactic medium include (1) the convective scale of the circumgalactic medium, which is approximately equal to the pressure scale height of the stratified medium and is the typical physical scale over which a parcel of fluid may move and bulk mixing might occur; (2) the scale on which turbulence is driven, which also affects mixing of metals and gas of different temperatures; and (3) the Field length, which is the physical length scale where the thermal conduction of heat into a cold gas cloud is balanced by radiative cooling, and which defines the minimum scale of cold gas clouds (Field 1965; Sharma et al. 2010). The pressure scale height of a virialized Milky Way–sized galaxy halo is ~ 20 kpc, suggesting that the convective scale is not setting the size of the small-scale structures in the circumgalactic medium. The driving scale of turbulence, and thus the largest physical objects whose properties are likely to be dominated by this phenomenon, depends on the source of the driving. In the case of cosmological accretion, that scale would be comparable to the diameter of the cosmological filament, or a significant fraction of the virial radius. In the case of turbulence driven by galactic winds/outflows, it is comparable in size to the width of the outflows—on the order of kiloparsecs to a few tens of kiloparsecs, depending on the galaxy’s recent star formation

rate or AGN behavior. Likewise, the Field length can have values spanning orders of magnitude: in 10^6 K diffuse gas it is tens of kiloparsecs (substantially larger than the cooling length), but in the denser 10^4 K gas the Field length is orders of magnitude smaller, comparable to the cooling length. In principle, the Jeans or Bonner–Ebert length scales—the physical scale where gravity dominates over thermal pressure—may be important as well, but given that the local gravity field is dominated by the host dark matter halo, it is likely that gas clouds are not self-gravitating (unlike in the intergalactic medium; Schaye 2001; Peeples et al. 2010a, 2010b), and thus hydrodynamic processes will dominate (see, e.g., Liang & Remming 2018). The physical scale of a cloud is likely to be smaller than the convective and turbulent driving scales and smaller than the cooling length scale, as these drive fragmentation and the onset of the thermal instability, respectively. The cloud must also be larger than the Field length, or else it will evaporate. If these scales are inverted (i.e., if the Field length is larger than these scales), the physical processes that drive precipitation will be suppressed. Most simulations to date either explore the small-scale structure of circumgalactic gas in highly idealized scenarios (Armillotta et al. 2017; Fielding et al. 2018) or explore regimes more akin to the disk–halo interface (i.e., much denser gas) than the diffuse circumgalactic medium (Schneider & Robertson 2017), although see Churchill et al. (2015) and van de Voort et al. (2019).

In this paper, we explore the consequences of resolving the circumgalactic medium to subkiloparsec scales, regardless of what physical processes may be responsible for generating structure on these small scales. In other words, we explore the effects of numerical resolution without changing the adopted feedback, to assess them independently. We describe our simulations and improved resolution technique in Section 2; our method for extracting and analyzing synthetic spectra is described in Section 2.3. The Keck Observatory Database of Ionized Absorption toward Quasars (KODIAQ; Lehner et al. 2014; O’Meara et al. 2017) sample to which we compare these synthetic spectra is described in Section 3. In Section 4, we show the impact our improved resolution has on the physical properties of the circumgalactic medium, and in Section 5 we show that we resolve the predominant structures responsible for most circumgalactic absorption at these redshifts. In Section 6, we compare the kinematic properties of the synthetic spectra to the KODIAQ absorbers, and we provide our concluding thoughts in Section 7. As the bulk of the high signal-to-noise ratio (S/N) KODIAQ absorbers with a wide range of available line transitions are at $2 \leq z \leq 2.8$, we focus our simulation resolution comparison here to $z = 2$ and $z = 2.5$. All lengths and distances are given in physical units unless specifically stated otherwise.

2. The Simulations

2.1. Basic Properties

The Figuring Out Gas & Galaxies In Enzo (FOGGIE) simulations are cosmological hydrodynamic simulations evolved with the block-structured adaptive mesh refinement (AMR) code Enzo (Bryan et al. 2014) using a flat Planck Collaboration et al. (2014) Λ CDM cosmology ($1 - \Omega_\Lambda = \Omega_m = 0.285$, $\Omega_b = 0.0461$, $h = 0.695$) and a Piecewise Parabolic Method (PPM) hydro solver. The simulation domain is a $(100 h^{-1} \text{cMpc})^3$ box. We generated the initial conditions with a simple 1024^3 -particle dark-matter-only simulation on the NCSA Blue Waters supercomputer. From this DM-only run we

¹¹ We note that most of the densities considered in McCourt et al. (2018) are higher than the circumgalactic densities we are interested in here; these higher physical densities translate to much smaller cooling lengths than typical for the circumgalactic medium.

selected a halo, which we have named “Tempest,” that has a Milky Way mass at $z = 0$ ($M_{\text{halo}} \simeq 1.5 \times 10^{12} M_{\odot}$) and a relatively quiescent merger history at $z < 1$. The initial conditions for the Tempest halo were then regenerated using the “cosmological zoom” method, with a 256^3 grid cell/particle base resolution and an effective resolution of 4096^3 particles ($M_{\text{dm}} = 1.39 \times 10^6 M_{\odot}$) in the region of interest, which is a Lagrangian region encompassing all of the particles within two virial radii of the galaxy at $z = 0$. These “zoom” simulations were then evolved on the NASA Pleiades supercomputer with a maximum of 11 levels of adaptive mesh refinement (cell sizes of $\simeq 190 \text{ pc } h^{-1}$ comoving). The Tempest halo has $R_{200} = 37 \text{ kpc}$ and $M_{200} = 3.9 \times 10^{10} M_{\odot}$ at $z = 2$.

We use Enzo’s Cen & Ostriker (2006) thermal supernova feedback model, forming stars in gas exceeding a comoving number density of $\simeq 0.1 \text{ cm}^{-3}$ with a minimum star particle mass of $2 \times 10^4 M_{\odot}$. Supernova feedback is composed of purely thermal energy that is deposited into the 27 nearest cells surrounding the star particle, after 12 gas dynamical times have elapsed since the star particle formed. The total energy imparted is $1.0 \times 10^{-5} m_{*} c^2$, the total mass ejected is $0.25 m_{*}$, and the total metal mass ejected is

$$0.025 m_{*}(1 - Z_{*}) + 0.25 Z_{*}, \quad (1)$$

where all metals are tracked as a single field. As a result, particular elemental abundances throughout the paper are calculated assuming solar relative abundances scaled to the local metallicity computed at runtime. The effects of Type Ia supernovae are not included.

The simulation includes metallicity-dependent cooling and a metagalactic background (Haardt & Madau 2012) using the Grackle¹² chemistry and cooling library (Smith et al. 2017). The code simultaneously solves a nonequilibrium six-species chemical reaction network (tracking H I, H II, D I, D II, He I, He II, He III, and e^{-}), computing cooling directly from these species. No correction is made for self-shielding. A metal field is followed separately, with cooling calculated assuming ionization equilibrium and solar composition.

2.2. Forcing Resolution in the circumgalactic medium

In order to isolate the effects of circumgalactic resolution on observable absorption-line spectra, we consider multiple runs with the same subgrid physics and interstellar resolution, varying only the refinement scheme that is applied in the circumgalactic medium gas. To impose uniform *spatial* resolution on the circumgalactic medium, we require a method that overrides the normal criteria for refinement on physical criteria such as gas density, dark matter overdensity, or metallicity. In its place, we use a “forced refinement” scheme, which imposes a fixed resolution on a specific subregion of the computational domain that can move over time. This region is specified in a parameter file read by Enzo at runtime. For this first generation of FOGGIE simulations, we adopt a cubic forced refinement region that runs from $-100 h^{-1}$ to $+100 h^{-1}$ comoving kpc from the moving center of mass of the targeted galaxy in each spatial dimension. The path of the halo of interest through the domain is obtained by tracking the halo’s center of mass (including gas and dark matter) from a lower-resolution run evolved to $z = 0$ with “standard” refinement. This run consumes minimal resources and serves as an important controlled comparison to the forced refinement run.

Once the 3D track through the domain is determined, we rerun the calculation with the forced refinement box tracking the galaxy, starting at $z = 4$, when the galaxy has formed $< 5\%$ of the stars it has by $z = 2$. When forced refinement is on, its level of refinement is specified as a parameter and fixed. Outside the forced refinement box, standard refinement is used with a gas overdensity criterion. Within the tracked box, refinement above the fixed level can be specified to follow a density criterion. In this way, our primary production run uses a fixed refinement of $n_{\text{ref}} = 10$ levels ($380 h^{-1} \text{ pc}$ comoving) over the $(200 h^{-1} \text{ ckpc})^3$ box, with additional refinements one level further to $n_{\text{ref}} = 11$ ($190 h^{-1} \text{ cpc}$), where the density criterion is met. We find that the forced $n_{\text{ref}} = 10$ refinement increases to 11 only in the galaxy’s disk and the ISM of satellites. This uniform method of refinement is generally *much* more computationally efficient than a standard AMR run with a similar total number of resolution elements because the grid patches are typically larger and more uniformly sized, meaning that there are fewer overall ghost zones and better cache utilization.

Figure 1 illustrates this technique with slices of the cell size (i.e., effective resolution) through the center of the halo at $z = 2$; note that the interstellar resolution of $190 h^{-1}$ comoving pc (91 physical pc, shown in green) is the same in both simulations. Slight differences in grid structure outside of the refinement region are due to the algorithm used for cell flagging and AMR grid creation. We alternately refer to the standard-resolution run as the “density-refined” run and the high-resolution simulation as the “refined” or “forced refinement” simulation. Both simulations have been evolved to $z = 0$.

As a demonstration and test of convergence, we have evolved a third simulation to $z = 2.5$ with a uniform 290 comoving pc ($n_{\text{ref}} = 11$) resolution within the 200 ckpc volume. We show the results from this $n_{\text{ref}} = 11$ everywhere simulation and how it compares to the $z = 2.5$ in the other two simulations in Appendix A. Broadly, we find that these two forced resolution simulations are much more similar in the metrics we consider throughout this paper than either is to the standard-resolution simulation.

Critically, enforcing small spatial scales to be resolved translates into very low *masses* per resolution element in the low-density gas. Figure 2 shows the masses of the cells projected along a position axis through the forced refinement region in the two simulations in temperature bins. The first clear difference between the two simulations is that there are *many* more resolution elements in the high-resolution simulation, particularly in the circumgalactic medium; at $z = 2$, the high-resolution run has $\sim 400\times$ as many, with $\sim 1.4 \times 10^8$ cells within the forced refinement region, while the standard-resolution simulation has $\sim 350,000$ cells within the corresponding volume. In the standard-resolution simulation, most of the circumgalactic gas is traced by cells with masses of $\gtrsim 10^4 M_{\odot}$, while in the high-resolution case, essentially all of the gas at temperatures $\gtrsim 10^{4.5} \text{ K}$ is resolved by cells with masses $< 1000 M_{\odot}$, as is most of the gas at lower temperatures. As we show in Section 5.2, resolving these small masses allows the FOGGIE simulations to resolve the small structures responsible for most circumgalactic absorption, even when that gas is at relatively low density.

2.3. Synthetic Absorption Spectroscopy

We extract synthetic spectra from the simulations using the Trident package (Hummels et al. 2017), adopting the Morton (2003) atomic data. Trident first uses Cloudy (last described in

¹² <https://grackle.readthedocs.io/en/latest/>




Figure 1. Slices showing the resolution of the standard-resolution (left) and high-resolution (right) simulations at $z = 2$. The cell sizes are given in physical kpc. The galaxy is located at the center of each slice; note that the ISM resolution does not change.

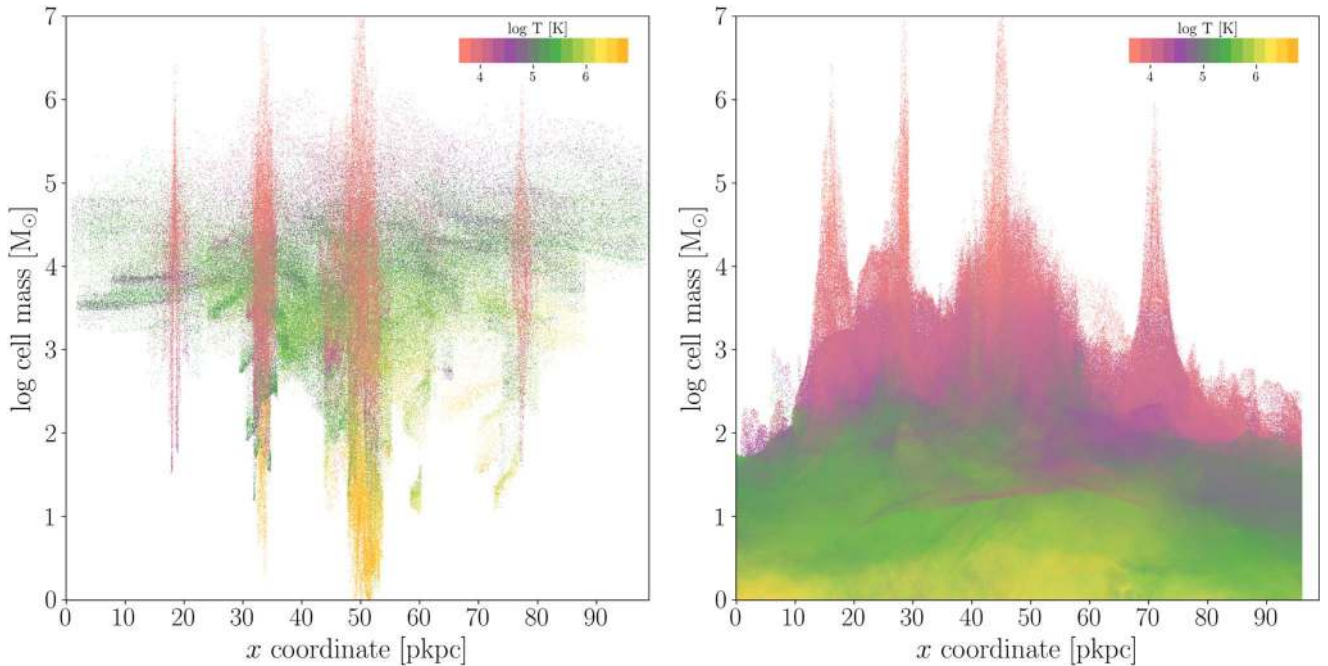


Figure 2. Distribution of cell masses within the $200h^{-1}$ cKpc forced refinement region at $z = 2$, color-coded by the most common temperature at that location and cell mass. In both panels, the cold, high-cell-mass peaks correspond to the dense ISM of the main galaxy and its satellites, while the lower-cell-mass gas at intervening positions is the circumgalactic medium. The left panel shows the cell masses in the standard simulation, where cells vary in size in proportion to the local density, and most of the circumgalactic medium is traced by structures $\gtrsim 10^4 M_{\odot}$. The right panel shows the same for the high-resolution simulation, where the uniform cell size in the circumgalactic medium is 548 comoving pc ($380/h$) and thus circumgalactic medium cell masses are $\sim 1\text{--}100 M_{\odot}$.

Ferland et al. 2013) to calculate the ionization fraction of the species of interest based on the cell-by-cell density, temperature, and metallicity; as when the simulations were evolved, we assume a Haardt & Madau (2012) metagalactic background (though the HI fraction is tracked natively by Enzo). To calculate the effective redshift of absorption features, Trident assumes a smooth Hubble flow along a line of sight. This is then added to the 3D velocity field of each AMR grid cell intersected by the sightline. The absorption produced by a grid cell is represented by a single Voigt profile at an instantaneous

velocity v and a Doppler b parameter specified by the temperature in the cell (i.e., the Hubble flow across the grid cell is neglected, no model is applied for turbulence on scales smaller than the cell size, etc.). Trident returns the optical depth and normalized flux as a function of observed wavelength along the line of sight; in order to ease the comparison with the observations, we consider each ionic transition individually so that we can uniquely convert the observed wavelength back to the effective redshift and thus relative velocity of the system. We then resample the spectra to 2 km s^{-1} per pixel, the

Table 1
Minimum Considered Column Densities for Each Ion of Interest

Species	log (Minimum Column Density) (cm^{-2})
H I	16
Si II	11
Si IV	12
C IV	13
O VI	13

approximate binning of Keck’s High Resolution Echelle Spectrometer (HIRES; see Section 3).¹³ Finally, we convolve the spectra with a 7 km s^{-1} full-width half-maximum (FWHM) Gaussian profile to mimic a representative HIRES line-spread function (LSF).¹⁴

The resampling, LSF handling, and characterization of the absorption features are part of the `spectacle`¹⁵ package, which is still under development and which we will release publicly in the near future (N. Earl & M. S. Peeples 2019, in preparation). Here, the only analysis feature we consider is the number of minima below a given flux threshold (Section 6). In practice, this threshold also serves to set the effective S/N of the synthetic spectra to be similar to the level of detected features in the real spectra without the need for developing an automated fitting routine that can also properly account for noisy data and uncertainties. As the S/N of the KODIAQ data varies from sightline to sightline (Section 3), we adopt a threshold in the normalized flux of 0.95 (requiring that the minima absorb more than 5% of the unattenuated normalized flux) in order to mimic the effects of noise in the real spectra.

For analyses based on these spectra (Sections 5, 6), we consider only simulated sightlines with $10^{16} \text{ cm}^{-2} < N_{\text{H I}} < 10^{21} \text{ cm}^{-2}$, taking 100 lines of sight along each orthogonal axis at $z = 2$ and $z = 2.5$, for a total of 600 sightlines per simulation. Each sightline has a pathlength of $200h^{-1}$ comoving kpc, which is wholly contained within the forced refinement regions (or, for the case of the standard-resolution simulation, the corresponding volume). While this pathlength is certainly smaller than the pathlengths giving rise to the Lyman-Limit Systems (LLSs) seen in KODIAQ (Section 3), this selection enables a comparison that is wholly caused by the difference in the underlying physical resolution in the two simulations. For each metal species, we further consider only sightlines with a minimum column density, roughly corresponding to the minimum column density at which it is observed in the KODIAQ data, as given in Table 1.

3. Observed Absorption Spectroscopy: The KODIAQ Sample

To facilitate comparison with real data, we require a sample of high resolution, high signal-to-noise data covering H I and

¹³ Note that because Trident extracts spectra in even bins of wavelength instead of even bins of velocity, this resampling requires nonuniform interpolations. However, this does not affect our results because of the relatively small wavelength ranges considered and the extremely high resolution at which the initial spectra are extracted from the simulation.

¹⁴ Though it is more physical to apply the LSF and then resample the spectrum, we do this order of operations in reverse because Trident extracts the spectra in wavelength space; again, our tests show that at these spectral resolutions the order of operations does not affect our results.

¹⁵ <https://spectacle-py.readthedocs.io/en/latest/>




Figure 3. Example of a typical KODIAQ absorber: normalized profiles of low and high ions as a function of the rest-frame velocity for the absorber toward Q1009+2956. Each ion was fitted simultaneously (only one transition for Si II is shown, but three were used in the profile fitting), but each species was fitted independently. The red lines indicate the global component model, and the blue lines show the individual components. The vertical orange dotted lines represent the velocity centroids for Si II, while the red tick marks show the velocity centroids for each species. Absorption that is not fitted is attributed to a contamination from an unrelated absorber.

metal-line absorption across multiple ionization states. We therefore compare to spectra from the DR2 release of KODIAQ (O’Meara et al. 2017), a large sample of Keck HIRES data on $z > 2$ QSOs. We select those spectra that can provide an accurate measurement of $N_{\text{H I}}$, which in turn requires that the absorption is at $z \gtrsim 2.5$ if $N_{\text{H I}}$ is determined from the termination of the H I Lyman series, or $z \gtrsim 1.6$ if the $N_{\text{H I}}$ is determined from damping wings in the H I Ly α line. In general, we favor absorption at $z > 2$

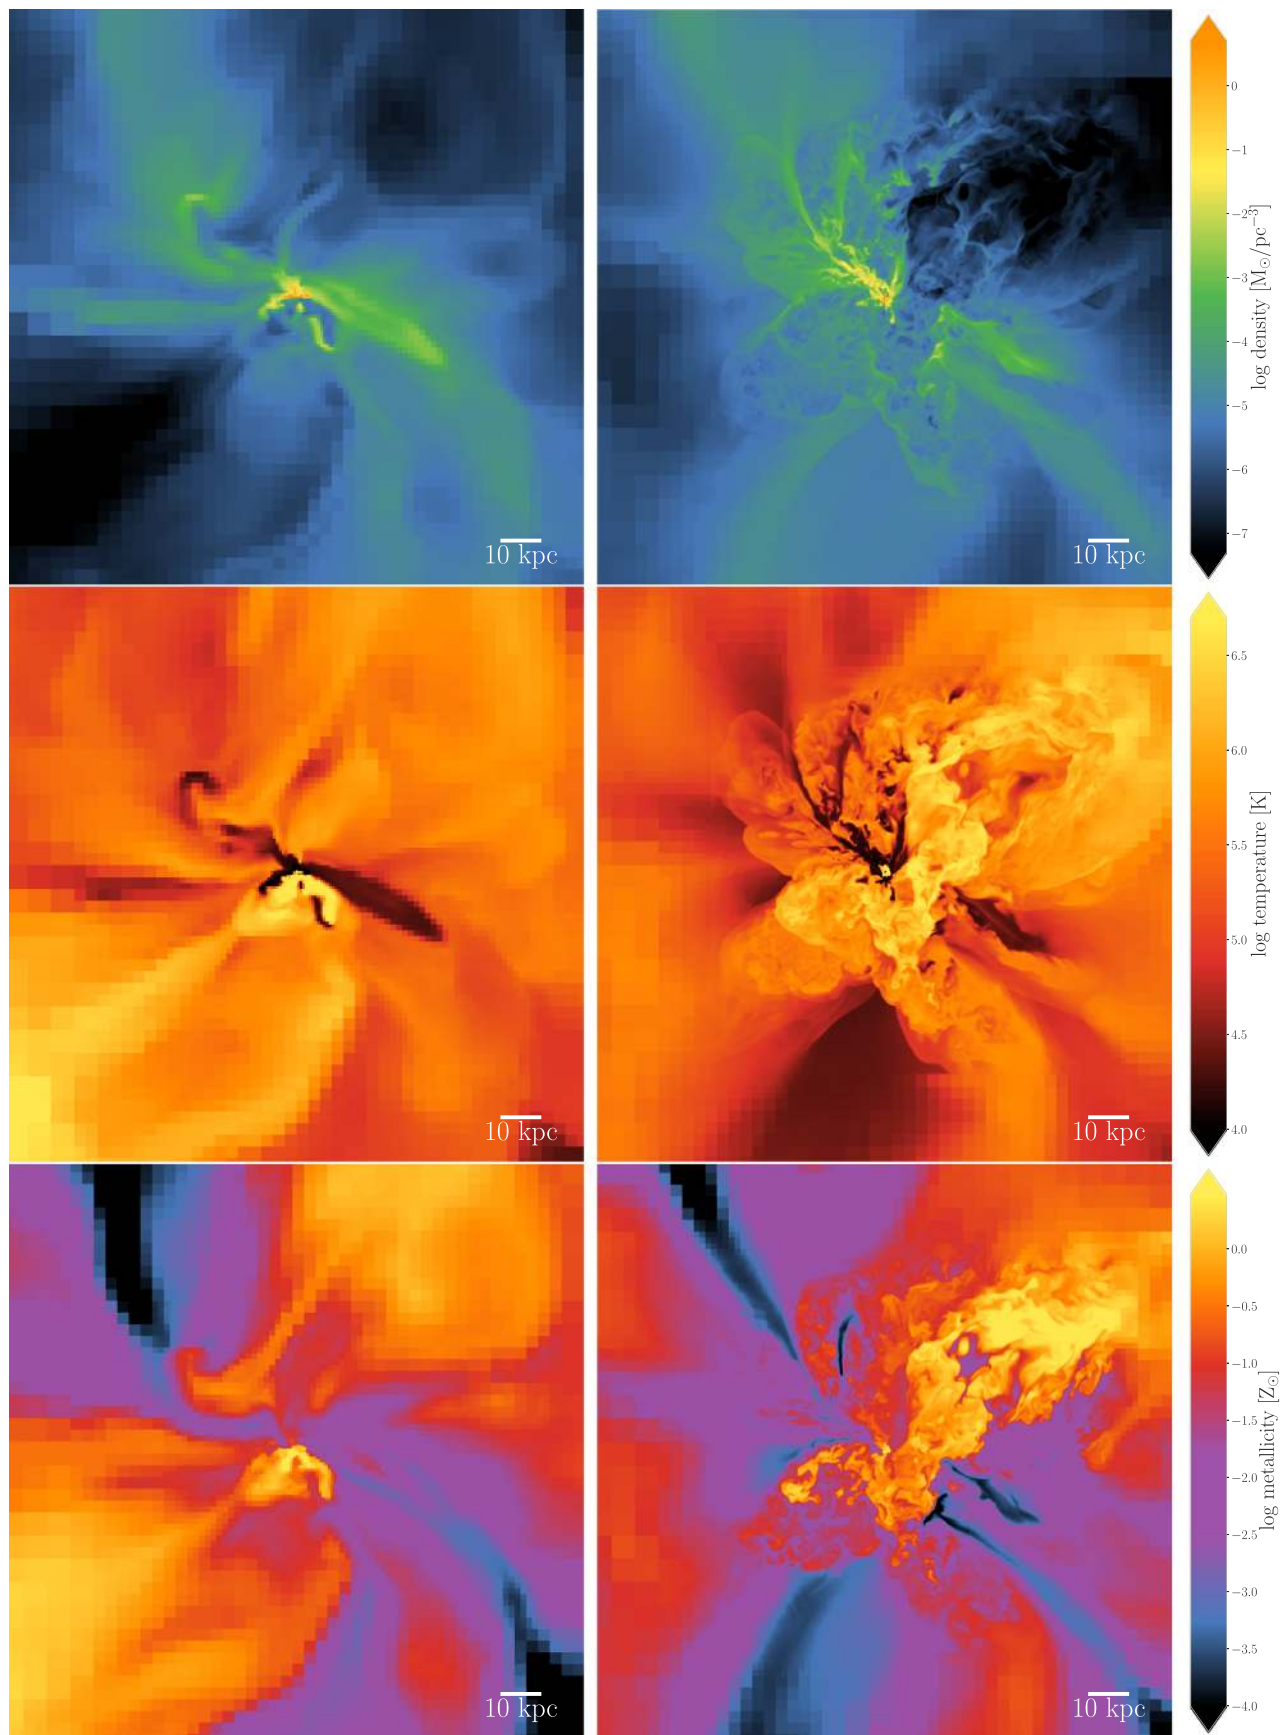


Figure 4. Slices of density (top), temperature (middle), and metallicity (bottom) through the central halo in the standard-resolution (left) and high-resolution (right) simulations at $z = 2$.




Figure 5. H I, Si II, Si III, Si IV, C IV, and O VI projections at $z = 2$ in the standard-resolution (left) and high-resolution (right) simulations. Each panel is $200h^{-1}$ ckpc (~ 97 pkpc) across and deep. The color map for H I is chosen such that Lyman limit gas is black/blue and DLA gas is red/orange, with optically thin gas in gray scale.

so that, in addition to H I, key temperature and ionization diagnostic ions such as O VI are covered in the spectra. The full KODIAQ DR2 sample includes hundreds of objects, but most have S/Ns too low to provide useful comparisons to our simulated spectra (typically $S/N > 30$ per pixel at $Ly\alpha$). Thus, we

further constrain our sample to a subset of previously studied KODIAQ absorbers in Lehner et al. (2014, 2016).¹⁶ To facilitate

¹⁶ The only exception is the absorber toward Q2126–158 at $z = 2.90731$, which was selected based on its $N_{H I}$ from an unpublished survey (Burns 2014).

Table 2Mass of All Gas and Different Ionic Species within the $(200 \text{ ckpc}/h)^3$ Forced Refinement Volume in the Two Simulations at $z = 2$

Species	High Resolution	Standard Resolution	Ratio
Stars	1.289×10^{10}	1.118×10^{10}	1.15
All Gas	1.082×10^{10}	1.215×10^{10}	0.890
H I	3.412×10^9	4.003×10^9	0.852
Si II	4.686×10^6	5.727×10^6	0.818
Si IV	1.421×10^4	5.048×10^2	28.2
C IV	4.119×10^4	3.529×10^4	1.17
O VI	4.400×10^4	4.429×10^4	0.993

Note. Masses are given in M_\odot . The rightmost column gives the ratio of the mass in the high-resolution simulation to that in the standard-resolution simulation. The total baryonic masses within the refined region (stars + gas) differ by $<2\%$.

a study of O VI absorption in Lyman limit and damped Ly α systems, the spectra in Lehner et al. (2014) were selected to have strong ($\log N_{\text{HI}} > 17$) H I absorption and little to no contamination of the O VI $\lambda\lambda 1032, 1037$ lines by the Ly α forest. The spectra in Lehner et al. (2016) were selected purely on H I to have $16.2 < \log N_{\text{HI}} < 19$ and be at $z > 2$.

We adopt the Lehner et al. (2014) fits for the high ions (C IV, Si IV, O VI, N V). However, we also want to compare FOGGIE to the observed low ions, which were not fitted in Lehner et al. (2014 or 2016). By including the low ions, we can directly determine how the kinematics of the absorption profiles change over a large range of ionization states. We select Si II as the main low ion because we can use several transitions (Si II $\lambda\lambda 1190, 1193, 1260, 1304, 1526$) in our profile fitting, which allows us to assess any contamination reliably and hence accurately model the absorption profiles. To fit the individual components of Si II, we follow the overall methodology undertaken by Lehner et al. (2014). In short, we use a modified version of the software described in Fitzpatrick & Spitzer (1997), which can simultaneously fit several transitions of the same ion or atom. The best-fit values describing the gas are determined by comparing the data to composite Voigt profiles convolved with an instrumental line-spread function (LSF). The LSFs are modeled as a Gaussian with an FWHM derived from the resolution of the KODIAQ spectra (which typically vary between 4 and 10 km s^{-1} FWHM; see Lehner et al. 2014, 2016; O’Meara et al. 2017). The three parameters—the column density, N_i , the Doppler parameter, b_i , and the central velocity, v_i —for each component, i , are input as initial guesses and are subsequently freely allowed to vary to minimize the χ^2 goodness of fit. We always start each fit with the smallest number of components that reasonably model the absorption profiles of Si II, and we only include an additional component if it improves the reduced χ^2 .

The results from these fits are summarized in Table 4 in Appendix B. In Figure 3, we show an example of these fits comparing low and high ions. Note that while we only show Si II $\lambda 1260$ in this figure, the fit of the Si II is actually based on three transitions for this absorber ($\lambda\lambda 1260, 1304, 1526$). The fitted composite profiles give us a noise- and contamination-free description of the data. We use these parameters within `spectacle` to construct arrays of flux versus velocity, with the same 2 km s^{-1} binning as the synthetic data. We subsequently use `spectacle` to analyze both the real and simulated spectra using the same tools and assumptions.




Figure 6. Cumulative H I covering fraction for strong systems along each of the three cardinal axes at $z = 2$ and $z = 2.5$ in the two runs. The high-resolution simulation has much stronger H I lines of sight than the standard-resolution simulation.

4. The Effects of Resolution on the Physical State of the Gas

We show slices of density, temperature, and metallicity through the halo at $z = 2$ in both the standard- and high-resolution runs in Figure 4; as in Figure 1, each panel is $1.5 \times 200h^{-1} \text{ ckpc}$ across to fully encompass the forced refinement region. In the standard-resolution simulation it is evident that the resolution directly tracks the density of the gas, while the boundaries of the forced refinement volume stand out clearly in the high-resolution case. While the same structures are likely no longer exactly the same in the two runs owing in part to the “butterfly effect” (Genel et al. 2019), the high-resolution gas is visually much more turbulent, with more well-defined large-scale bubbles and shocks from outflowing gas than in the standard-resolution simulation. The metallicity adds an extra dimension to these flows: in the high-resolution simulation, the high-metallicity outflowing gas is not forcibly “overmixed” with its environs because it is not derefining as it expands and decreases in density. The high metallicity of the outflowing gas is maintained to much larger radii, suggesting that the difficulty of simulations to reproduce the observed pockets of high-metallicity circumgalactic medium gas (Prochaska et al. 2017) owes in part to unnatural overmixing. Likewise, the low-metallicity inflowing gas is able to penetrate to much closer to the galaxy when well resolved than in the standard-resolution simulation. We speculate that these low-metallicity flows are related to the rare low-metallicity Lyman limit systems observed at $z > 2$ (Fumagalli et al. 2011; Lehner et al. 2016) and may be naturally generated in larger simulations with improved circumgalactic (or intergalactic) resolution.

The physical differences shown in Figure 4 translate to quantitative differences in how the dynamic multiphase circumgalactic medium is physically structured and subsequently manifested in different ionic tracers. Though the *average* properties of the circumgalactic medium at the two resolutions are broadly the same, the ionic absorption-line tracers used to diagnose the CGM’s physical structure are highly sensitive to the *distribution* of gas temperature, density, and metallicity (see, e.g., Figure 6 of Tumlinson et al. 2017). The higher levels of both cold *and* hot gas in turn have important consequences for the column density distributions and covering fractions of commonly observed transitions (as illustrated in Figure 7 and discussed in more detail in Section 4.1 of Paper II). This effect is evident in Figure 5, where we

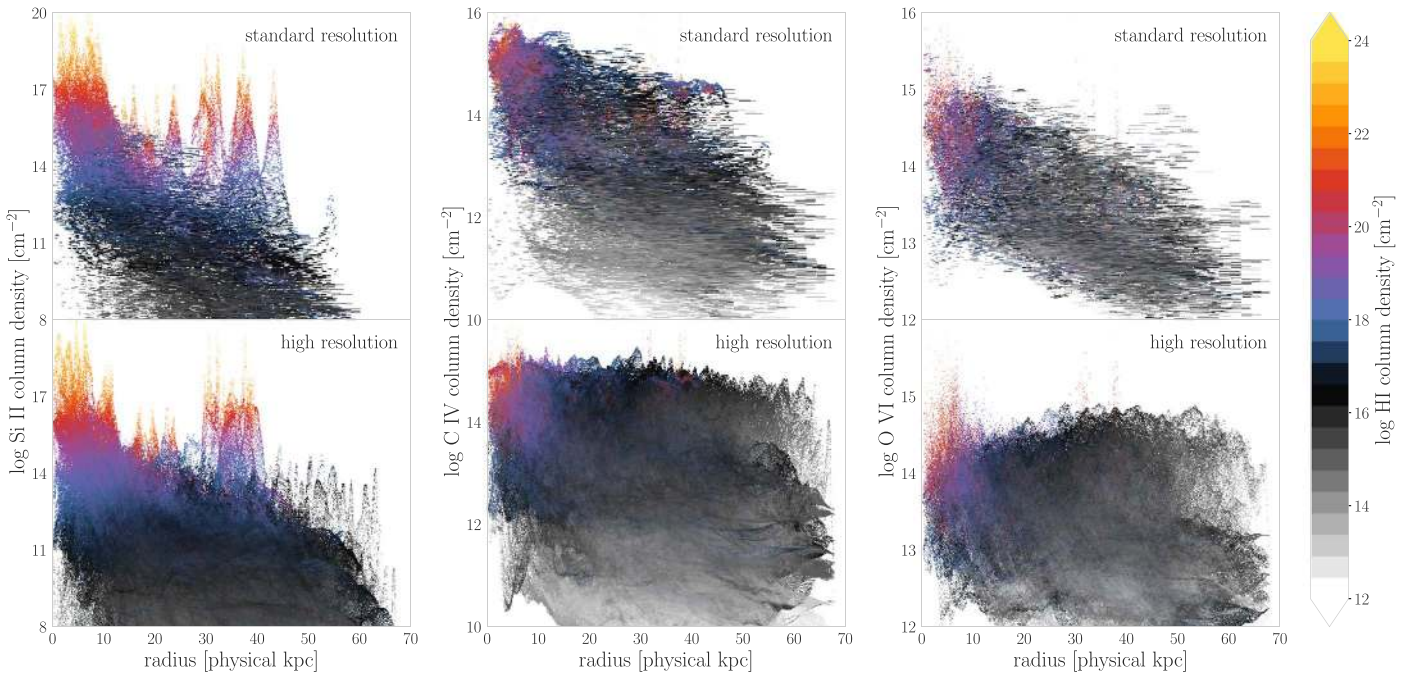


Figure 7. Radial profiles of Si II (left), C IV (middle), and O VI (right) column density for the standard-resolution (top) and high-resolution (bottom) simulations along each of the three cardinal axes at $z = 2$ and $z = 2.5$ in the two runs, color-coded by the most frequent H I column at each given radius and ionic column density. The standard-resolution simulation shows a much steeper radial trend for all species. Si II and H I column are clearly correlated, whereas outside of the main halo the higher ions are largely uncorrelated with H I.

show projected H I, Si II, Si III, Si IV, C IV, and O VI column densities at $z = 2$ in the two simulations; the central halo is in the bottom middle panel. For all species, the high-resolution simulation reaches higher column densities out to higher impact parameter, and small-scale clouds and filamentary features with internally resolved structure are clearly visible in the high-resolution projections.

We plot the cumulative $N_{\text{H I}}$ covering fractions in Figure 6, considering the three cardinal axes at both $z = 2$ and 2.5. We find that the high-resolution simulation produces a higher covering fraction of strong H I absorbers than the standard-resolution simulation, though at $\log N_{\text{H I}} > 18$ this effect is subtle. We summarize the masses of different species of interest within the forced refinement region at $z = 2$ in Table 2; despite the higher prevalence of strong H I absorbers, the total H I mass within the forced refinement region in the high-resolution simulation is $\sim 15\%$ lower than in the standard-resolution simulation. This difference, however, can be mostly explained by the difference in the total amount of gas in the two boxes, which in turn owes to some combination of more star formation and (perhaps) more efficient outflows in the high-resolution simulation.

Using the AREPO code, van de Voort et al. (2019) similarly found that improving circumgalactic resolution increased the amount of H I and the incidence of strong H I absorbers, bringing their simulated galaxies more in line with the COS-Halos $z \sim 0.25$ H I observations (Thom et al. 2012; Tumlinson et al. 2013; Prochaska et al. 2017). This similarity of results given the different hydrodynamic codes, initial conditions, and redshifts suggests that resolving smaller circumgalactic structures generically increases incidence of strong H I absorbers.

While we do not find a significant difference in the column density distribution functions of many of the metal ions, the

spatial distribution of the metal ions and their association with strong H I are significantly different in the two simulations. We show in Figure 7 column density versus radius for the two simulations, again considering projections along each of the three cardinal axes at both $z = 2$ and 2.5. The median radial profiles of metallic ion column densities are steeper in the standard-resolution simulation, especially for the higher ions. The horizontal features in the standard-resolution simulation panels are a manifestation of its coarse resolution: its column densities do not change over relatively large spans of radius. Figure 8 shows the cumulative covering fractions of Si II, C IV, and O VI for both all sightlines and sightlines with $\log N_{\text{H I}} > 16$. Likely owing to the differences in the radial profiles and in the higher prevalence of strong H I in the high-resolution simulation, LLS-selected sightlines in the high-resolution simulation counterintuitively have slightly weaker metal-line absorption than in the standard-resolution simulation, but these differences are well within the typical uncertainties in covering fraction measurements.

Finally, we plot in Figure 9 the Si II, C IV, and O VI column densities versus H I column, independent of radius, for the two simulations. We hypothesize that the more detailed structures seen in the high-resolution panels of Figure 9 are caused by multiple sightlines probing the same underlying physical structure, whereas the patchiness seen in the standard-resolution panels is caused by small numbers of resolution elements being responsible for most of the line-of-sight column density. Unfortunately, much of the parameter space traced by KODIAQ (Section 3) is not probed by this simulation, so our comparison to the data will be somewhat qualitative.¹⁷

¹⁷ Our preliminary tests using the new self-shielding method in Grackle indicate that, as expected, the H I fraction and column densities are higher when self-shielding is taken into account.




Figure 8. Si II, C IV, and O VI covering fraction for both all sightlines (top) and sightlines with $\log N_{\text{HI}} > 16$ (bottom) along each of the three cardinal axes at $z = 2$ and 2.5 in the two runs. While the high-resolution simulation has a similar amount of strong metal lines to that in the standard simulation, these systems are not as confined to the strong H I sightlines. While these differences are somewhat subtle, they are quite noticeable in the comparison plots we discuss in Section 6.




Figure 9. Si II (left), C IV (middle), and O VI (right) vs. H I column densities along each of the three cardinal axes at $z = 2$ and 2.5 in the high-resolution (top) and standard (bottom) simulations, as compared to the KODIAQ data (Section 3). While, broadly speaking, the relationship between strong H I and the metal ions does not change, the high-resolution simulation more smoothly samples the available parameter space.

5. The Effects of Resolution on Circumgalactic Absorption

Circumgalactic observations, like those in most of astronomy, suffer from the projection of the third spatial dimension onto a velocity scale. In a medium like the circumgalactic medium that is highly structured in density, ionization, and metallicity, this projection effect causes a significant loss of diagnostic information. Clouds with large physical separations along the line of sight can appear close to each other in velocity space. Weak components can be concealed by stronger ones. Low-metallicity gas can be screened out by high-metallicity gas. These kinds of issues greatly complicate the interpretation of observed absorption-line profiles and are a large source of systematic uncertainty in conclusions about the circumgalactic medium. A major advantage of using a simulated universe to help interpret real data is that in simulations we can deproject the observable spectra and investigate in detail how complex,

multiphase absorption profiles are produced. We show how dynamic multiphase circumgalactic gas translates to absorption-line systems in Section 5.1, arguing in Sections 5.2 and 5.3 that the bulk of the absorption is resolved in our high-resolution simulation but not in the standard-resolution simulation.

5.1. Velocity–Line-of-sight Phase Space

Figures 10 and 11 show diagrams mapping the 3D spatial fields of temperature and metallicity to the line-of-sight velocity phase space and resulting absorption-line systems for example sightlines in the high- and standard-resolution simulations, respectively. We show further illustrative examples in Appendix C. In each diagram, the shaded panels at left show a “core sample” extending $200h^{-1}$ ckpc along the line of sight and $\pm 10h^{-1}$ ckpc to either side of the ray in two orthogonal axes. These core samples place the infinitesimally

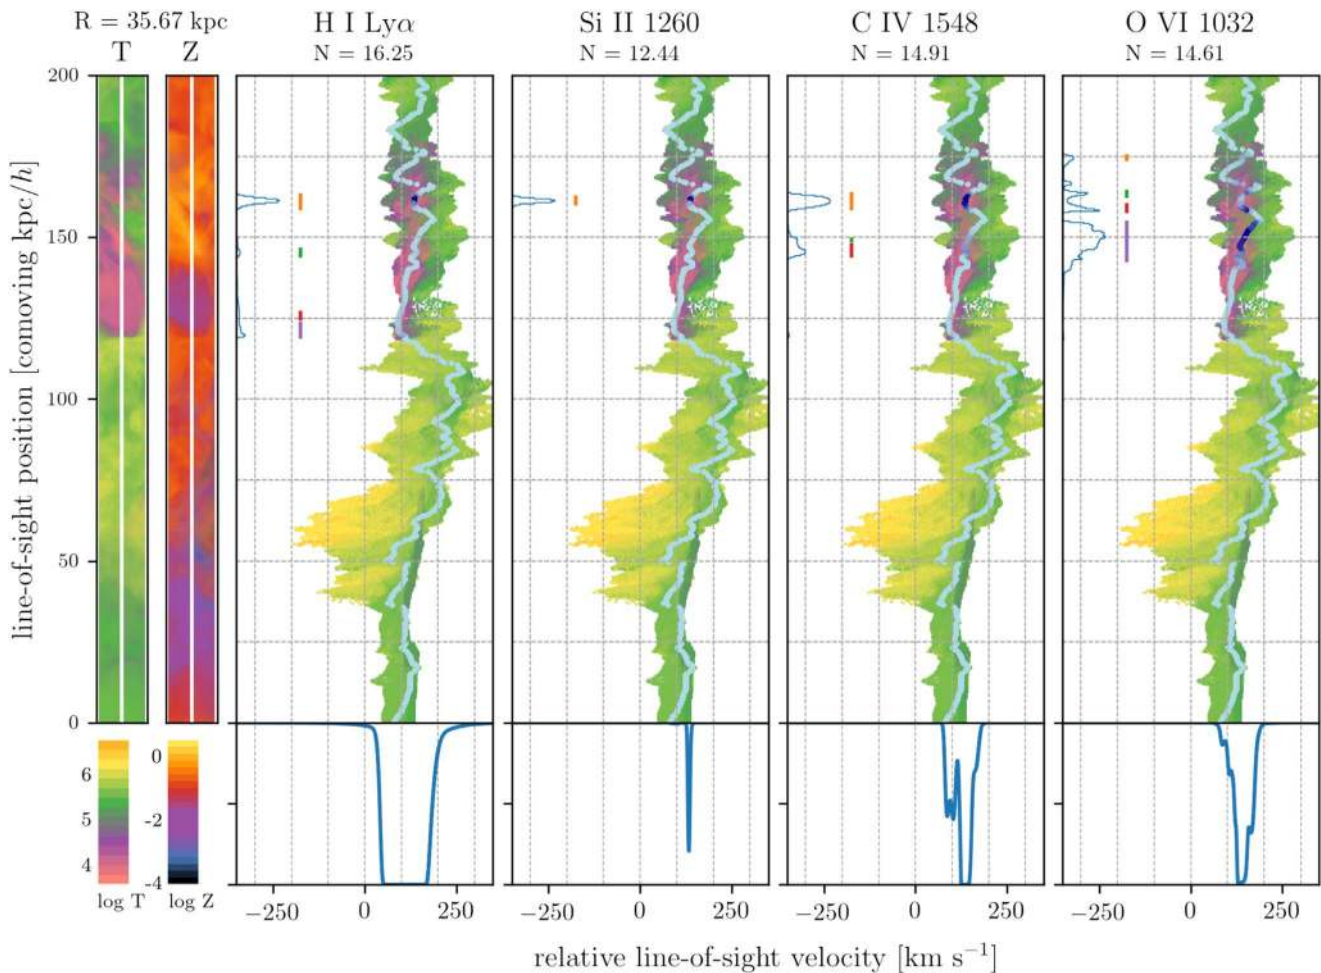


Figure 10. How multiphase gas becomes multi-species absorption, here for a sightline through the high-resolution simulation at $z = 2$. The two thin windows at left render the temperature and metallicity in a rectangular prism $\pm 10 h^{-1}$ ckpc to either side of a $200 h^{-1}$ ckpc sightline; the impact parameter with respect to the central halo is given in the upper left corner in physical kpc. As in Figure 2, the color at each position in these panels is given by the most common temperature of a cell at that location. In the four main panels, the same cells are rendered again with the temperature shading but with the line-of-sight velocity on the lower axis. The blue traces denote the path of the ray through the velocity–line-of-sight phase space, with darker-blue regions denoting higher ionization fractions in each ion; the “ $N =$ ” at the top of these panels gives the total column density for that ion along the sightline in $\log \text{cm}^{-2}$. The ionization fractions are also shown as the interior left histograms (arbitrarily scaled) of each of these panels, with identified “clouds” marked with colored bars. The spectra in the bottom panels are aligned with the same velocity axis and have had no rebinning or LSF applied; the dashed gray lines are at fixed 100 km s^{-1} intervals to guide the eye. See Section 5.1 for a more thorough description.

narrow sightline (shown by the white lines) in the context of the structures it intercepts.

The four main panels render the same cells from the core sample, color-coded by temperature, with the horizontal axis now denoting the line-of-sight velocity along the ray, v_{los} . The most obvious behavior seen here is that structures at many densities, temperatures, and metallicities extending over the $200 h^{-1}$ ckpc of the ray are still quite tightly constrained in velocity space, with most of the gas ranging over $\pm 200 \text{ km s}^{-1}$. Each of these phase-space panels corresponds to an individual ion (here H I, Si II, C IV, and O VI). The thin trace marked in shades of blue (a color absent from the temperature scale, to provide contrast) shows the locations in phase space occupied by the cells directly along the line of sight, enabling one to immediately read off the interplay between the physical position and line-of-sight velocity for a given sightline. The darkness of blue shades along these traces is in proportion to the number density of the ion species of interest in that cell; the darkly shaded cells therefore generally line up with the absorption peaks in the simulated spectra shown in the bottom panels. In the left of each phase-space panel, there is a thin blue

histogram showing the distribution of each ion’s number density (arbitrarily scaled for visibility). Overlaid on these are the colored bars marking the extent of the identified “clouds”;¹⁸ we explore the properties of these clouds in detail in Sections 5.2 and 5.3. Finally, we mark the locations of the spectral minima analyzed in Section 6 in the bottom panels. These panels are arranged such that reading over from the core samples to the phase-space panels and down to the spectra ties together the temperature, metallicity, ionic number density, and absorption signature of each cell along the line of sight.

These diagrams enable us to immediately draw several conclusions about the simulated circumgalactic medium. First, the strongest absorption is mostly confined to within $\pm 200 \text{ km s}^{-1}$ of the systemic velocity, with the hot gas having the highest velocity dispersion and the cold gas generally the lowest. We can see also that the strongest absorption—those components that make detectable impressions on the spectrum—are often confined to only a few physical locations, with the

¹⁸ Though these absorbing regions are clearly continuously fluctuating regions, we will for simplicity refer to them as “clouds.”

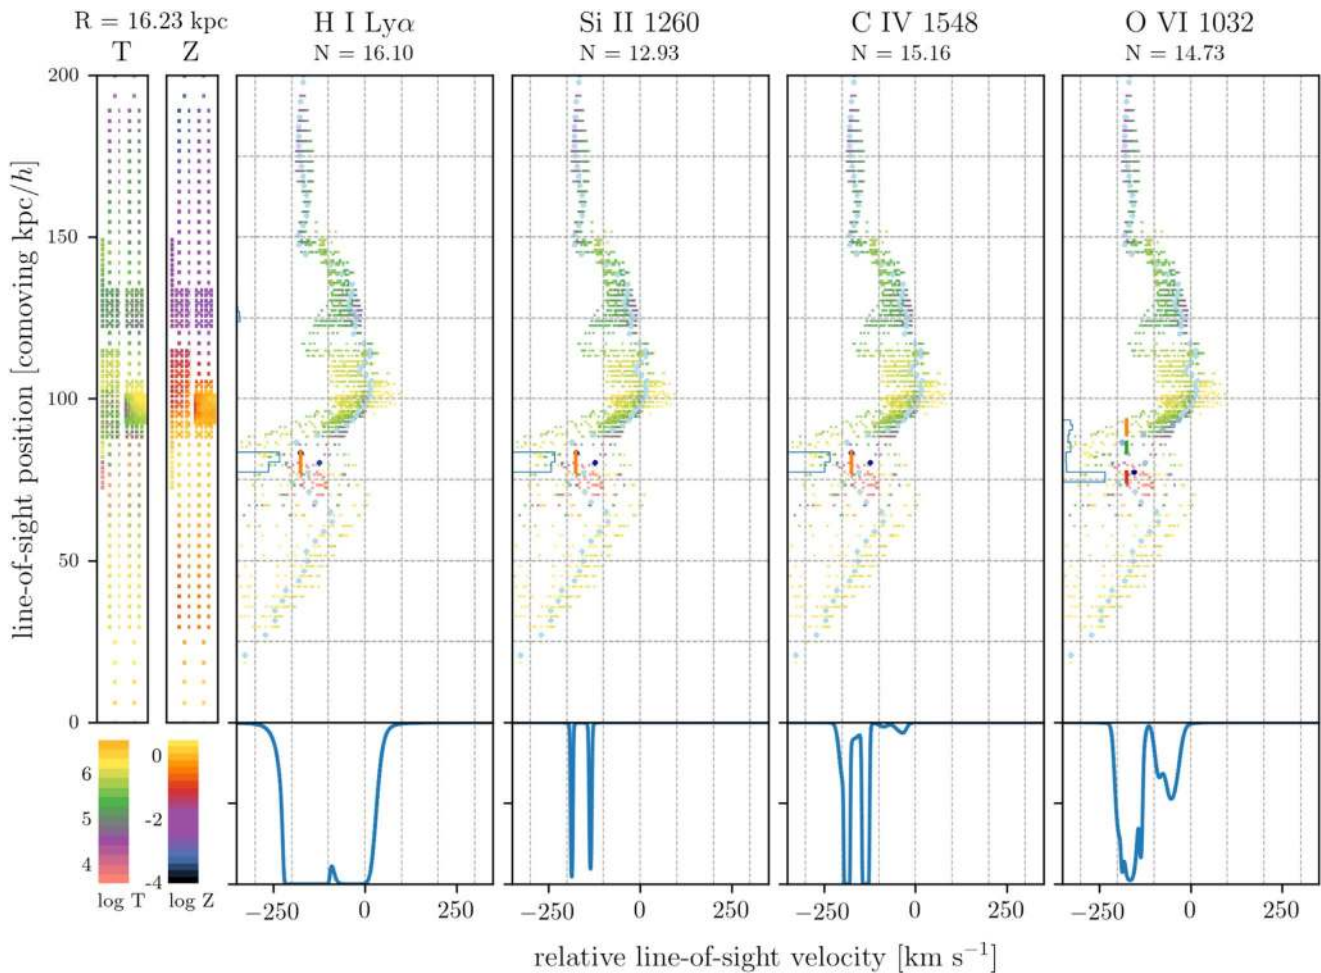


Figure 11. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.5$. The contribution of individual cells is clearly evident in the temperature and metallicity renderings.

higher ions arising from gas that is more spatially extended than the lower ions. Generally, these locations correspond to where the cells have significant metal enrichment (with the notable exception of the H I-traced gas) and a temperature that is suitable for the species to have a significant ionization fraction.

Contrasting Figures 10 and 11, we immediately notice how sparsely sampled both the underlying physical *and* velocity spaces are in the standard-resolution simulation. As expected from Section 4, the high-resolution simulation shows much more variation along the line of sight than in the standard-resolution simulation. As we show in Section 5.2, the bulk of the gas giving rise to the “observed” absorption in the standard-resolution simulation is often from only a few resolution elements, while these structures are more fully resolved in the high-resolution simulation.

5.2. Resolving the Absorbers Spatially

To quantify our comparisons of absorption structures generated from simulations at different resolutions, we define a relatively unbiased, empirical definition of what constitutes a “cloud” that is independent of the simulated absorption. For each ion, we find the threshold in the ionic number density n_{ion} above which 80% of the total sightline column density is produced. Using all cells with number densities above this

threshold value, we find the positions along the line of sight where n_{ion} crosses this threshold on its way up and then down again to define a single cloud. Between up-and-down threshold crossings, n_{ion} can take on any value, and so the clouds can be complex in shape and size. While our choice of 80% of the total sightline column density is somewhat arbitrary, we find that much larger percentages often lead to the near entirety of the sightline contributing to single clouds, which, while interesting, is less useful as a diagnostic tool. The identified clouds and their extents are denoted with the vertical colored bars next to the line-of-sight histograms in the velocity phase-space plots (Figures 10 and 11).

The immediate conclusion from this analysis is that only a few clouds dominate the simulated column density of each sightline, and that the number of distinct clouds responsible for 80% of the column density increases as we go up the scale of ionization potential from H I to O VI. For example, though the ions shown in Figure 10 are mostly at the same line-of-sight velocity, the H I gas is largely confined to a single ~ 5 kpc cloud, while the O VI-traced gas is spread out over ~ 35 kpc, most of which is not directly associated with the H I-bearing gas. Moreover, the low-ionization Si II largely (though not entirely) traces the H I, while the intermediate C IV ion has one cloud largely associated with the H I and Si II and one cloud overlapping with the most predominant O VI-bearing cloud. This finding by itself calls into question the general tendency to

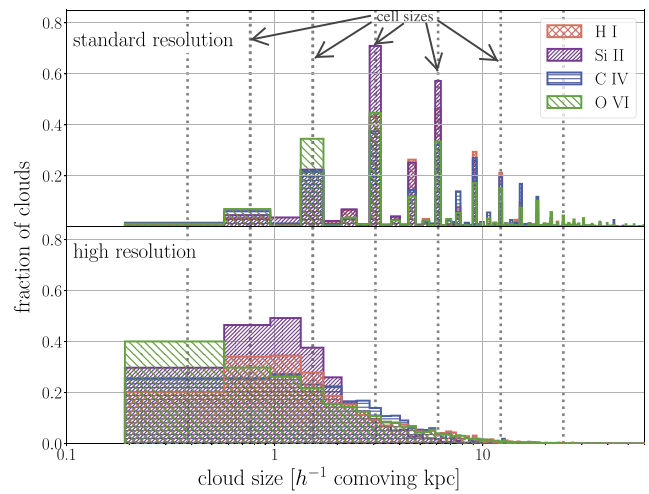



Figure 12. Cumulative (left) and differential (right) distributions of H I (pink), Si II (purple), C IV (blue), and O VI (green) cloud lengths in the standard-resolution (top) and high-resolution (bottom) simulations. We plot the differential distributions in comoving coordinates to show that in the standard-resolution simulation cloud sizes tend to cluster around the cell sizes (vertical gray dashed lines in the right panels), indicating that the clouds are unresolved.




Figure 13. Cumulative distribution of the number of resolution elements giving 80% of the absorption along the line of sight per absorber for H I (pink), Si II (purple), C IV (blue), and O VI (green).

treat observed absorbers as single objects with homogeneous conditions.

In Figure 12, we show the distribution of the cloud sizes in the standard and forced refinement runs for H I, Si II, C IV, and O VI. The dashed gray lines in the differential histogram denote the cell sizes at different refinement levels found in the circumgalactic medium in the standard-resolution simulation: the gas contributing to the bulk of the absorption in this simulation is clearly often arising from single resolution elements and is therefore unresolved. In contrast, the highly resolved circumgalactic medium has a more power-law distribution of cloud sizes, as would be more generally expected if the cloud size is set by turbulence (in analogy to the mass distribution of star-forming clouds driven by turbulence; Padoan & Nordlund 2002). Figure 13 further quantifies the differences in cloud sizes by showing the distribution of the numbers of cells along the line of sight per cloud in the two simulations. In the standard-resolution simulation, $\sim 75\%$ of the sightlines with $\log N_{\text{H I}} > 16$ have only ≤ 5 cells giving rise to $> 80\%$ of their H I column density,

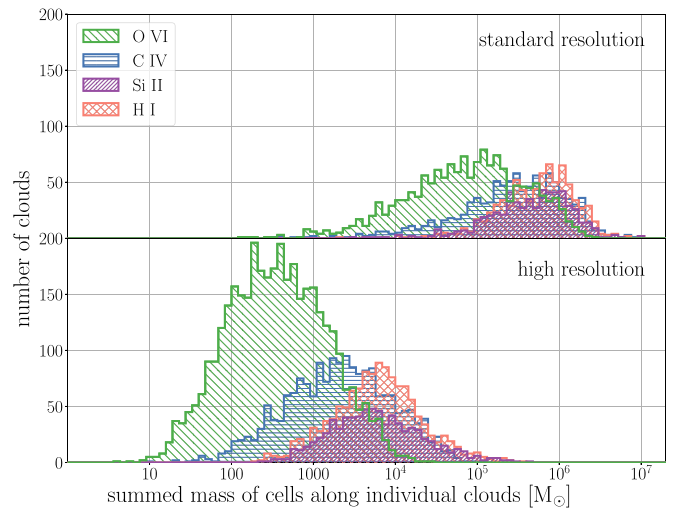


Figure 14. Distributions of 1D cloud masses for H I (pink), Si II (purple), C IV (blue), and O VI (green) in the standard-resolution (top) and high-resolution (bottom) simulations.

whereas this fraction is $< 20\%$ in the high-resolution simulation. This difference is still stark for the highly ionized O VI: $\sim 17\%$ of the sightlines with $\log N_{\text{O VI}} > 13$ have $> 80\%$ of this column density in only ≤ 5 cells, whereas this fraction is $\sim 1\%$ in the high-resolution simulation.

In Figure 14, we show how the improvement in *spatially* resolving the absorbing clouds translates back into the *mass* resolution commonly enforced by Lagrangian codes. In the standard-resolution simulation, we find that most of the absorbing clouds have masses of $\gtrsim 10^4 M_{\odot}$, consistent with the masses of resolution elements in FIRE (Hopkins et al. 2014, 2018), Illustris (Vogelsberger et al. 2014), or EAGLE (Schaye et al. 2015). In the high-resolution simulation, however, the absorbing clouds typically have masses two to three orders of magnitude lower, suggesting that the shift to lower cell masses seen in Figure 2 results in lower-mass structures being resolved. The more highly ionized O VI tends to arise in larger clouds, more clouds along the line of sight, and much less mass per cloud than the lower-ionization metal

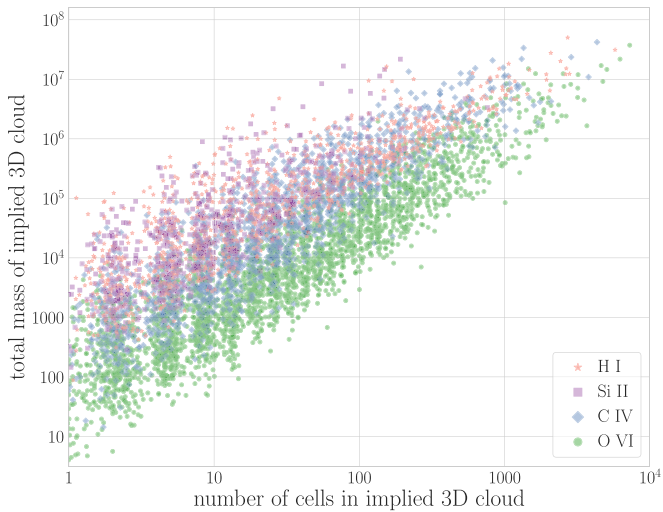


Figure 15. Implied 3D cloud masses (see Equation (2)) of H I (pink), Si II (purple), C IV (blue), and O VI (green) vs. the implied 3D number of cells in the high-resolution simulation. Small random offsets have been applied to the cell numbers to make the trends easier to see. At the typical mass resolution of most current Lagrangian simulations ($\sim 10^4$ – $10^6 M_\odot$), FOGGIE resolves most absorbers with tens to hundreds of resolution elements.

species, which is roughly what is expected from the more highly ionized species tracing lower-density gas. This combination of results highlights the fact that to adequately resolve low-density gas—even over large spatial scales—low-mass scales must also be achievable: in the high-resolution simulation, $>75\%$ of the clouds responsible for $>80\%$ of the absorption in the $>10^{13} \text{ cm}^{-2}$ O VI sightlines have 1D cloud masses of $<1000 M_\odot$ (and $\sim 20\%$ of the clouds have masses $<100 M_\odot$).

We calculate an estimated 3D cloud mass, m_{cloud} , by calculating the average density along the sightline and assuming spherical geometry, that is,

$$m_{\text{cloud}} = \frac{4}{3}\pi \left(\frac{\text{cloud length}}{2} \right)^3 \times \sum_{\text{cells in cloud}} \frac{\text{cell mass}}{\text{cell volume}}. \quad (2)$$

Likewise, for a uniform grid, the estimated number of cells in a 3D cloud is $(\pi/6) \times (\text{the number of cells in the 1D cloud})^2$. We plot the 3D cloud masses versus estimated number of cells in the 3D cloud for the high-resolution simulation in Figure 15. As the cell sizes are effectively uniform within the forced refinement region, the trend in cloud mass for a fixed number of cells between the different ions is a tracer of the underlying density, with the lower-ionization species being traced by higher-density gas, as expected. We also find that for 3D cloud masses of $\gtrsim 10^5 M_\odot$, FOGGIE resolves the structures giving rise to most circumgalactic absorption with $\gtrsim 10$ – 1000 resolution elements.

5.3. Kinematically Resolving Circumgalactic Absorbers

In real absorption-line observations, we are unable to measure distance along the line of sight; instead, we attempt to disentangle the physical positions of absorbing clouds from their line-of-sight velocity. While coarsely resolved simulations can be used to understand observed large-scale kinematic trends (Turner et al. 2017), we are interested here in what sets

the kinematics within individual absorption complexes. As can be seen from Figures 10 and 11 (and the other examples in Appendix C), absorbing gas that is many tens of kiloparsecs apart may be coincident in line-of-sight velocities. This complicates the interpretation of single absorption-line “components” and the comparison of absorption seen at the same velocity in species with different ionization levels. Likewise, a flow along the line of sight can imprint itself on an absorption feature by spreading nearly cospatial gas out over a broad range in velocity. Observationally, this means that bulk motions contribute significantly to the nonthermal widths of broad absorption features, and it is generally the hotter, more thermally broadened gas that has higher relative velocities.

It is the latter case that is particularly tricky to simulate: the underlying velocity field is discretized by the sampling of the individual resolution elements but then smeared out by thermal broadening when creating the mock spectrum. While a high-resolution simulation may better resolve the dynamic structure of the circumgalactic medium, a low-resolution simulation may have more kinematic components in the *spectra* because of single cells giving rise to individual absorption-line components (see Figure 16), whereas the same feature when fully resolved may be blended out into one larger, more complicated feature. The Si II and C IV panels of Figure 11 show good examples of this phenomenon: though the identified “clouds” in these panels comprise multiple cells, it is clear that the sharp absorption features are arising from individual cells—and that if the underlying flow was better resolved, the resulting spectrum would be more blended. We explore the impacts of resolution on observable kinematic structure and how they compare to the KODIAQ data in Section 6, but we first examine here the interplay between the number of kinematic components in an “observed” spectrum and how well resolved the underlying clouds are.

Consider an individual 1D cloud composed of some number of cells, N_{cells} . For a given ion, the relevant average temperature of the cloud is the ion-density-weighted temperature

$$\bar{T}_{\text{ion,cloud}} = \frac{1}{N_{\text{ion,cloud}}} \times \sum_{\text{cells}} T_{\text{cell}} N_{\text{cells}}, \quad (3)$$

where $N_{\text{ion,cell}}$ is the column density of the ion in that cell and $N_{\text{ion,cloud}}$ is the column density of the ion in the entire cloud. From here, we define the average velocity dispersion from thermal motions per cell in the cloud as

$$\sigma_{\text{th}} = \sqrt{\frac{k\bar{T}_{\text{ion,cloud}}}{m_{\text{ion}}}}, \quad (4)$$

where k is the Boltzmann constant and m_{ion} is the mass of the ion. Taking the velocity width of the cloud Δv_{cloud} to be the difference between the maximum and minimum line-of-sight velocities of cells in the cloud, we can define a cloud velocity sampling as

$$\text{cloud velocity sampling} = \frac{\Delta v_{\text{cloud}}}{\sigma_{\text{th}}}. \quad (5)$$

This cloud velocity sampling then gives us an estimate of how many *kinematic* resolution elements are required to resolve a cloud. In order to meaningfully compare the cloud velocity sampling to N_{cells} , we define a flow to be “resolved” if the average spacing between Gaussian profiles results in a fully blended profile, i.e., if there is only one resulting peak with no

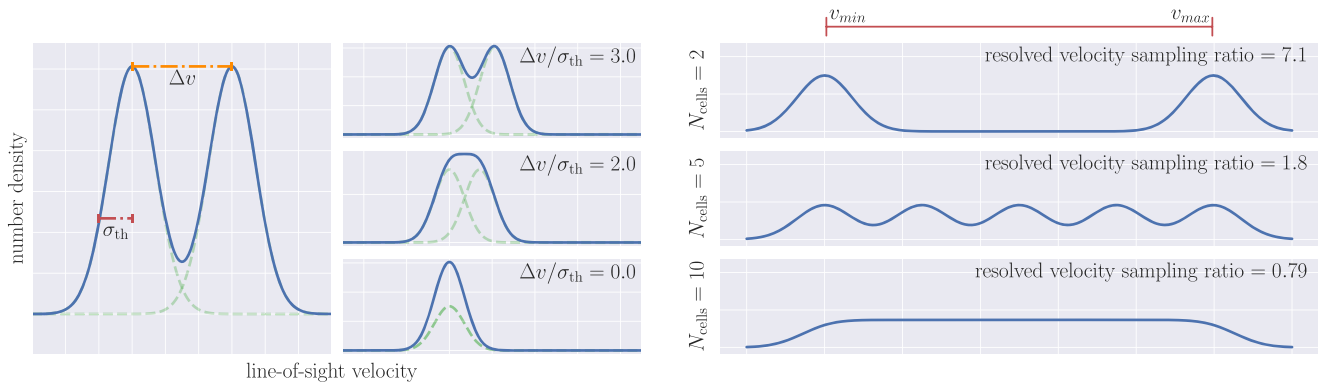


Figure 16. Left: how the cloud velocity width Δv_{cloud} and average cell thermal broadening σ_{th} correspond to the cloud velocity sampling $\Delta v_{\text{cloud}}/\sigma_{\text{th}}$ (Equation (5)). In the case of two components, systems with $\Delta v_{\text{cloud}}/\sigma_{\text{th}} \leq 2$ show only one extremum, while less well-sampled systems are less blended. Right: how better-resolved velocity flows translate into smaller resolved velocity sampling ratios (Equation (6)) for a range of N_{cells} ; here, the total area under the curves in each panel is the same.




Figure 17. Cumulative distributions of the resolved velocity spacing ratio (Equation (6)) in the standard-resolution (top) and high-resolution (bottom) simulations.

troughs in the combined profile. For identical Gaussian profiles, this occurs when the spacing between centroids is $\lesssim 2\sigma$, as we illustrate in Figure 16. Extrapolating to many cells, we define the “resolved velocity sampling ratio” as

$$\text{resolved velocity sampling ratio} \equiv \frac{\text{cloud velocity sampling}}{2(N_{\text{cells}} - 1)}. \quad (6)$$

Here a cloud that is better sampled (and thus better kinematically resolved) has a lower resolved velocity sampling ratio, whereas clouds with ratios > 1 likely have individual cells as the predominant contributor to individual absorption-line components, as illustrated in the right panels of Figure 16. (In cases with $N_{\text{cells}} = 1$ [and thus $\Delta v \equiv 0$], we set the resolved velocity sampling ratio to 10^6 , as these clouds are clearly unresolved.) As an example, consider the Si II absorbers in Figures 10 and 11. In the high-resolution case, there is a single Si II cloud composed of five cells with $\Delta v = 3.17 \text{ km s}^{-1}$ (and length $1.9 h^{-1} \text{ ckpc}$) and $\sigma_{\text{th}} = 1.8 \text{ km s}^{-1}$, yielding a cloud velocity sampling of 1.7 and a resolved velocity sampling ratio of 0.21. The velocity structure of this cloud is clearly resolved. In contrast, the example standard-resolution Si II absorber is visibly unresolved. Quantitatively, this absorption arises from a

single “cloud” with only two cells separated by 51 km s^{-1} (and $6.1 h^{-1} \text{ ckpc}$) with $\sigma_{\text{th}} = 2.2 \text{ km s}^{-1}$, yielding a resolved velocity sampling ratio of 11.7.

Figure 17 shows the cumulative distributions of the resolved velocity sampling ratio for the two simulations. We draw several conclusions from this analysis. First, for the metallic ions in both simulations, the kinematics of the lower-ionization gas are less well resolved than for the more highly ionized gas, an effect that here is likely further exaggerated by the relatively high atomic mass of silicon (which enters via σ_{th} as defined in Equation (4)). Conversely, the low atomic mass of hydrogen (and our high $\log N_{\text{HI}}$ selection) leads to relatively well-resolved HI kinematics. Second, while the internal kinematics of absorbing clouds is better resolved in the high-resolution simulation, there is still a significant fraction of clouds that likely have individual cells dominating individual absorption-line components, especially for Si II: 39% of Si II clouds in the high-resolution simulation and 66% of sightlines in the standard-resolution simulation have resolved velocity spacing ratios > 1 . We find a large improvement in the resolution of C IV kinematics, with 55% of C IV clouds with resolved velocity spacing ratios < 1 in the standard-resolution simulation increasing to 82% in the high-resolution case.

In the observable universe, we of course cannot directly measure number-density-weighted velocity profiles, but instead we observe number density weighted by the oscillator strength and wavelength of the relevant atomic transition. The effects of saturation in particular wash out the underlying velocity structure, as is especially obvious when using the locations of low-ionization components to guide fitting saturated HI lines. Nonetheless, as the velocity phase diagrams shown in Figures 10 and 11 (and in Appendix C) demonstrate, how well resolved—or not—the underlying gas flow is can have a dramatic impact on the kinematic structure in the resulting absorption-line spectrum. We explore these effects and how they complicate comparisons to the KODIAQ data in Section 6.

6. Quantifying Kinematic Structure and Comparisons to KODIAQ

We turn now to comparing the kinematic structure of the synthetic absorbers from the two simulations to observed absorption-line systems in high-resolution spectra at $2 \lesssim z \lesssim 2.5$.

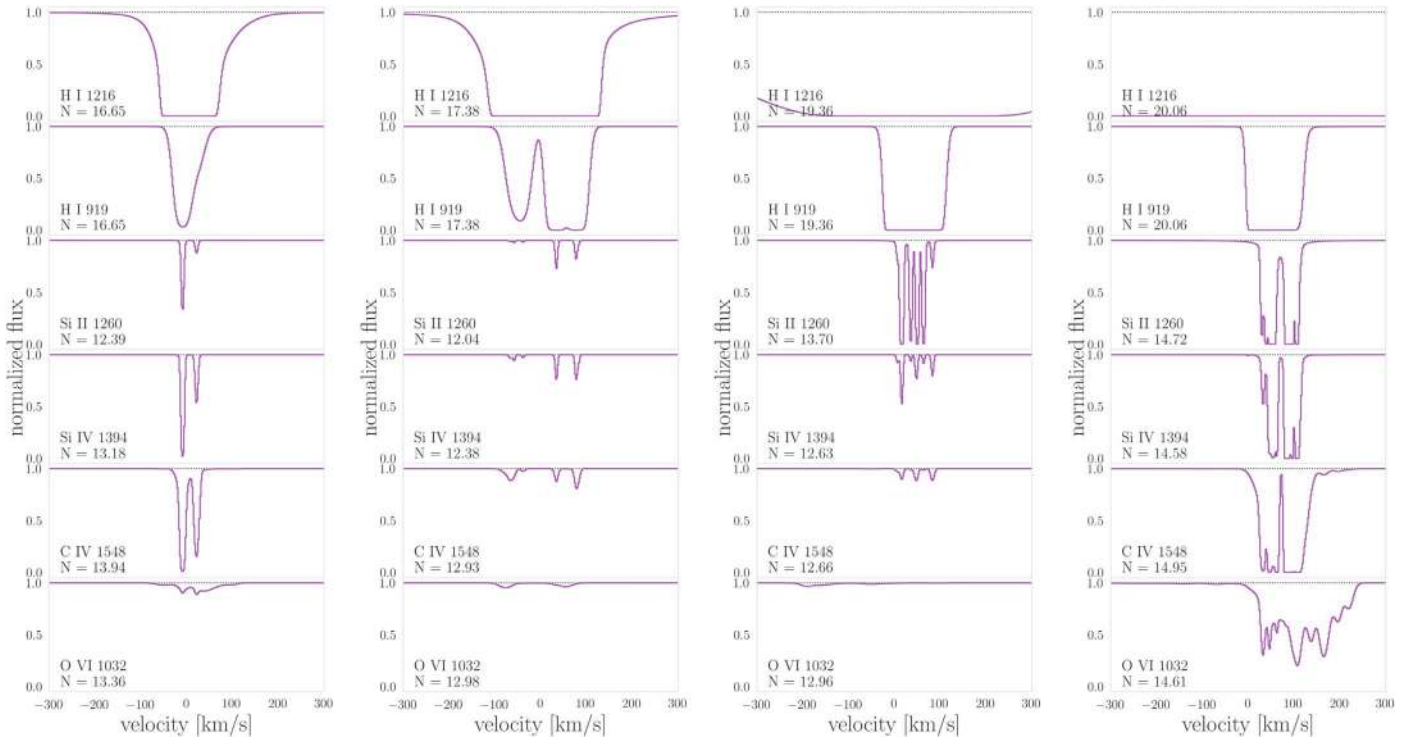


Figure 18. Typical absorbers from the standard-resolution simulation, with $\log N_{\text{H I}}$ increasing to the right. Each panel gives the species and rest-frame wavelength and the log column density in cm^{-2} as “ $N =$ ”; we include H I $\lambda 919$ to highlight some of the structure that is lost owing to the saturation of Ly α .




Figure 19. Typical absorbers from the high-resolution simulation, with $\log N_{\text{H I}}$ increasing to the right. Each panel gives the species and rest-frame wavelength and the log column density in cm^{-2} as “ $N =$ ”; we include H I $\lambda 919$ to highlight some of the structure that is lost owing to the saturation of Ly α .

The larger spreads of physical properties emergent in the higher-resolution simulation also correspond to larger spreads in the line-of-sight velocities. Because of the changes to the underlying structure of the gas, we compare sightlines in the

two simulations at fixed H I column density rather than fixed position in the box. We show in Figures 18 and 19 sample spectra, at fixed H I column density, in the two runs, as compared to a KODIAQ set of absorbers at similar H I column.

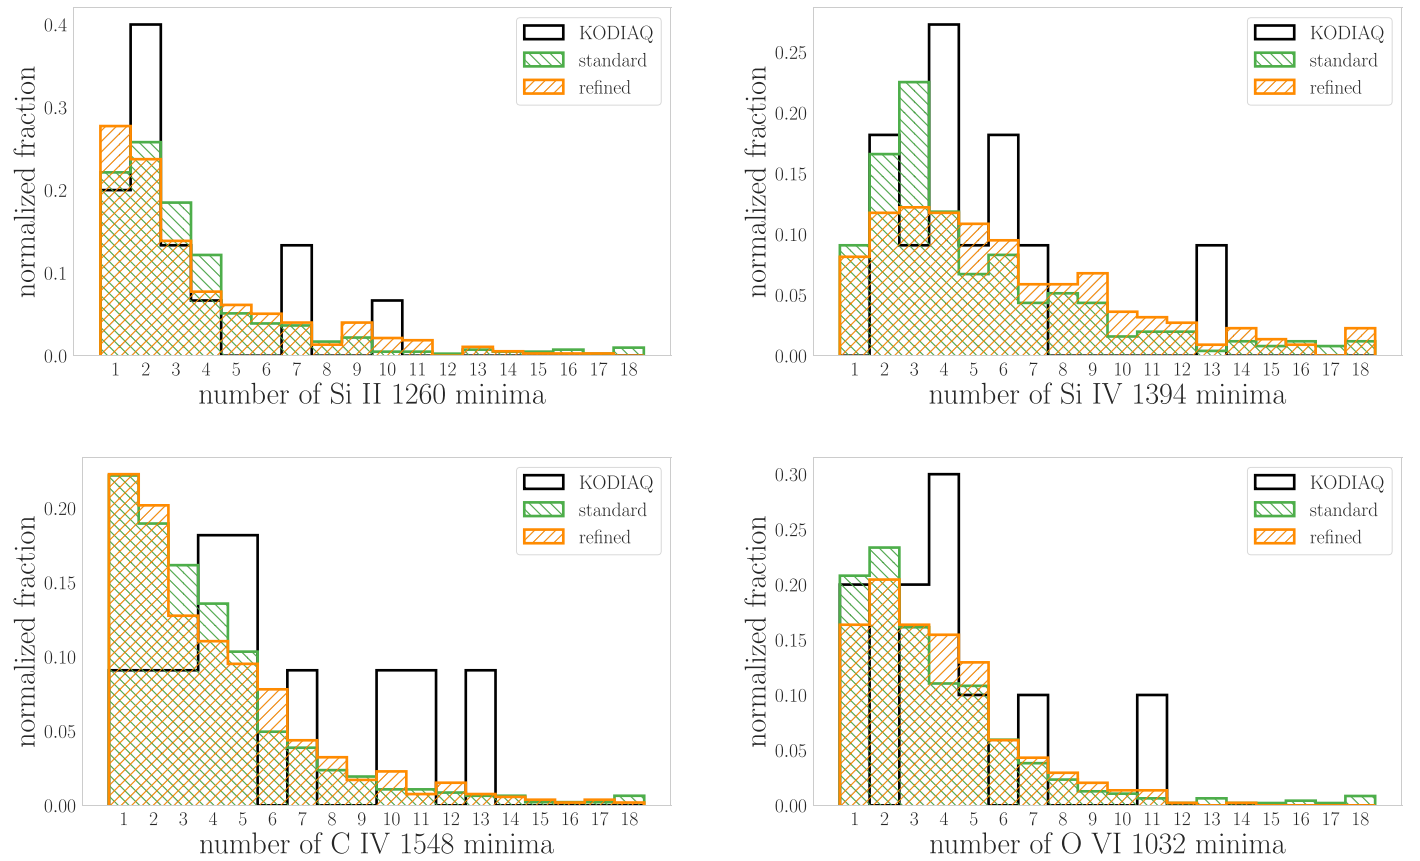


Figure 20. Relative number of minima with flux < 0.95 in LLS-selected synthetic sightlines chosen to have at least the minimum column as seen in the KODIAQ data.

By eye, the high-resolution simulation yields a much richer kinematic structure. First, the absorbers from the high-resolution simulation are much more likely to have many components, whereas the standard-resolution simulation yields absorbers that are more frequently a single component. Second, when the standard-resolution simulation does have multiple components, they are very often unblended, with the flux reaching or nearly reaching the continuum between minima. While the first effect may be in part because of the different covering fractions (and conditional covering fractions) in the two simulations, both effects are likely caused by the artificial discretization of velocity from the larger resolution elements in the standard-resolution simulation.

By eye, the spectra from the high-resolution simulation seem to have many more components than those from the standard-resolution simulation. Moreover, when the standard-resolution spectra *do* have many components, they tend to be fairly unblended, i.e., they have a comb-like structure with distinct lines that go back up near the continuum before absorbing again. We postulated in Section 5.3 that these unblended components are predominantly caused by single grid cells, and while these cases do occur in spectra generated from the high-resolution simulation, they are less frequent.

When the absorption-line systems in the KODIAQ data have many components, they tend to be more blended than even found in the high-resolution simulation. We stress that the comparisons for this generation of the FOGGIE simulations to the KODIAQ data should be taken only qualitatively: the real spectra are probing pathlengths of $\gg 200h^{-1}$ ckpc and a much wider range of galaxy masses, environments, and impact parameters than offered by our single halo. As we show in

Figure 8, the column density parameter space traced by our synthetic absorbers and the KODIAQ data is partially non-overlapping—there is simply a small parameter space of LLS gas probed by this one halo. Also, as we analyze the synthetic spectra without adding noise, we are able to “detect” more minima in these spectra than we would be able to fit for in even the highest-S/N KODIAQ spectra. Nonetheless, we find it instructive to compare the synthetic data to the real universe where possible.

In Figure 20 we show the distribution of the number of minima in the two simulations for Si II, Si IV, C IV, and O VI as compared to the KODIAQ sample. While the differences between the standard- and high-resolution simulations are subtle, there are generally more systems with $4 \leq N_{\min} \leq 10$ in the high-resolution simulation than in the standard-resolution simulation. Two-sided Kolmogorov–Smirnov tests rule out the Si IV distributions as being drawn from the same sample at a 9.6×10^{-3} level; the other distributions cannot be ruled out as being drawn from the same parent sample.

The distributions of N_{\min} in Figure 20 are complicated by the fact that higher column density systems tend to have more minima, as we show in Figure 21, and that the high-resolution simulation has lower metal column densities for LLS systems (Figures 8 and 9). Figure 21 also shows that, at fixed column density, the high-resolution simulation tends to have more minima than the standard-resolution simulation, and for Si II and C IV, the high-resolution simulation is in better agreement with the KODIAQ data.

Finally, in both simulations we find that the absorbers with the highest number of minima are preferentially found at low impact parameters, as shown in Figure 22. This prediction is

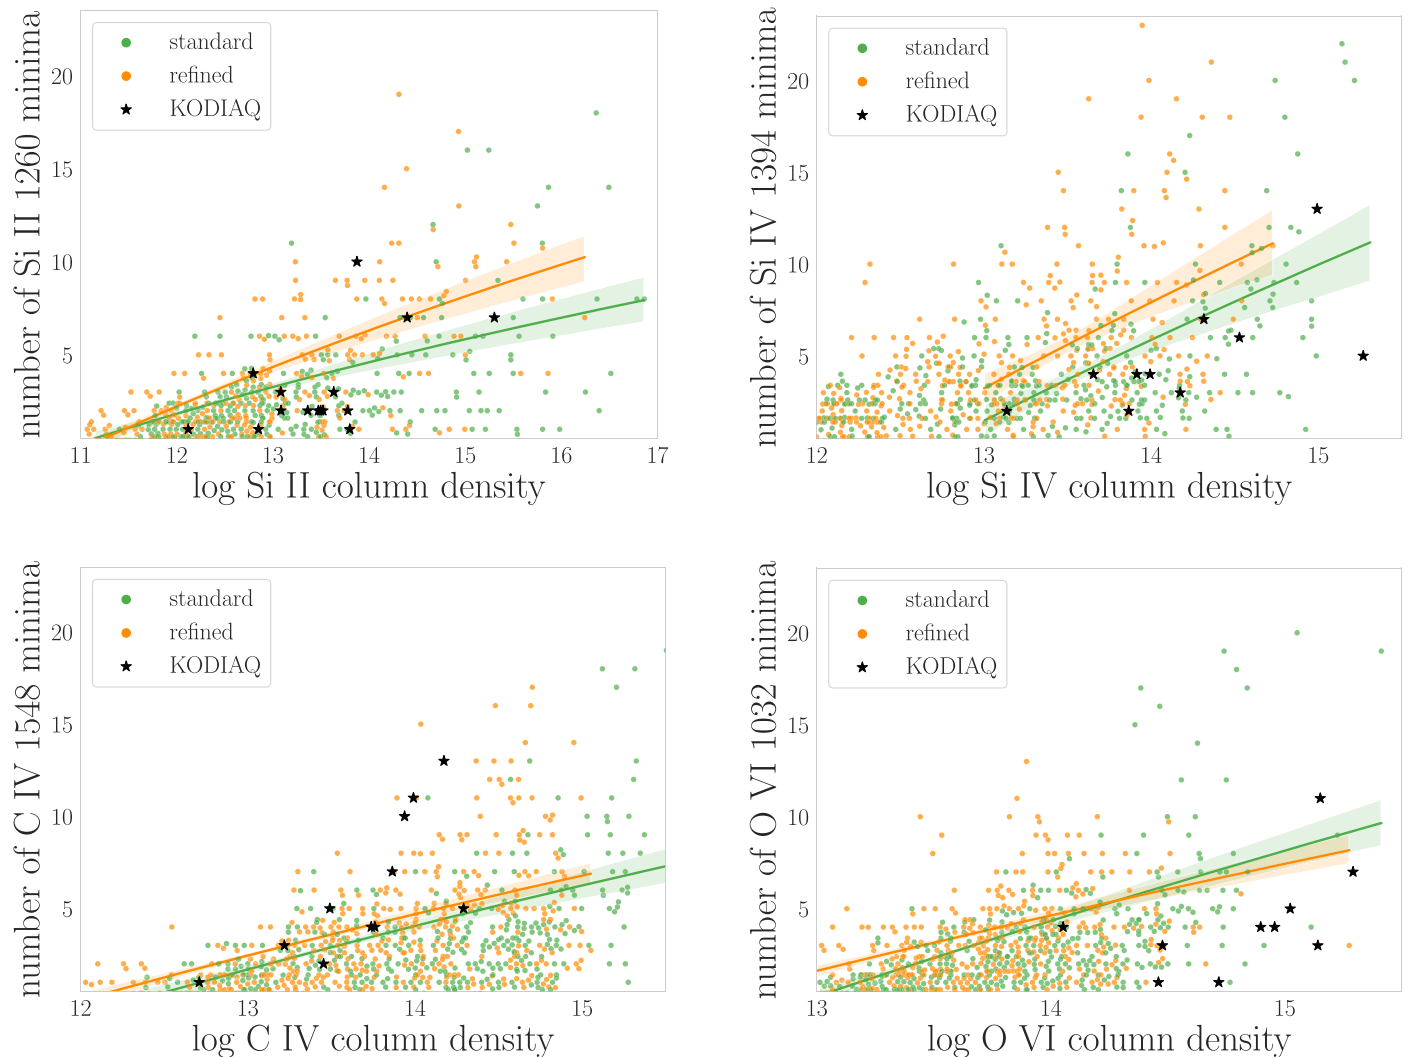


Figure 21. Number of minima with flux < 0.95 for LLS sightlines vs. column density. The lines are linear regressions to guide the eye; in general, higher column densities have more absorption minima, and for most ions at most column densities, the high-resolution “refined” simulation has more minima than the standard-resolution simulation.

consistent with the COS-Halos data of L^* galaxies at $z \sim 0.25$ (Werk et al. 2014, 2016; cf. Borthakur et al. 2016). As we gather the associated galaxy information for the KODIAQ absorbers in the future, it will be interesting to see whether the more kinematically complex absorbers have galaxies at small impact parameters.

7. Conclusions and Future Directions

Testing theories of galaxy evolution with inter- and circumgalactic absorption-line spectroscopy is a long-standing goal of extragalactic astrophysics that has motivated many substantial investments in observing time and computing resources. Both observations and physical arguments point to the circumgalactic medium having physical structure on scales smaller than those typically resolved by cosmological hydrodynamic simulations aimed at understanding galaxy evolution. Simulations using a wide range of numerical techniques applied at a range of scales have addressed this problem with mixed success. By forcing the circumgalactic resolution to subkiloparsec scales, we have shown here that the gas structures responsible for most of the circumgalactic absorption

are resolved—and that this is not the case in a standard-resolution density-refined simulation.

Many studies have assumed that readily observable quantities and statistics derived from them—such as column densities, covering fractions, and system-level velocities—are adequate to provide detailed constraints on models. Under this assumption, feedback prescriptions and other input physics are varied to match the adopted data, and recovery of the observables is taken as an indication that the input physics is “right” (e.g., Ford et al. 2013; Hummels et al. 2013; Oppenheimer et al. 2018). FOGGIE shows that the resolution of a simulation has a major effect on the physical configuration of the circumgalactic medium gas—cloud sizes, masses, and kinematics—even if the underlying “subgrid” prescriptions do not change. Figure 8 shows that FOGGIE’s covering fractions of several commonly observed ions do not change dramatically when moving from standard to forced resolution. Cumulative covering fractions are typically within 10%–20% of one another for Si II, C IV, and O VI. However, the “cloud” sizes that produce these covering fractions (Figures 12–14) change by orders of magnitude. The forced refinement simulations yield more structure in the temperature, density, and metallicity




Figure 22. Number of Si II and C IV minima vs. impact parameter for LLS sightlines in the two simulations. The high- N_{\min} systems tend to be found at smaller impact parameters.

phase space than standard refinement, yet somehow the bulk, projected quantities such as mass and ion surface densities change only marginally, at a level often within observational uncertainties. This result suggests that such integrated quantities are poor diagnostics of the actual physical processes that give rise to the detected absorption.

FOGGIE’s findings on the kinematic structure of simulated absorbers point to the same conclusion: low-resolution simulations—that is, the currently cutting-edge cosmological simulations with Lagrangian-like fixed mass resolution—cannot adequately reproduce the kinematic complexity of real circumgalactic medium absorption. To quantify the changes in kinematic resolution, we have defined a new figure of merit (Equation (6)) that expresses how well simulations resolve the velocity field of detected absorption. Figure 17 shows that the standard-resolution simulations fail this criterion at rates 2–5 times higher than the high-resolution simulations. These improvements are driven by better sampling of the true velocity fields and by placing more resolution elements across physical structures, with the net effect being that there are more—and better-resolved—velocity components. Since bulk quantities such as projected column densities have proven to

have relatively little diagnostic power over small-scale gas structure, the detailed kinematic structure of the multiphase ions and their interrelationships may prove our best hope of actually constraining the small-scale physics of the circumgalactic medium. Quantifying the kinematic structure of absorption-line measurements in a meaningful and robust way is therefore an open challenge for both future simulations and future observational analyses.

This study is, however, limited in that we have simulated a single, low-mass halo and have a limited set of observational data from KODIAQ. For future FOGGIE comparisons to KODIAQ and $z \geq 2$ galaxy samples, we will simulate a larger sample of halos from a broad range of halo masses and analyze a larger sample of real LLS spectra. Moreover, the feedback scheme used here only includes the thermal feedback from supernovae, which may lead to stellar-to-halo mass ratios that are too high (Hopkins et al. 2014); likewise, outflows from purely thermal feedback are by design not multiphase, so any cool gas in the circumgalactic medium must arise in situ in this current generation of simulations (though, in principle, cold ISM gas could be entrained by the hot outflows, this is less likely; Schneider & Robertson 2017). For these reasons, future FOGGIE studies will improve the star formation and AGN feedback schemes in Enzo. Sources of nonthermal pressure from, e.g., cosmic rays (Butsky & Quinn 2018) have been shown in lower-resolution simulations to have a dramatic impact on the structure of the circumgalactic medium (Salem et al. 2016) such that the models are brought more in line with the observations; it will be important to quantify how these processes are affected by better resolving the circumgalactic medium.

Finally, the forced refinement scheme used here, while simple, is both computationally costly and limited. Future FOGGIE runs will take advantage of a new scheme we are calling “cooling refinement,” where, in addition to the forced refinement region used here, we allow the gas to refine based on its cooling length. Our initial tests show that this hybrid approach allows the denser gas to cool more efficiently while not introducing unnatural mixing in the low-density phases seen in more standard refinement schemes.

In Paper II, we show that resolving gas structures in the circumgalactic medium is vital for predicting signatures of the circumgalactic medium in emission. We conclude that future simulations aimed at understanding the coevolution of galaxies and their gaseous halos should attempt to resolve the circumgalactic gas in addition to the denser interstellar gas.

This study was primarily funded by the National Science Foundation via NSF AST-1517908, which helped support the contributions of L.C., B.W.O., N.L., J.O.M., and J.C.H. L.C. was additionally supported in part by *HST* AR #15012. B.W.O. was supported in part by NSF grants PHY-1430152, AST-1514700, and OAC-1835213, by NASA grants NNX12AC98G and NNX15AP39G, and by *HST* AR #14315. N.L. was also supported by NASA ADAP grant NNX16AF52G. N.E. was supported by *HST* AR #13919, *HST* GO #14268, and *HST* AR #14560. B.D.S. was supported by NSF AST-1615848. J.H.W. was supported by NSF grants AST-1614333 and OAC-1835213, NASA grant NNX17AG23G, and *HST*-AR-14326. C.B.H. was supported by *HST* AR #13917, *HST* AR #13919, and an NSF AAFP. Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the

NASA Advanced Supercomputing (NAS) Division at Ames Research Center and were sponsored by NASA’s Science Mission Directorate; we are grateful for the superb user support provided by NAS. Resources were also provided by the Blue Waters sustained-petascale computing project, which is supported by the NSF (award No. ACI 1238993 and ACI-1514580) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its NCSA. This work benefited from the dancing penguin and all the things emojis on Slack. The data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. This research has made use of the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExScI), under contract with the National Aeronautics and Space Administration. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain. We have made extensive use of the Python libraries Astropy, a community-developed core Python package for Astronomy (The Astropy Collaboration et al. 2018, <http://www.astropy.org>), datashader and holoviews from Anaconda (datashader.org), pandas (McKinney 2010), and seaborn. Computations described in this work were performed using the publicly available Enzo code, which is the product of a collaborative effort of many independent scientists from numerous institutions around the world.

Facilities: NASA Pleiades, Blue Waters, Keck (HIRES).

Software: astropy (The Astropy Collaboration et al. 2018), Cloudy (Ferland et al. 2013), Enzo (Bryan et al. 2014), grackle (Smith et al. 2017), yt (Turk et al. 2011), Trident (Hummels et al. 2017), ytree (Smith & Lang 2018).

Appendix A

Resolution Test: The circumgalactic medium at <100 pc

In addition to the $n_{\text{ref}} = 10$ “high-resolution” simulation we evolved to $z = 0$ and discuss in the main text, we also evolved a simulation enforcing $n_{\text{ref}} = 11$ within the same forced

refinement volume to $z = 2.5$. (More detailed properties of this simulation at $z = 3$ are provided in Paper II.) We show here the same analyses as in the main text, but for $z = 2.5$ only. At this redshift, the $n_{\text{ref}} = 11$ simulation has a fixed circumgalactic (and interstellar) resolution of 78.4 physical pc; the $n_{\text{ref}} = 10$ simulation has a circumgalactic medium resolution of 157 pc. We refer in this appendix to these simulations by these physical resolutions. Given the complex multiphysics nature of the simulations, a formal convergence test is difficult to define precisely, so instead we choose to examine more meaningful metrics, i.e., convergence as it applies to the physical phenomena that are relevant to this investigation. Broadly speaking, we find that improving the resolution by another factor of two spatially (and thus $8\times$ in volume) does result in yet smaller clouds, but the change is much less dramatic than between the standard-resolution simulation and the $n_{\text{ref}} = 10$ simulation.

Figure 23 shows slices of density, temperature, and metallicity, and Figure 24 shows projected column densities at $z = 2.5$, for both the 78 and 157 pc resolution simulations. These large-scale structures are much more similar to one another than to the standard-resolution simulation (Figure 5), though finer structure, smaller-scale turbulence, and more well-defined bubbles are visible in the 78 pc resolution projections. We give the total masses of stars, gas, and the ions of interest within the forced refinement volume at $z = 2.5$ in Table 3; these values are more converged for the high-resolution simulations than for the standard-resolution simulations. Following the analysis techniques outlined in Section 5, we show the 1D cloud sizes and masses in Figure 25; while the clouds are somewhat smaller in the 78 pc resolution simulation, they are much more similar in both extent and mass to those from the 157 pc resolution simulation than either are to the standard-resolution simulation. While the fraction of clouds larger than an $n_{\text{ref}} = 10$ cell (~ 157 pc) does not significantly change, many of the smallest clouds fragment to smaller scales, again as expected in a turbulent medium. The masses of the low-ionization clouds do not shift much, but essentially all of the O VI clouds in the 78 pc resolution simulation have masses $\lesssim 1000 M_{\odot}$. Figure 26 shows how this improved resolution translates to many more cells contributing to 80% of the column density along the line of sight, but that the combination of these changes does *not* translate to a significant change in the *kinematic* structure of the absorbing clouds.

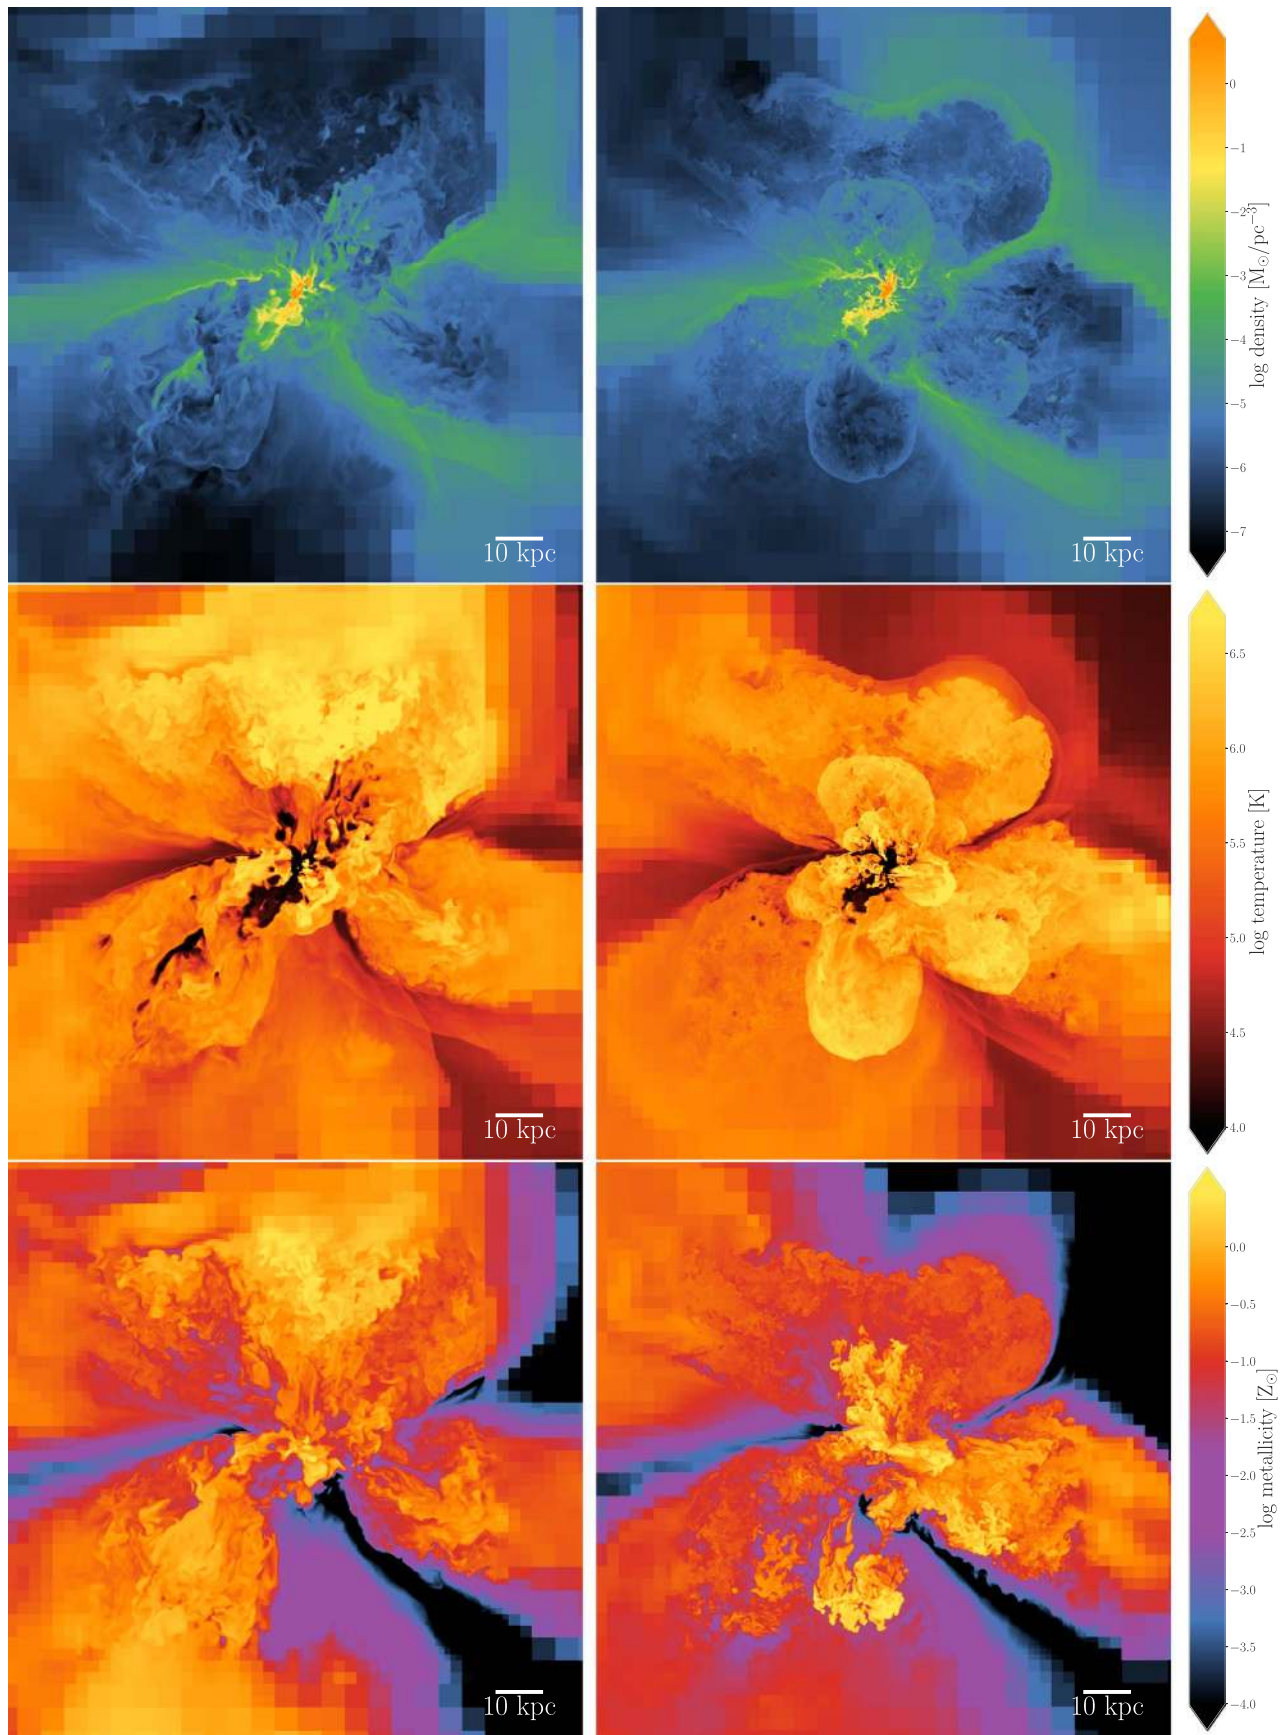


Figure 23. Slices of density (top), temperature (middle), and metallicity (bottom) through central halo in the 157 pc (left) and 78 pc (right) resolution simulations at $z = 2.5$.




Figure 24. H I, Si II, Si III, Si IV, C IV, and O VI projections at $z = 2.5$ in the 157 pc (left) and 78 pc resolution (right) simulations. Each panel is $200h^{-1}$ cKpc (~ 82 pkpc) across and deep. The central galaxy is in the top center of each panel. The color map for H I is chosen such that Lyman limit gas is black/blue and DLA gas is red/orange, with optically thin gas in gray scale. Both the bulk and detailed properties of these two simulations are much more similar than the 157 pc and standard-resolution simulations are (Figure 5).





Figure 25. Distributions of 1D cloud sizes (left) and masses (right) for H I (pink), Si II (purple), C IV (blue), and O VI (green) in the standard-resolution (top) and 157 pc resolution (middle), and 78 pc resolution (bottom) simulations at $z = 2.5$.





Figure 26. Distributions of number of cells in 1D clouds (left) and the resolved velocity sampling ratio (right) for H I (pink), Si II (purple), C IV (blue), and O VI (green) in the standard-resolution (top), 157 pc resolution (middle), and 78 pc resolution (bottom) simulations at $z = 2.5$. Note that for the left panels the x-axis has been extended by a factor of two compared to Figure 13.

Table 4
(Continued)

Target	z_{abs}	$\log N_{\text{HI}}$	$\sigma_{\log N_{\text{HI}}}$	[X/H]	$\sigma_{\text{metal}}^{\text{up}}$	$\sigma_{\text{metal}}^{\text{low}}$	Ion	i	v_i	σ_{v_i}	b_i	σ_{b_i}	$\log N_i$	$\sigma_{\log N_i}$	Reference
Q2126–158	2.90731	16.45	0.25	–999	999	999	C IV	1	–87.8	1.5	5.3	2.3	12.02	0.14	B14
Q2126–158	2.90731	16.45	0.25	–999	999	999	C IV	2	–30.9	2.4	14.4	2.4	12.90	0.16	B14
Q2126–158	2.90731	16.45	0.25	–999	999	999	C IV	3	–0.0	2.1	23.6	2.5	13.48	0.04	B14
Q2126–158	2.90731	16.45	0.25	–999	999	999	C IV	4	4.2	0.9	4.2	1.8	12.33	0.19	B14
Q2126–158	2.90731	16.45	0.25	–999	999	999	O VI	1	–44.1	1.5	15.4	1.9	13.50	0.07	B14
Q2126–158	2.90731	16.45	0.25	–999	999	999	O VI	2	0.8	1.8	29.0	2.7	13.96	0.03	B14

Note. The full data set is available as an electronic table. The references are Lehner et al. (2014) as L14, Lehner et al. (2016) as L16, and Burns (2014) as B14. (This table is available in its entirety in machine-readable form.)

Appendix C

Additional High-resolution Phase-space Diagrams

We show here additional examples of the velocity phase-space diagrams introduced in Section 4 and Figures 10 and 11. Figures 27–33 are sightlines from the high-resolution simulation, highlighting examples of where the additional resolution either

increases or decreases the number of apparent spectral components and where the low and high ions are aligned or misaligned in interesting ways. Figures 34–39 are from the standard-resolution simulation, highlighting the effects of coarse or variable resolution and cases where the bulk of the absorption for some of the ions is from individual cells.

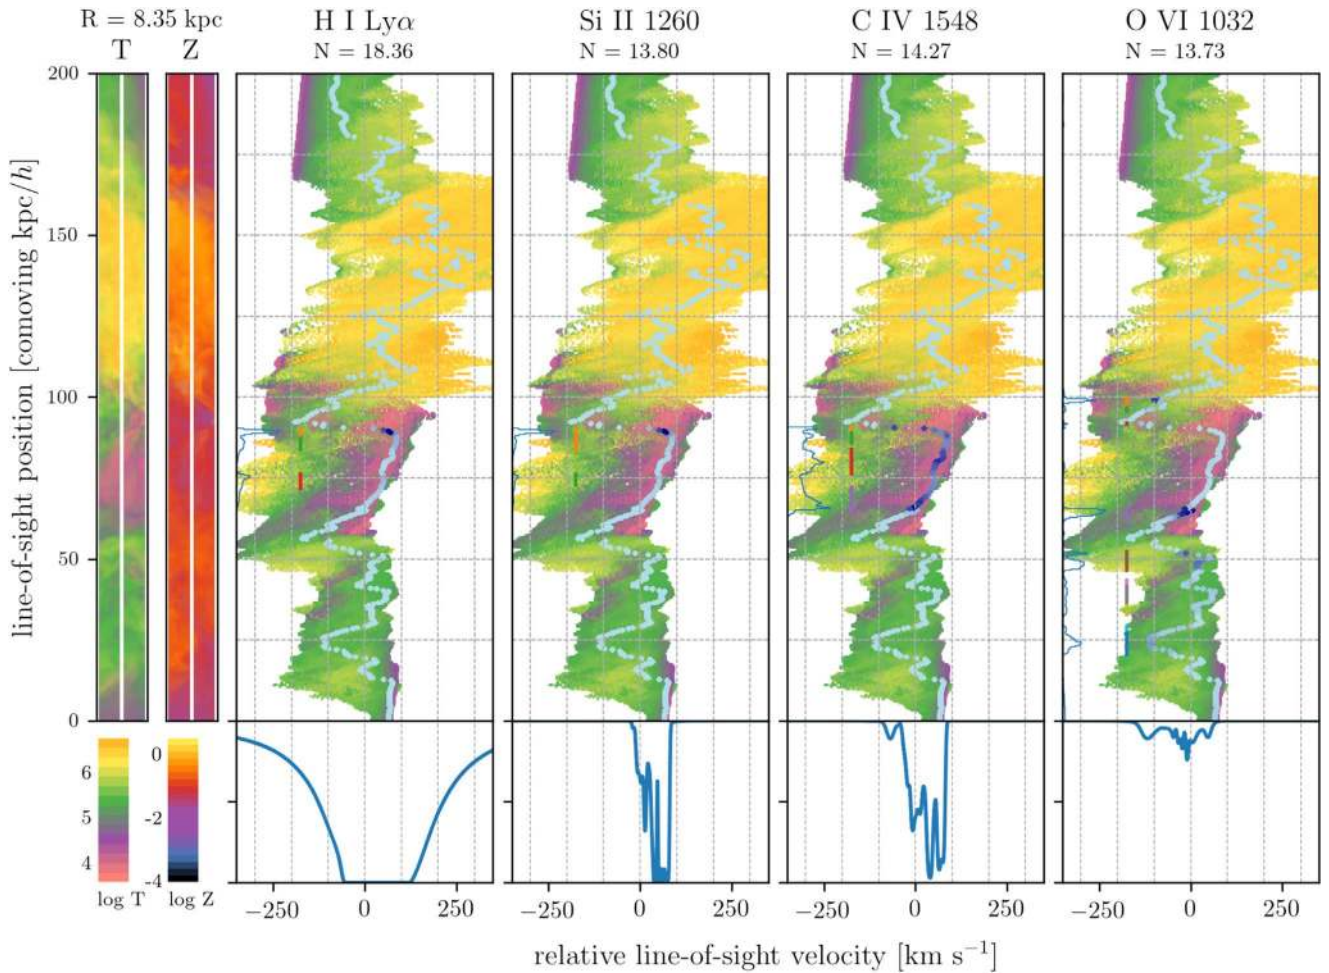


Figure 27. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.0$. Here the O VI is at physically distinct locations from the lower ions, despite being at similar velocities.

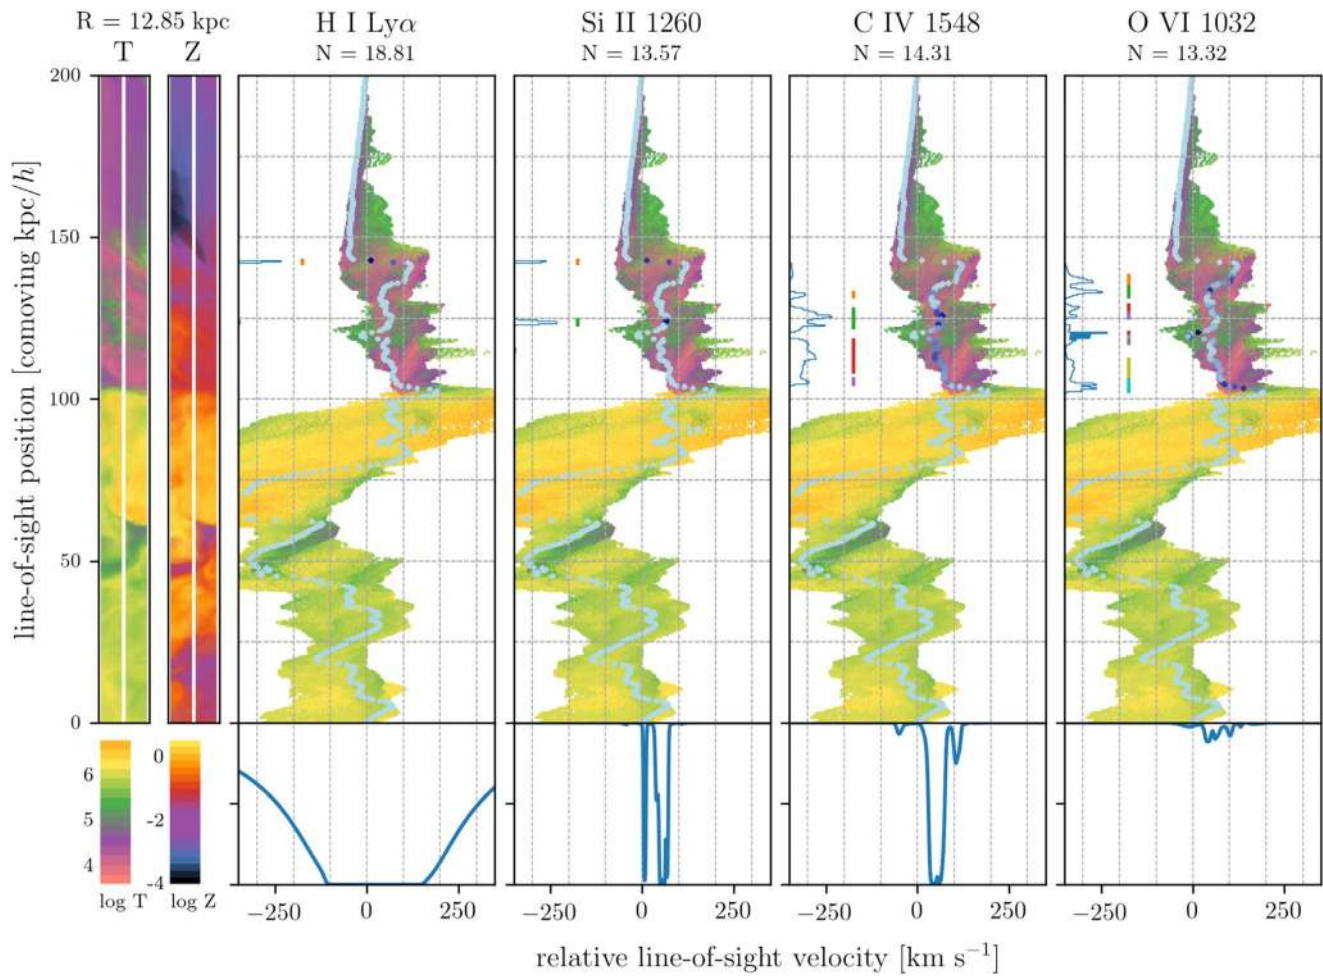


Figure 28. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.0$. Here a velocity shear along the line of sight at $\sim 140 h^{-1}$ ckpc leads to an H I- and Si II-bearing cloud being resolved with only two cells with a large $\Delta v_{\text{cloud}}/\sigma_{\text{th}}$. The Si II absorption only appears blended because of another cloud at the same velocity but $\sim 20 h^{-1}$ ckpc away. The C IV- and O VI-bearing clouds are both more extended and better resolved. This sightline is also a good example of the very hot, high-velocity, high-metallicity gas along the line of sight having essentially no associated absorption in these ionic tracers.

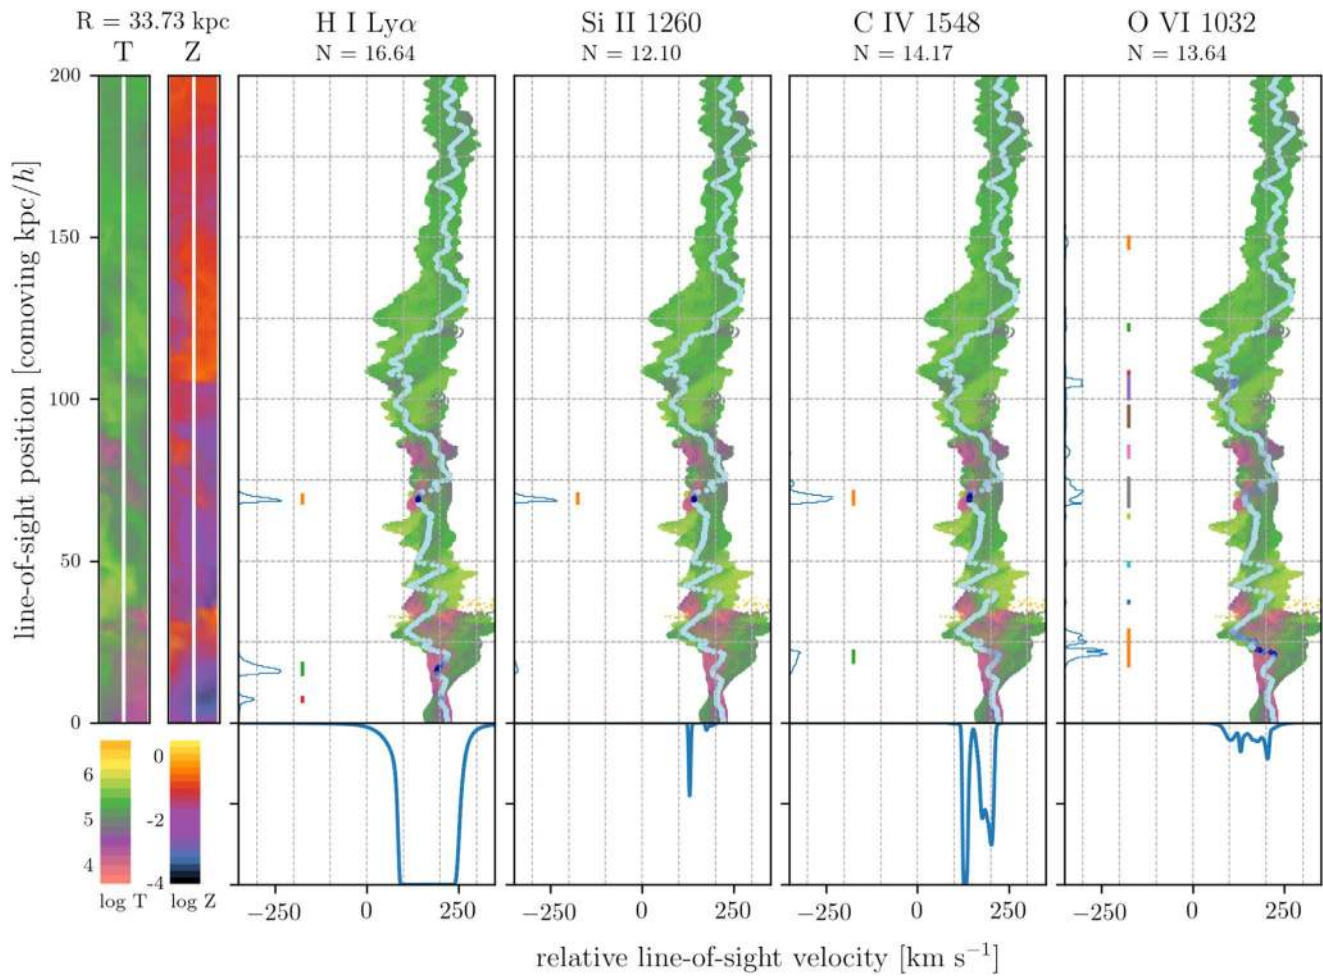


Figure 29. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.0$. The H I absorption primarily arises from two distinct clouds $\sim 50h^{-1}$ ckpc apart, with the Si II absorption being primarily associated with only one of the clouds; an Si II/H I ratio for this sightline would therefore misestimate the metallicity of the individual clouds. There are many more O VI clouds along the sightline than Si II clouds.

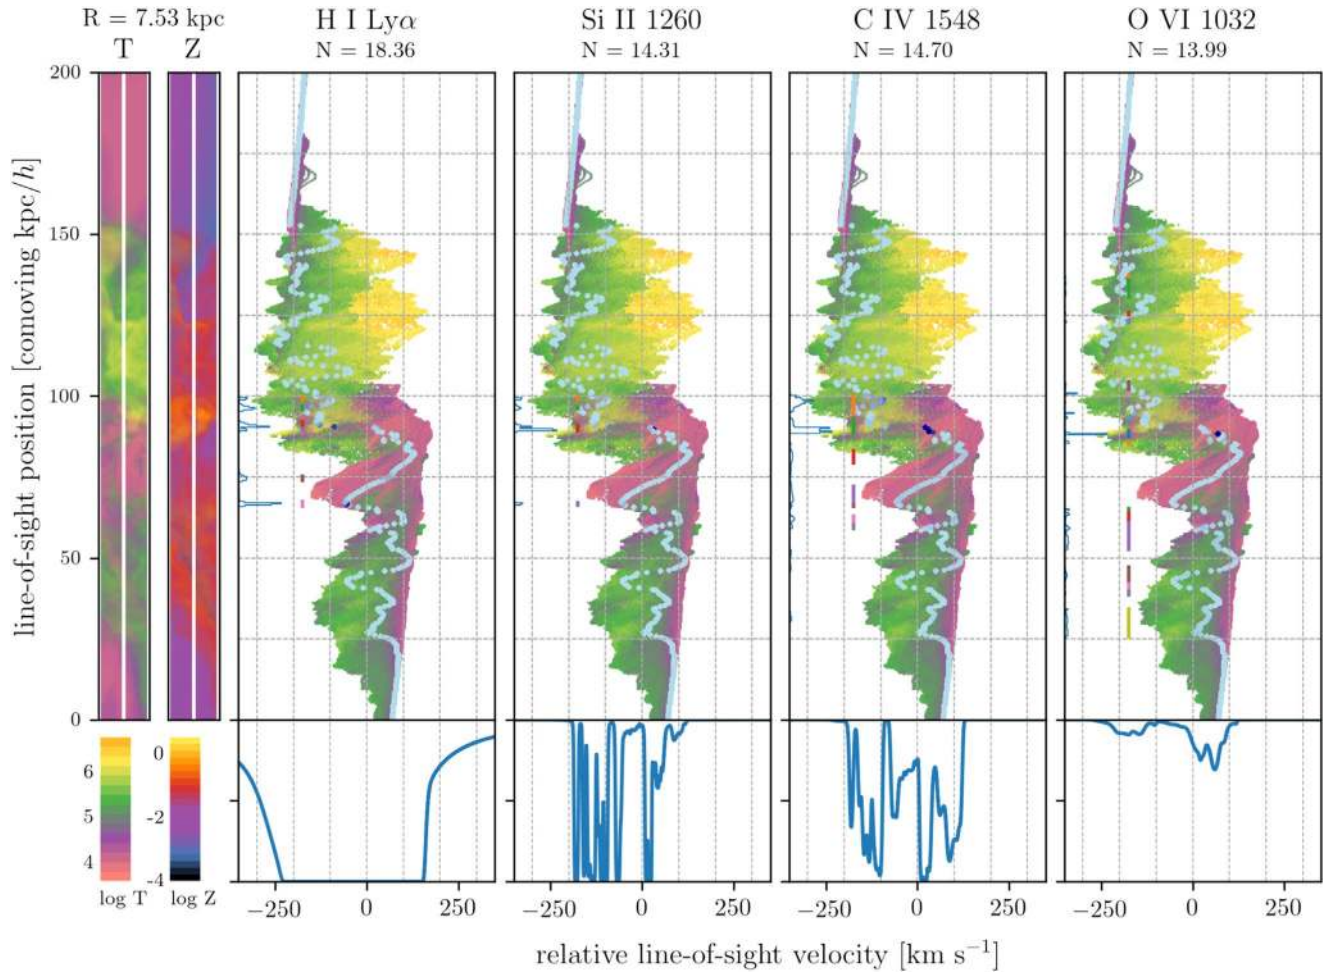


Figure 30. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.0$. This example shows a wide range of velocity and temperature structures along the line of sight.

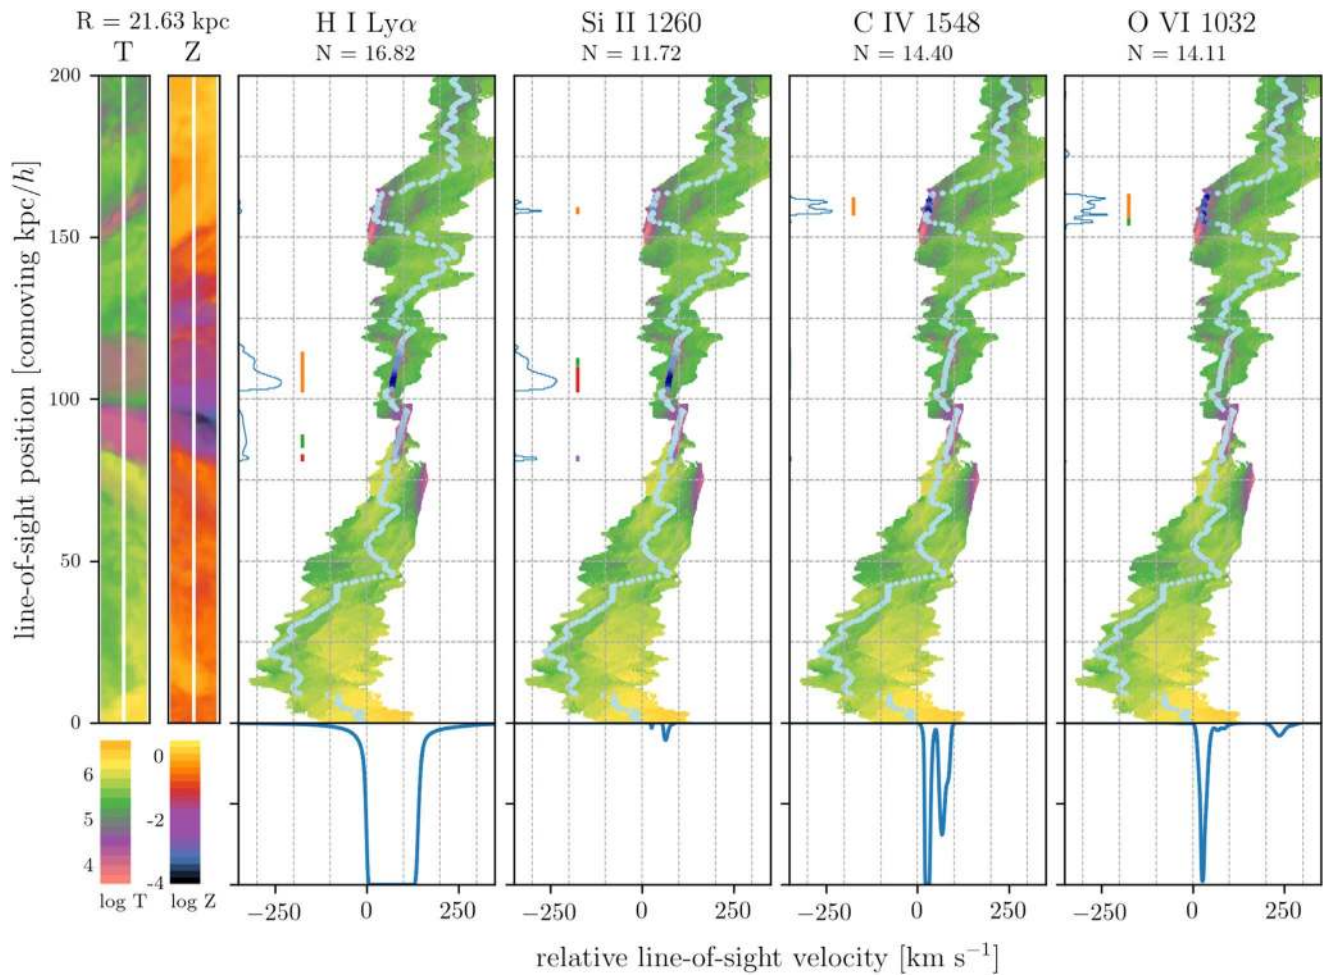


Figure 31. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.5$. This is a good example of O VI and C IV having absorption at velocities coincident with H I, but arising from gas $\sim 50h^{-1}$ ckpc away.

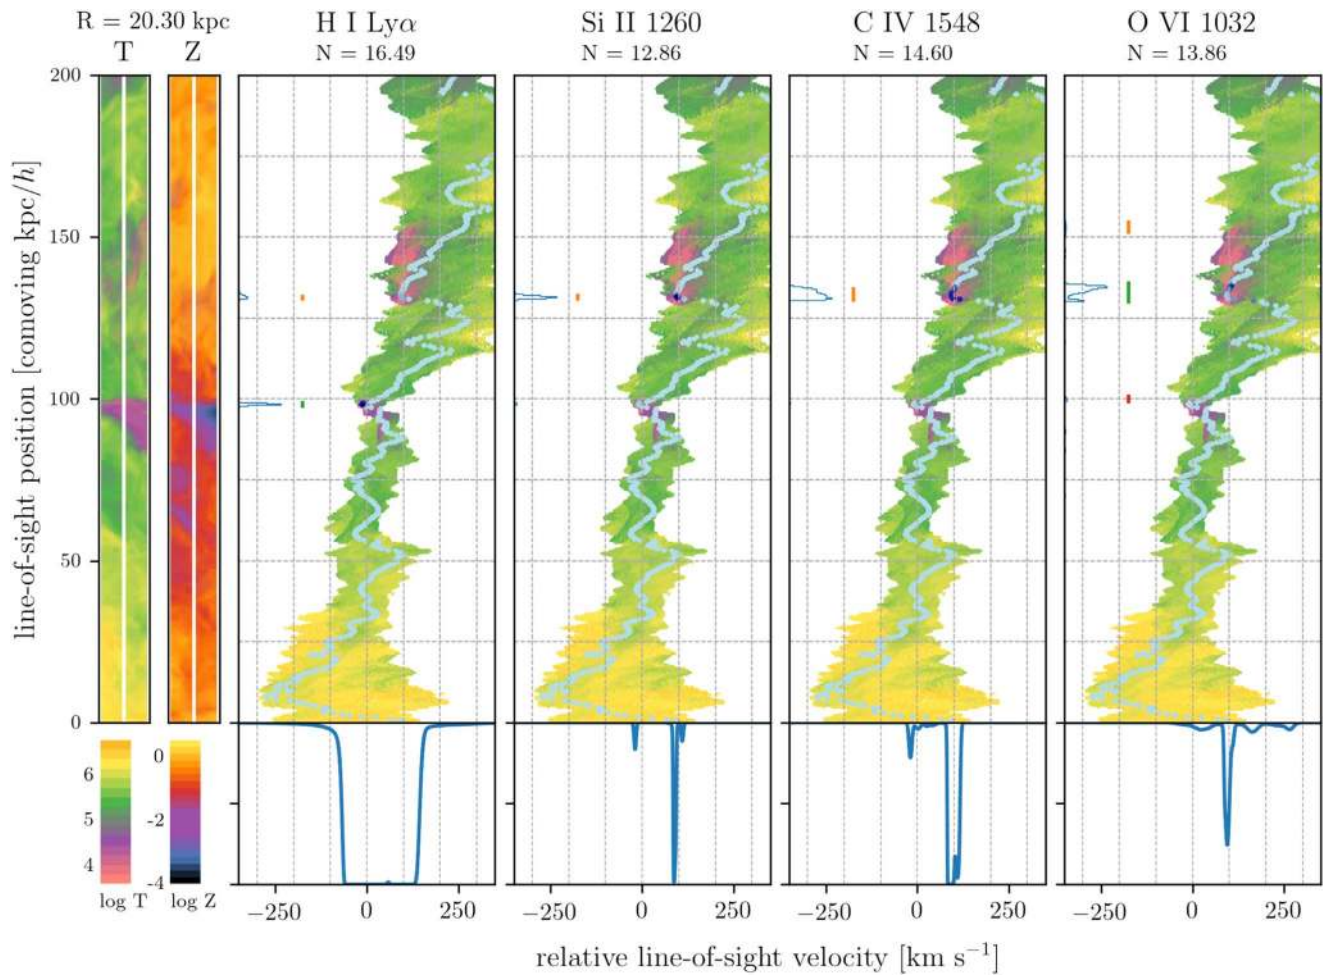


Figure 32. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.5$. This is a good example of the metal ions arising from spatially colocated gas, while most of the H I-traced gas is physically distinct; because the observed H I is saturated, this configuration may lead to a misestimation of the gas metallicity.

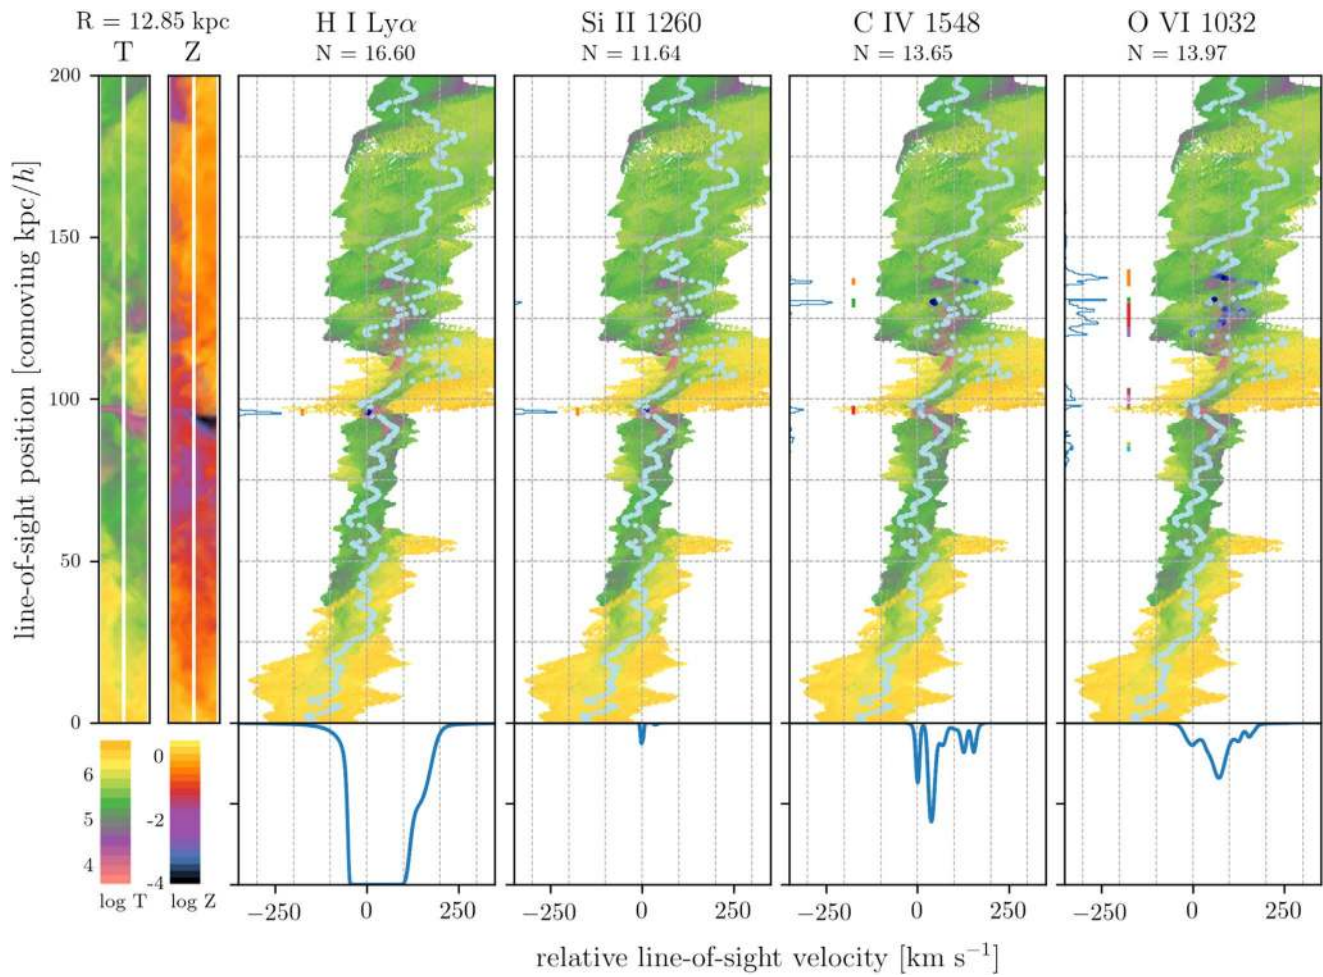


Figure 33. Same as Figure 10, but for a sightline through the high-resolution simulation at $z = 2.5$. This example shows an O VI absorber with a broad velocity width that is predominantly caused by the bulk gas kinematics rather than thermal broadening.




Figure 34. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.0$. This is a good example of essentially all of the absorption arising from $\lesssim 3$ cells.

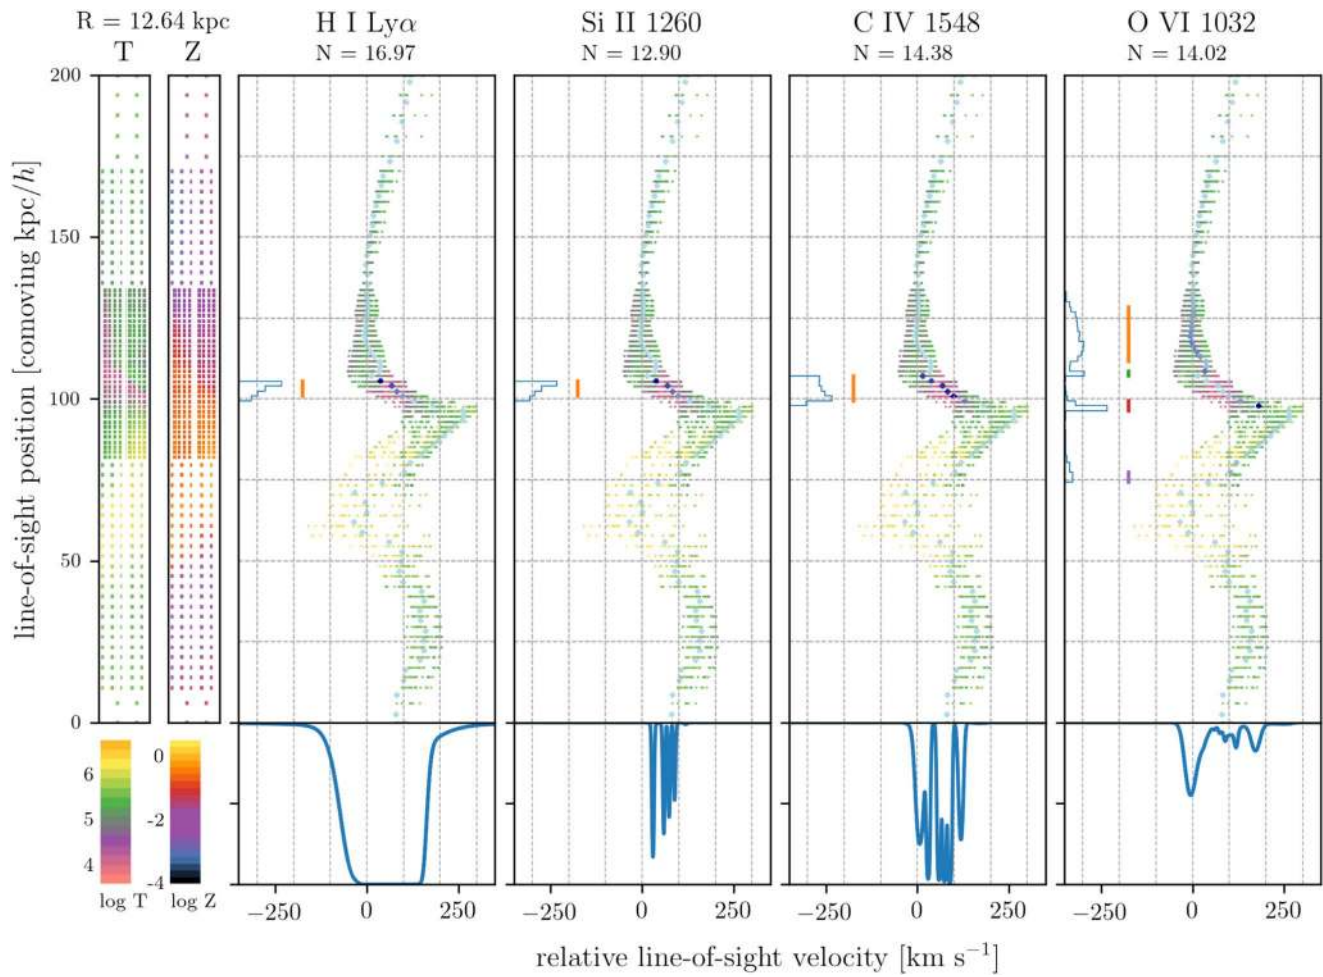


Figure 35. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.0$. Absorption “lines” from individual cells are particularly evident in Si II and C IV.

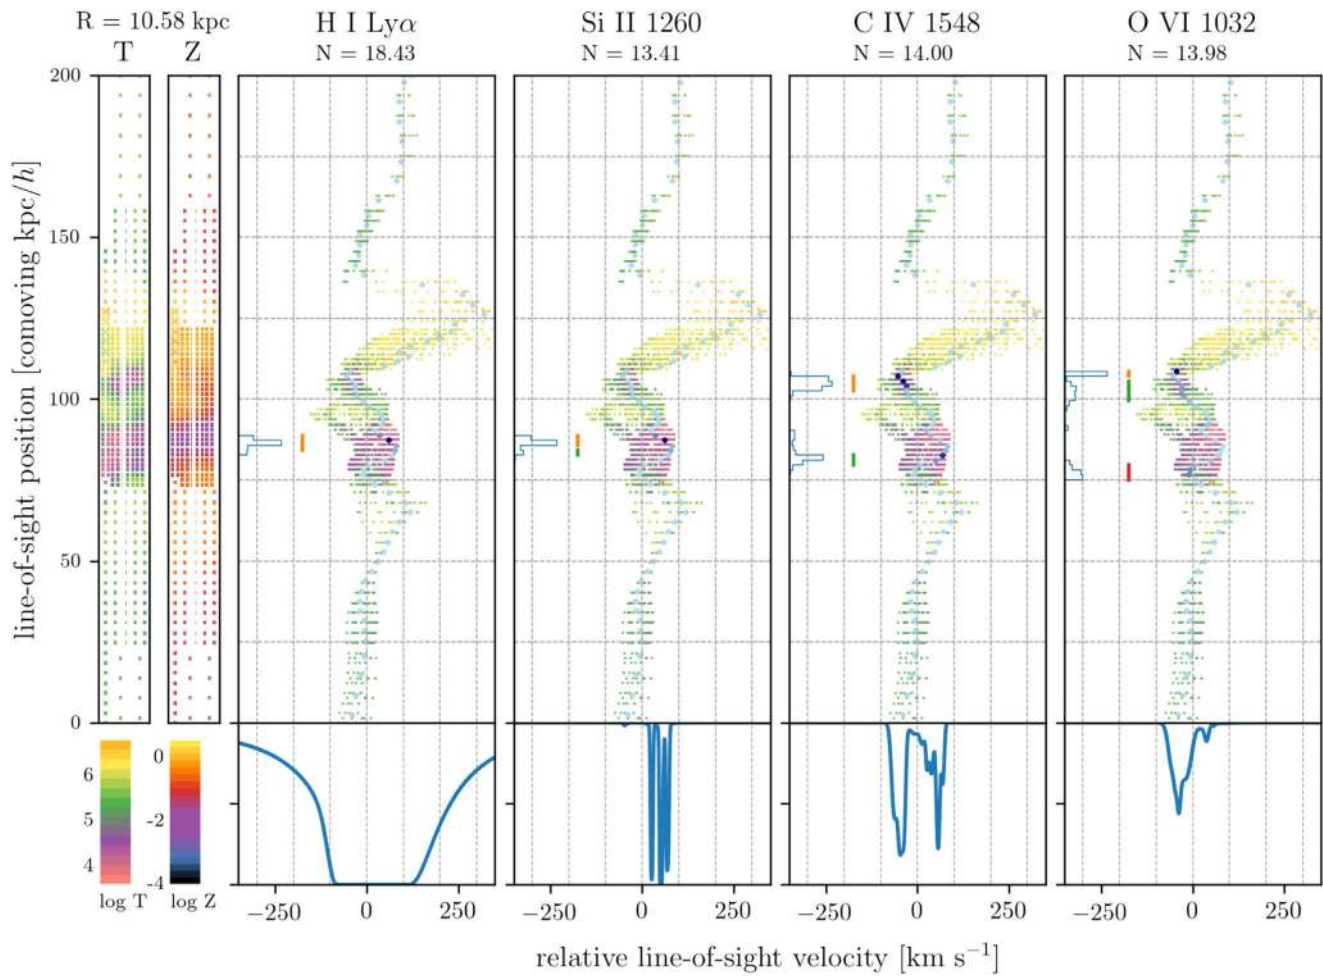


Figure 36. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.0$. This is a good example of absorption preferentially arising in more refined regions in the standard-resolution simulation; even so, this is another good example of multiple Si II minima each arising from individual cells.

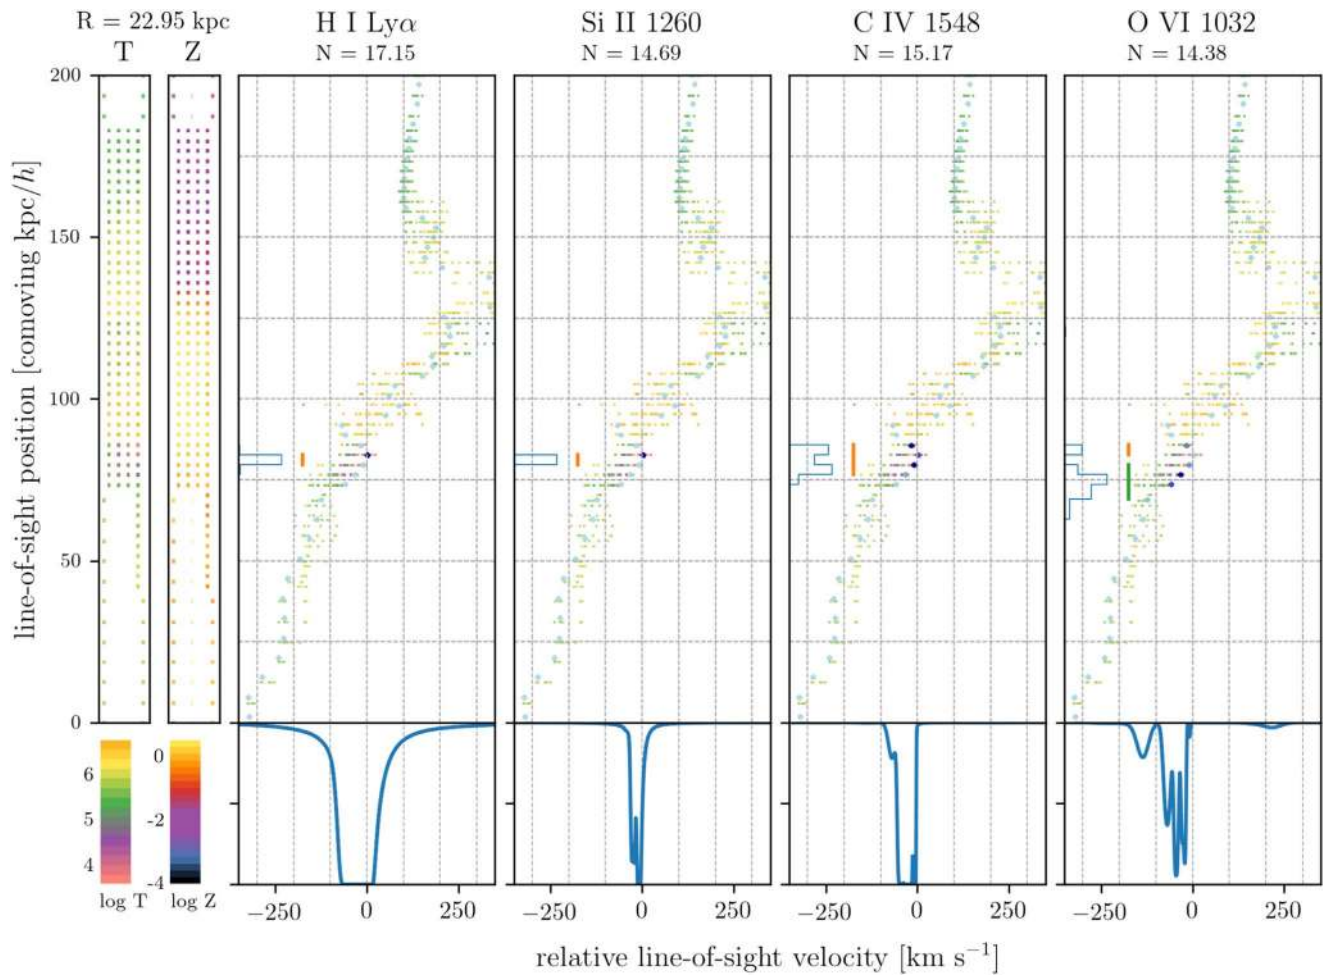


Figure 37. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.5$. Absorption from individual cells is clearly visible for each ion.

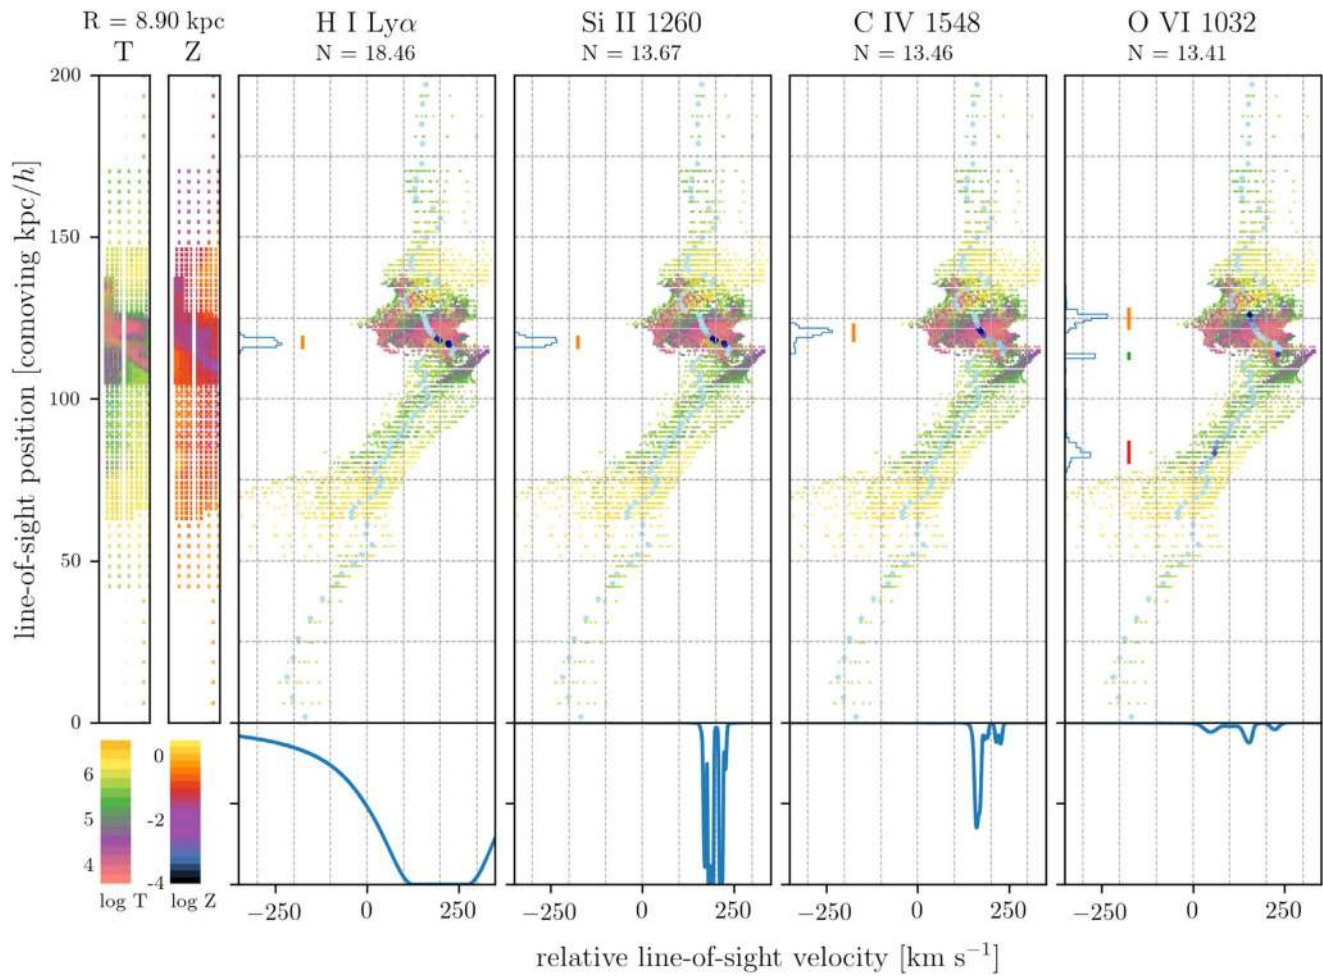


Figure 38. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.5$. Many of the standard-resolution simulations show the higher column density gas clearly favoring the more refined regions—which are, for these runs, by construction higher-density gas.

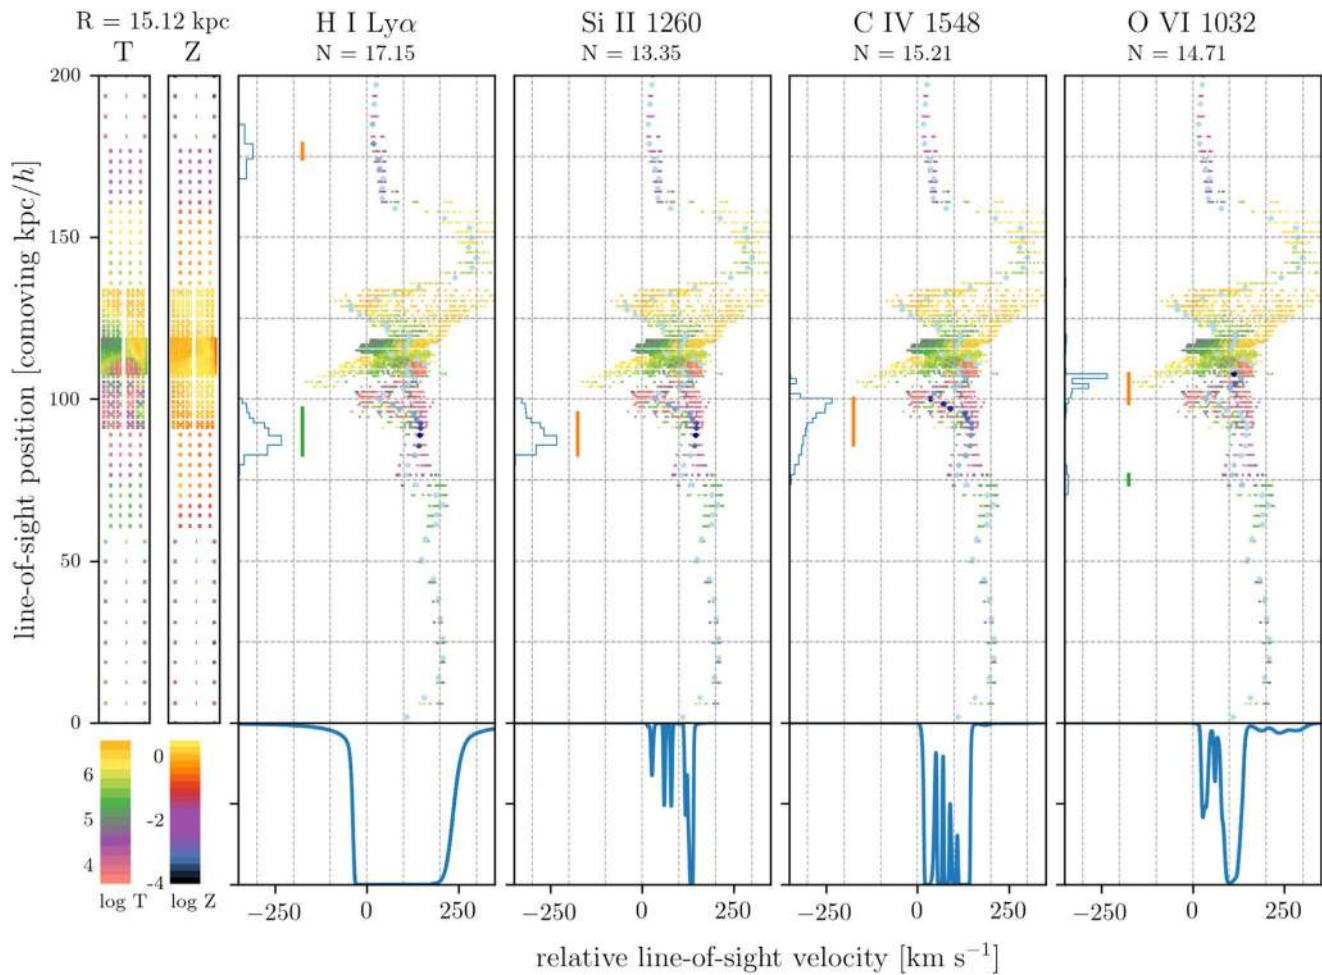


Figure 39. Same as Figure 10, but for a sightline through the standard-resolution simulation at $z = 2.5$. This example shows relatively large ($\sim 15h^{-1}$ ckpc) clouds that still have unresolved velocity structure; this is most evident in the C IV absorption.

ORCID iDs

Molly S. Peebles <https://orcid.org/0000-0003-1455-8788>
 Lauren Corlies <https://orcid.org/0000-0002-0646-1540>
 Jason Tumlinson <https://orcid.org/0000-0002-7982-412X>
 Brian W. O’Shea <https://orcid.org/0000-0002-2786-0348>
 Nicolas Lehner <https://orcid.org/0000-0001-9158-0829>
 John M. O’Meara <https://orcid.org/0000-0002-7893-1054>
 J. Christopher Howk <https://orcid.org/0000-0002-2591-3792>
 Nicholas Earl <https://orcid.org/0000-0003-1714-7415>
 Britton D. Smith <https://orcid.org/0000-0002-6804-630X>
 John H. Wise <https://orcid.org/0000-0003-1173-8847>
 Cameron B. Hummels <https://orcid.org/0000-0002-3817-8133>

References

Armillaotta, L., Fraternali, F., Werk, J. K., Prochaska, J. X., & Marinacci, F. 2017, *MNRAS*, 470, 114
 Borthakur, S., Heckman, T., Tumlinson, J., et al. 2016, *ApJ*, 833, 259
 Bryan, G. L., Norman, M. L., O’Shea, B. W., et al. 2014, *ApJS*, 211, 19
 Burns, V. 2014, Undergraduate Honors thesis, Univ. Notre Dame
 Butsky, I., & Quinn, T. R. 2018, *ApJ*, 868, 108
 Cen, R., & Ostriker, J. P. 2006, *ApJ*, 650, 560
 Christensen, C. R., Dave, R., Brooks, A., Quinn, T., & Shen, S. 2018, *ApJ*, 867, 142
 Christensen, C. R., Davé, R., Governato, F., et al. 2016, *ApJ*, 824, 57

Churchill, C. W., Vander Vliet, J. R., Trujillo-Gomez, S., Kacprzak, G. G., & Klypin, A. 2015, *ApJ*, 802, 10
 Corlies, L., Peebles, M. S., Tumlinson, J., et al. 2018, arXiv:1811.05060
 Ellison, S. L., Ibata, R., Pettini, M., et al. 2004, *A&A*, 414, 79
 Ferland, G. J., Porter, R. L., van Hoof, P. A. M., et al. 2013, *RMxAA*, 49, 137
 Field, G. B. 1965, *ApJ*, 142, 531
 Fielding, D., Quataert, E., & Martizzi, D. 2018, *MNRAS*, 481, 3325
 Fitzpatrick, E. L., & Spitzer, L., Jr. 1997, *ApJ*, 475, 623
 Ford, A. B., Oppenheimer, B. D., Davé, R., et al. 2013, *MNRAS*, 432, 89
 Ford, A. B., Werk, J. K., Davé, R., et al. 2016, *MNRAS*, 459, 1745
 Fumagalli, M., O’Meara, J. M., & Prochaska, J. X. 2011, *Sci*, 334, 1245
 Genel, S., Bryan, G. L., Springel, V., et al. 2019, *ApJ*, 871, 21
 Haardt, F., & Madau, P. 2012, *ApJ*, 746, 125
 Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, *MNRAS*, 445, 581
 Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018, *MNRAS*, 480, 800
 Hummels, C. B., Bryan, G. L., Smith, B. D., & Turk, M. J. 2013, *MNRAS*, 430, 1548
 Hummels, C. B., Smith, B. D., & Silvia, D. W. 2017, *ApJ*, 847, 59
 Lehner, N., & Howk, J. C. 2011, *Sci*, 334, 955
 Lehner, N., Howk, J. C., Thom, C., et al. 2012, *MNRAS*, 424, 2896
 Lehner, N., O’Meara, J. M., Fox, A. J., et al. 2014, *ApJ*, 788, 119
 Lehner, N., O’Meara, J. M., Howk, J. C., Prochaska, J. X., & Fumagalli, M. 2016, *ApJ*, 833, 283
 Liang, C. J., & Remming, I. S. 2018, arXiv:1806.10688
 Lopez, S., Tejos, N., Ledoux, C., et al. 2018, *Natur*, 554, 493
 McCourt, M., Oh, S. P., O’Leary, R., & Madigan, A.-M. 2018, *MNRAS*, 473, 5407
 McKinney, W. 2010, in Proc. 9th Python in Science Conf., SciPy2010, ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 51, <https://conference.scipy.org/proceedings/scipy2010/pdfs/proceedings.pdf>

- Morton, D. C. 2003, *ApJS*, 149, 205
- O’Meara, J. M., Lehner, N., Howk, J. C., et al. 2017, *AJ*, 154, 114
- Oppenheimer, B. D., Crain, R. A., Schaye, J., et al. 2016, *MNRAS*, 460, 2157
- Oppenheimer, B. D., & Davé, R. 2008, *MNRAS*, 387, 577
- Oppenheimer, B. D., Segers, M., Schaye, J., Richings, A. J., & Crain, R. A. 2018, *MNRAS*, 474, 4740
- Padoan, P., & Nordlund, Å. 2002, *ApJ*, 576, 870
- Pagel, B. E. J. 2008, in ASP Conf. Ser. 390, Pathways through an Eclectic Universe, ed. J. H. Knapen, T. J. Mahoney, & A. Vazdekis (San Francisco, CA: ASP), 483
- Peebles, M. S., Weinberg, D. H., Davé, R., Fardal, M. A., & Katz, N. 2010a, *MNRAS*, 404, 1281
- Peebles, M. S., Weinberg, D. H., Davé, R., Fardal, M. A., & Katz, N. 2010b, *MNRAS*, 404, 1295
- Peebles, M. S., Werk, J. K., Tumlinson, J., et al. 2014, *ApJ*, 786, 54
- Péroux, C., Rahmani, H., Arrigoni Battaia, F., & Augustin, R. 2018, *MNRAS*, 479, L50
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, 571, A16
- Prochaska, J. X., Werk, J. K., Worseck, G., et al. 2017, *ApJ*, 837, 169
- Putman, M. E., Peek, J. E. G., & Joung, M. R. 2012, *ARA&A*, 50, 491
- Rauch, M., Sargent, W. L. W., & Barlow, T. A. 2001a, *ApJ*, 554, 823
- Rauch, M., Sargent, W. L. W., Barlow, T. A., & Carswell, R. F. 2001b, *ApJ*, 562, 76
- Rubin, K. H. R., Diamond-Stanic, A. M., Coil, A. L., & Crighton, N. H. M. 2018a, *ApJ*, 868, 142
- Rubin, K. H. R., O’Meara, J. M., Cooksey, K. L., et al. 2018b, *ApJ*, 859, 146
- Salem, M., Bryan, G. L., & Corlies, L. 2016, *MNRAS*, 456, 582
- Schaye, J. 2001, *ApJ*, 559, 507
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, 446, 521
- Schneider, E. E., & Robertson, B. E. 2017, *ApJ*, 834, 144
- Sharma, P., Parrish, I. J., & Quataert, E. 2010, *ApJ*, 720, 652
- Smith, B., & Lang, M. 2018, Ytree: Merger-tree Toolkit, Zenodo, doi:10.5281/zenodo.1174374
- Smith, B. D., Bryan, G. L., Glover, S. C. O., et al. 2017, *MNRAS*, 466, 2217
- Stocke, J. T., Keeney, B. A., Danforth, C. W., et al. 2013, *ApJ*, 763, 148
- Suresh, J., Rubin, K. H. R., Kannan, R., et al. 2017, *MNRAS*, 465, 2966
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
- Thom, C., Peek, J. E. G., Putman, M. E., et al. 2008, *ApJ*, 684, 364
- Thom, C., Tumlinson, J., Werk, J. K., et al. 2012, *ApJL*, 758, L41
- Trujillo-Gomez, S., Klypin, A., Colín, P., et al. 2015, *MNRAS*, 446, 1140
- Tumlinson, J., Peebles, M. S., & Werk, J. K. 2017, *ARA&A*, 55, 389
- Tumlinson, J., Thom, C., Werk, J. K., et al. 2013, *ApJ*, 777, 59
- Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, 192, 9
- Turner, M. L., Schaye, J., Crain, R. A., et al. 2017, *MNRAS*, 471, 690
- van de Voort, F., Springel, V., Mandelker, N., van den Bosch, F. C., & Pakmor, R. 2019, *MNRAS*, 482, L85
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *MNRAS*, 444, 1518
- Voit, G. M., Bryan, G. L., O’Shea, B. W., & Donahue, M. 2015, *ApJL*, 808, L30
- Wakker, B. P., York, D. G., Howk, J. C., et al. 2007, *ApJL*, 670, L113
- Wakker, B. P., York, D. G., Wilhelm, R., et al. 2008, *ApJ*, 672, 298
- Werk, J. K., Prochaska, J. X., Cantalupo, S., et al. 2016, *ApJ*, 833, 54
- Werk, J. K., Prochaska, J. X., Tumlinson, J., et al. 2014, *ApJ*, 792, 8