



Revealing the Dark Threads of the Cosmic Web

Joseph N. Burchett¹, Oskar Elek², Nicolas Tejos³, J. Xavier Prochaska^{1,4}, Todd M. Tripp⁵, Rongmon Bordoloi⁶, and Angus G. Forbes²

¹ Department of Astronomy & Astrophysics, University of California, 1156 High Street, Santa Cruz, CA 95064, USA; burchett@ucolick.org

² Department of Computational Media, University of California, 1156 High Street, Santa Cruz, CA 95064, USA

³ Pontificia Universidad Católica de Valparaíso, Chile

⁴ Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU), 5-1-5 Kashiwanoha, Kashiwa, 277-8583, Japan

⁵ University of Massachusetts—Amherst, 710 North Pleasant Street, Amherst, MA 01003-9305, USA

⁶ North Carolina State University, 2401 Stinson Drive, Raleigh, NC 27695, USA

Received 2019 November 26; revised 2020 January 10; accepted 2020 January 23; published 2020 March 10

Abstract

Modern cosmology predicts that matter in our universe today has assembled into a vast network of filamentary structures colloquially termed the “cosmic web.” Because this matter is either electromagnetically invisible (i.e., dark) or too diffuse to image in emission, tests of this cosmic web paradigm are limited. Wide-field surveys do reveal web-like structures in the galaxy distribution, but these luminous galaxies represent less than 10% of baryonic matter. Statistics of absorption by the intergalactic medium (IGM) via spectroscopy of distant quasars support the model yet have not conclusively tied the diffuse IGM to the web. Here, we report on a new method inspired by the *Physarum polycephalum* slime mold that is able to infer the density field of the cosmic web from galaxy surveys. Applying our technique to galaxy and absorption-line surveys of the local universe, we demonstrate that the bulk of the IGM indeed resides in the cosmic web. From the outskirts of cosmic web filaments, at approximately the cosmic mean matter density (ρ_m) and ~ 5 virial radii from nearby galaxies, we detect an increasing HI absorption signature toward higher densities and the circumgalactic medium, to $\sim 200\rho_m$. However, the absorption is suppressed within the densest environments, suggesting shock-heating and ionization deep within filaments and/or feedback processes within galaxies.

Unified Astronomy Thesaurus concepts: [Cosmic web \(330\)](#); [Intergalactic medium \(813\)](#); [Intergalactic filaments \(811\)](#); [Voids \(1779\)](#); [Quasar absorption line spectroscopy \(1317\)](#); [Circumgalactic medium \(1879\)](#); [Galaxy evolution \(594\)](#); [Galaxy quenching \(2040\)](#)

1. Introduction

After decades of study, the Λ cold dark matter (Λ CDM) cosmological framework is now significantly refined, and many basic parameters of this theory are constrained by observations. At this stage, investigations of many key cosmology tests and galaxy evolution questions require computationally intensive large-scale simulations that attempt to incorporate a wide variety of physical processes with adequate spatial and mass resolution to capture essential behaviors of galaxies and the universe. These cosmological simulations are highly complex, generally relying on varying assumptions or “prescriptions” for dealing with these unresolved physical processes hydrodynamically (e.g., Oppenheimer & Davé 2006; Schaye et al. 2015). In principle, the intergalactic medium (IGM), which should primarily involve hydrogen ionization and recombination, provides one of the most straightforward tests of the simulations (Davé et al. 1999). However, the IGM has extremely low densities and is therefore challenging to detect. Moreover, galaxies are surrounded by a gaseous circumgalactic medium (CGM; e.g., Bergeron & Boissé 1991; Lanzetta et al. 1995; Keeney et al. 2013; Tumlinson et al. 2017, for a review) that is affected by a more complicated mixture of galactic inflows, outflows, and energy “feedback,” and it is not immediately clear how to distinguish hydrogen in the CGM from the intergalactic hydrogen so that the simulations can be tested in the relatively simple regime of the IGM.

As the most sensitive method for detecting the highly diffuse gas in the universe, absorption-line spectroscopy has

been the primary technique for probing the CGM and IGM. Generally, the technique involves using a bright background source such as a quasi-stellar object (QSO) or a galaxy to probe the halo of a foreground object projected near the line of sight. On scales beyond the CGM, the emergent cosmic web (Bond et al. 1996)—generically predicted in cosmological models of a CDM universe—should be permeated by low-density IGM gas (Cen et al. 1994). Efforts to correlate absorption signatures with galaxies on IGM scales have shown that gas and galaxies are indeed correlated over Megaparsec-scale distances (Morris et al. 1993; Chen & Mulchaey 2009; Tejos et al. 2012, 2014; Prochaska et al. 2019), and HI absorption tomography using multiple sightlines probing a single filament suggests a stronger absorption signal near the center of that filament (Wakker et al. 2015). Under the assumption that cosmic filaments connect massive galaxy clusters, it has also been shown that warm-hot (10^{5-6} K) gas is enhanced within filaments (Tejos et al. 2016; Pessa et al. 2018), indicating that the IGM is heated in these over-dense environments. However, it remains to be demonstrated over a cosmologically representative volume that the IGM absorbers routinely detected in QSO spectroscopy indeed trace the *same* large-scale cosmic web structure evident in galaxy redshift surveys (de Lapparent et al. 1986).

The principal challenge is identifying and quantifying the large-scale structure of the cosmic web, particularly in regions where the most diffuse gas is purported to exist. Several groups have developed automatic methods to identify large-scale




Figure 1. Monte Carlo Physarum Machine (MCPM) cosmic web reconstruction using 37,662 galaxies from Sloan Digital Sky Survey (SDSS). Top: large-scale visualization of the emergent structure identified by our MCPM algorithm. This intricate filamentary network is reconstructed given only the SDSS galaxy coordinates, redshifts, and masses. Bottom: three individual regions showing the underlying SDSS galaxies (G) and the superimposed MCPM filament density field (G+F).

structures and/or derive the underlying density fields (e.g., Aragón-Calvo et al. 2007; Sousbie et al. 2011; Cautun et al. 2013; Tempel et al. 2014; Chen et al. 2016; Libeskind et al. 2018). The difficulty of the problem for all of these methods is compounded by the lack of a “ground truth” solution for large-scale structures in the observed universe. These methods have been fruitfully applied to galaxy survey data to show that morphology, star formation rate, and cold gas content depend not only on their halo masses but also on a galaxy’s location within the large-scale structure (e.g., Stark et al. 2016; Chen et al. 2017; Kuutma et al. 2017). However, associating the IGM absorption with this underlying structure requires key features that are often missing from such methods. First, the cosmic web density values must be defined at points in the survey volume on scales of several Megaparsecs away from galaxies, where the QSO sightlines probe low-density regions of the IGM. Second, the environment must be quantified across a wide dynamic range of scale, from small scales of galaxy groups to the larger scales of cosmic web filaments. The former generally excludes methods such as nearest-neighbor distances, and the latter generally excludes those based on purely geometrical filament identification.

Here, we present and apply an alternative novel approach inspired by the *Physarum polycephalum* slime mold, which attempts to satisfy both requirements simultaneously, as shown in Figure 1. Our model implicitly traces the cosmic web structure by finding efficient continuous pathways between galaxies, which are known to be preferentially distributed along filaments.

2. Methods

2.1. Reconstructing the Cosmic Web

A key difficulty in identifying cosmic web structures such as filaments is the heavily under-constrained nature of the problem: how does one reconstruct a physically plausible 3D structure out of sparse and often heterogeneous observations? We approach this challenge from the perspective of finding a meaningful interpolator for such sparse data—put in other words, finding a generative model conditioned by the data and informed by an appropriate statistical prior.

Our generative model is heavily inspired by the work of Jones (2010), who proposed an agent-based algorithm to mimic the growth and development of the *Physarum polycephalum* “slime mold.” This model, a kind of virtual Physarum machine, is a computational counterpart to analog Physarum machines. The latter have been applied over the last 20 yr to solve mazes (Nakagaki et al. 2000), shortest path finding (Nakagaki et al. 2001), transportation network design (Tero et al. 2010), and many other optimal transport and spatial organization problems (Adamatzky 2010), some of which are NP-hard (i.e., among the most computationally difficult).

Jones’s model (Jones 2010) finds an optimal transport network via a feedback loop between a swarm of discrete agents (analogous to slime mold cells) and a continuous lattice of “chemo-attractant” deposited by the agents and the input data. Within the model domain, the agents move along paths determined by the presence of the attractant and emit their own deposit at each time step. We have adapted this model to three dimensions and performed a critical modification: a

probabilistic sampling and selection of the agents’ spatial trajectories. In its original form, Jones’s algorithm imposes that agents’ trajectories maximize the deposit values encountered; our version samples possible paths probabilistically, such that paths toward smaller deposits may still be traversed, although less frequently. This stochastic process formulation enables us to frame the extracted 3D network as a probability density of the presence of filamentary structures, rather than an absolute maximum likelihood solution. We call the resulting framework the Monte Carlo Physarum Machine (MCPM; see the [Appendix](#)).

We have applied the MCPM method to a sample of 37,662 spectroscopically observed galaxies from the SDSS at redshifts $z = 0.0138\text{--}0.0318$. This sample exhibits a number of cosmic web filaments as well as massive galaxy clusters (including Coma) and voids, clearly evident from our 3D visualization tool (Burchett et al. 2019); the volume is also pierced by >500 QSO sightlines observed with the Cosmic Origins Spectrograph (COS) on board the *Hubble Space Telescope* (HST) and that cover the H I Ly α transition over these redshifts. Figure 1 shows our resulting cosmic web reconstruction, where the galaxies effectively serve as “food” sources for a swarm of virtual “slime mold” agents released into a 3D space defined by the celestial coordinates of each galaxy and the luminosity distance implied by each galaxy’s redshift as the radial coordinate. The agents continually move through space and eventually reach an equilibrium state, tracing an approximate optimal transport network from galaxy to galaxy given a small number of model parameters (see the [Appendix](#)).

Essential to the equilibration of the model is a *deposit* emitted by both the galaxies and the agents at each time step; this deposit is sensed by nearby agents, enabling the creation and reinforcement of tentative pathways. Visualized in the top panel of Figure 1, for example, is the *trace*, which records the aggregate agents’ trajectories, and we in turn leverage this trace as a proxy of density values at each point in the 3D space. While the galaxies are directly observed, the purpose of the “agents” is to model the distribution of the (unobserved) dark matter that molds the cosmic web; the “trace” is the map of web filaments and structures found in the converged agent and galaxy distribution.

2.2. Validating the Model with Galaxy Populations

The well-established correlations between galactic star formation and local environment (Balogh et al. 1999; Kauffmann et al. 2004; Peng et al. 2010; Rasmussen et al. 2012; Treyer et al. 2018) and stellar mass (Baldry et al. 2004) provide a natural test for our density reconstruction (see the [Appendix](#) for additional model tests). Crossmatching the input galaxy catalog with our MCPM density field, we present the 2D locus of the “red” galaxy fraction (SDSS color $g-i > 0.8$) as a function of local environmental density and stellar mass in Figure 2. At any fixed stellar mass, the red fraction increases with increasing environmental density; at fixed density, the red fraction increases with stellar mass. A comparison with Figure 6 of Peng et al. (2010) shows qualitatively similar dependences of galaxy quenching on both mass and environment. Our results are also in agreement with more recent studies (Kututma et al. 2017) that report decreasing star formation rates for galaxies residing in the denser regions of filaments.




Figure 2. The dependence of star formation activity on galaxy environment and stellar mass for the galaxies within our SDSS volume. The color coding denotes the fraction of “red” galaxies in the population within each mass/environment bin, where the environmental density (ρ_{Phys}) is determined from our MCPM cosmic web reconstruction algorithm. A comparison with Figure 6 of Peng et al. (2010) shows a similarly increasing red fraction as functions of both mass at fixed density and density at fixed mass.

2.3. H I in the Cosmic Web

A key strength of the Physarum model is that it provides a density value (ρ_{Phys}) for every point in the 3D space (i.e., not only at the locations of galaxies), and each pixel of the HST/COS spectra skewering the volume can be assigned a corresponding density value. We then measure the absorption signal from a given spectral transition as a function of the local density of the cosmic web. Two of the three coordinates needed to localize each voxel within the Physarum datacube are provided by the celestial coordinates of each sightline, and the third is the luminosity distance of each would-be redshift assuming that an H I Ly α line falls at the wavelength of that pixel. Figure 3 shows the H I absorption signal as a function of the local density, where we have binned the HST/COS spectral pixels by the Physarum trace values probed by these spectra and calculated the median normalized flux for pixels in each bin. The absorption signal is expressed as a quantity similar to the optical depth:

$$\tau = \ln\left(\frac{f_b}{f_i}\right) \quad (1)$$

where τ is the absorption measure, f_b is the normalized baseline flux obtained from a linear fit to the absorption measure values at $\log_2 \rho_{\text{Phys}} < 0.8$ (see [Appendix](#)), and f_i is the normalized flux in a given pixel. The blue shading along the curve in Figure 3 indicates 1σ uncertainties on the median from bootstrap resampling. To calibrate the Physarum trace values ρ_{Phys} to cosmologically meaningful quantities, we employed the Bolshoi–Planck (BP) cosmological simulation (Klypin et al. 2016). We fitted the MCPM to a halo catalog (Behroozi et al. 2013) from the simulation and mapped the ρ_{Phys} values from this fit to the dark matter density field (in units of the cosmic mean matter density $\rho_m/\langle\rho_m\rangle$) from the particle data; the




Figure 3. Median H I absorption as a function of cosmic web density. The lower horizontal axis represents density (trace) values from our MPCM model, while the upper axis shows the corresponding density in terms of the $z = 0$ cosmic mean matter density (see the Appendix). Bootstrap errors in the median are shaded about the black median absorption measure curve. Three density regimes are highlighted according to the absorption measure: low density (blue), where the signal is essentially flat with the horizontal dotted line representing a linear fit to those values; intermediate density (green), where the absorption signal monotonically increases with increasing density; and high density (red), where the absorption trend abruptly halts and even reverses.

corresponding $\rho_m / \langle \rho_m \rangle$ for the SDSS ρ_{Phys} values are labeled on the upper axis of Figure 3.

3. Results

Three key features emerge from Figure 3: (1) the “low-density” regime at Physarum model densities $\log_2 \rho_{\text{Phys}} \lesssim 1.2$ (shaded in blue), where we show few statistically significant deviations from the baseline flux value; (2) the “moderate-density” regime at $1.2 < \log_2 \rho_{\text{Phys}} < 1.9$ (shaded in green), where the absorption signal steadily increases; and (3) the “high-density” regime at $\log_2 \rho_{\text{Phys}} > 1.9$ (shaded in red), where the absorption signal no longer increases and shows an *abrupt attenuation* at $\log_2 \rho_{\text{Phys}} \sim 2.3$. The moderate-density regime sets in at approximately the mean matter density of the universe. Figure 4 highlights the low-intermediate density transition, demonstrating that our absorption signal is indeed arising from the far outskirts of filaments and revealing the dark and diffuse fringes of the cosmic web.

Previous studies (e.g., Morris et al. 1993; Tripp et al. 1998; Penton et al. 2002; Stocke et al. 2007; Wakker & Savage 2009; Prochaska et al. 2011; Tejos et al. 2012) have suggested that weak H I absorbers (column densities of $N(\text{H I}) < 10^{13.5} \text{ cm}^{-2}$) may indeed arise from highly underdense regions of the cosmic web (voids). However, the definitions of voids vary widely and are typically expressed in terms of the distance to nearby galaxies, e.g., $>1.4 \text{ Mpc}$ to the nearest L^* galaxy (Stocke et al. 2007). Our method attempts to define structures in terms of the underlying density field. However, for comparison, we cross-matched our sightlines with potentially associated galaxies. In the density bin where our detected signal sets in (near $\langle \rho_m \rangle$), we find that the Ly α absorbing spectral regions are generally $>2.5 \text{ Mpc}$ from the nearest L^* galaxy within $\pm 500 \text{ km s}^{-1}$. We posit that some fraction of the few void absorbers previously

reported may indeed be outliers in the distribution of H I in underdense regions and represent rare large-density perturbations uncorrelated with cosmic web filaments traced by galaxies (Tejos et al. 2012, 2014). However, the monotonic rise of the absorption signal in Figure 3 through the intermediate-density regime suggests that weak Ly α forest absorbers trace the outskirts of filaments, as also shown in Figure 4, smoothly increasing in strength with the ambient density field as would be expected from results of hydrodynamical simulations (e.g., Davé et al. 2010).

In Figure 3, the absorption measure in the moderate-density regime increases with local density until it abruptly halts near $\log_2 \rho_{\text{Phys}} = 1.9$, corresponding to $\rho_m / \langle \rho_m \rangle \sim 200$. In the regime where the absorption measure is steadily increasing (green shaded region in Figure 3), the hydrogen absorption segues from arising in truly intergalactic material that has no physical connection with any nearby galaxies to circumgalactic material that plausibly interacts with the closest galaxies. Crossmatching sightlines probing this regime with nearby galaxies suggests that the high-density region sets in at ≈ 2 virial radii from nearby galaxies. Here, galaxy feedback and other hydrodynamical effects are likely modulating gas physical conditions within the CGM, heralding a physically distinct regime from the surrounding IGM. The CGM as traced by H I and heavy element ions is heavily dependent on host galaxy properties (Bordoloi et al. 2011; Tumlinson et al. 2011, 2013; Nielsen et al. 2013; Lan et al. 2014; Johnson et al. 2015; Burchett et al. 2016), such as mass, star formation rate, and environment. The same processes reflected in these CGM tracers plausibly impart the sudden deviation from the intermediate-density absorption trend.

Within the high-density regime (shaded red in Figure 4), the H I absorption remarkably decreases at $\log_2 \rho_{\text{Phys}} \sim 2.3$. This density regime includes galaxy clusters and the inner




Figure 4. A detailed inspection of the IGM transition region. Top left: a heatmap visualization of the full density field in a 3 Mpc thick slice from our MCPM fit to the SDSS galaxy sample. Top right: the same 3 Mpc slice segmented and color coded by the three density regimes: low-, moderate-, and high-densities are colored blue, green, and red, respectively, as in Figure 3. Bottom: highlighting three bins in density bracketing the low-moderate density transition shown in Figure 3 within a single spatial region. The deposits of the galaxies input to the MCPM algorithm are shown in magenta, and the MCPM trace is colored green in each panel with corresponding $\log_2 \rho_{\text{phys}}$ values indicated below each panel. The left panel shows that the density regime just before the onset of the HI absorption signal includes regions within putative voids, whereas the center and right panels demonstrate that the onset of significant HI absorption occurs in the outskirts of filaments.

regions (or “spines”) deep within filaments. In addition to the circumgalactic effects mentioned above, additional environmental factors are likely at play here. First, within galaxy clusters, HI absorption is highly suppressed, even in sightlines probing near galaxies (Yoon & Putman 2013, 2017; Burchett et al. 2018; Butsky et al. 2019). The intracluster medium (ICM) is mostly composed of hot $>10^6$ K gas, for which the neutral hydrogen fraction is extremely low; furthermore, the cool gas abundant in the CGM of galaxies outside clusters can be readily stripped upon falling into the ICM (Tonnesen et al. 2007; Bahé et al. 2013; Zinger et al. 2018). Second, the filamentary environment itself is predicted to heat the intergalactic gas, which in turn will further ionize the gas and/or broaden the HI absorption profiles, effectively reducing the absorption strength per pixel (Richter et al. 2006, 2008; Tepper-García et al. 2012; Wakker et al. 2015; Tejos et al. 2016; Tonnesen et al. 2017). An additional effect may occur in the moderate- and high-density regimes arising from high-peculiar velocity flows relative to the galaxies we are using to trace the large-scale structure. This may redistribute flux to pixels attributed to higher- or lower-density regimes than that of its physical origin.

By employing the MCPM cosmic web reconstruction technique, we have characterized the neutral hydrogen content in the low-redshift universe over four orders of magnitude in matter density. We have shown that HI absorption from the IGM tracks the large-scale structure of the universe and smoothly increases from the outskirts of cosmic web filaments toward higher densities until processes in the CGM, ICM, and dense cores of filaments suppress the signal. Regarding the application of MCPM (rooted in biology) within an astrophysical context, we emphasize that the Physarum behaviors modeled by MCPM do not attempt to model the physical origins of structure formation, i.e., gravity and dark energy. Rather, the outcomes of these disparate modeled processes are related by their formation of optimized network structures (Adamatzky 2010; Hong et al. 2016). Furthermore, the agents in our model follow trajectories throughout the structure, although these trajectories will differ from those traveled by galaxies in cosmic web filaments and galaxy clusters. However, we are analyzing these agent trajectories to glean further insight about the cosmic web structure (e.g., filament identification) for a subsequent publication. Future applications of MCPM will more deeply examine the intimate connections galaxies have

with their large-scale environments and the resulting symbiosis with their gaseous surroundings.

We are grateful to Jan Ivanecký, who authored an early prototype of the simulation code and remained a source of valuable technical advice, and Daisuke Nagai, who provided key helpful suggestions. The simulation art of Sage Jenson inspired the idea of possible connection between the morphology of *Physarum polycephalum* and the IGM structure. BP simulation data were kindly provided by Joel Primack and Doug Hellinger. Support for this work was provided by NASA through grant No. HST-AR 15009 from the Space Telescope Science Institute, which is operated by AURA, Inc., under NASA contract NAS 5-26555.

Appendix

A.1. Observations and Simulation Data

Our analysis employs two key observational data sets: (1) the NASA/Sloan Atlas (NSA), a value-added catalog of SDSS photometry and spectroscopy, and (2) the Hubble Spectroscopic Legacy Archive (HSLA; Peebles et al. 2017), a database of reduced and coadded spectra from COS on board *HST*. We estimated star formation rates, halo masses, and virial radii for all galaxies within the NSA to accompany the stellar masses provided in the original catalog (Burchett et al. 2016). We focus our analysis on a volume including a redshift range of $z = 0.0138\text{--}0.0318$ over the spring SDSS footprint. This sample was chosen to feature several attributes: (a) a rich diversity in large-scale structures including galaxy clusters, filaments, and voids; (b) a high spectroscopic completeness to faint dwarf galaxies; and (c) a large number of sightlines to background QSOs with *HST* archival spectra. Indeed, the volume that we analyze includes five Abell clusters (Abell 1965), including Coma, and several galaxy filaments apparent from visual inspection within SDSS (Figure 1). Our redshift slice, centered on Coma’s brightest cluster galaxies, includes approximately ± 3 times the velocity dispersion of Coma (Fuller et al. 2016) and is spectroscopically complete down to $0.05 L^*$ (assuming the nominal magnitude limit of SDSS galaxy spectra; Strauss et al. 2002). Because the HSLA does not include QSO redshifts within the main database, we crossmatched all HSLA QSOs with the Million Quasars Catalog (Flesch 2015) v5.2 to obtain their redshifts. We then selected those QSOs with sightlines within the survey volume and with $z > 0.0318 + \Delta v/c$, for $\Delta v = 5000 \text{ km s}^{-1}$ and where c is the speed of light, to select sightlines probing the full volume and mitigate confusing foreground absorption with outflow signatures of distant QSOs (Tripp et al. 2008). For the 346 QSO spectra meeting the criteria above and with a signal-to-noise ratio (S/N) > 3 , we fit continua over the wavelength range covering Ly α 1215.67 Å at $z = 0.0138\text{--}0.0318$ using an automated procedure, which iteratively fits series of Legendre polynomials to the data while sigma clipping to reject pixels likely due to physical absorption features from the fit. The order of the polynomial series fitted is determined by an F-test.

We ultimately identify large-scale structure and quantify the local density using our MCPM slime mold model. To calibrate the density values extracted from this *Physarum* model with cosmologically meaningful values, we employ the dark matter only BP cosmological simulation (Klypin et al. 2016; Roca-Fàbrega et al. 2019). The BP simulation volume comprises

$(250 \text{ h}^{-1} \text{ Mpc})^3$ and contains 2048³ particles; we utilize the $z = 0$ snapshot. Conveniently, the physical dimensions of the simulation exceed the luminosity distances of the most distant galaxies in our sample. Therefore, we truncated the simulation data to a width of approximately 70 Mpc in one direction centered on the median luminosity distance of our galaxy sample, which would ease the computational load when fitting the BP data with the MCPM algorithm. Specifically, we used two BP data products in our analysis: a halo catalog produced using the Rockstar algorithm (Behroozi et al. 2013) and a datacube containing the dark matter density field smoothed over 0.25 Mpc scales (Lee et al. 2017; Goh et al. 2019).

A.2. The MCPM Model

The employed MCPM model runs in a time-discrete loop, continually producing tentative fits constrained by the input data. We stop the fitting procedure when the model reaches a soft equilibrium, that is when the solution changes only at the level of process variance.

The core of MCPM consists of two alternating steps: *propagation* and *relaxation*. Being a hybrid model, each step is centered around one of the two MCPM’s modalities, as follows.

1. The *propagation step* (detailed in Figure 5) iterates over a discrete ensemble of agents, which are the model’s devices for exploring the simulation domain. Each agent’s state is represented by a position and a movement direction, which are updated in parallel to navigate through the *deposit field* (see Figure 6(a)). The deposit field is stored as a piecewise-continuous 3D lattice of scalar values, representing both the input data as well as the agents’ spatial density. The deposit effectively guides the agents, as these are propagated with higher likelihood to the regions where deposit values are greater. In addition to the deposit, we also maintain a scalar *trace field* (Figure 6(d)), which serves only to record the the agents’ spatio-temporal density, but does not participate in their guiding.
2. The *relaxation step* ensures that the simulation eventually reaches an equilibrium. The deposit field is diffused spatially by a small uniform kernel and decays—that is, each cell is multiplied by a value < 1 . The trace field also decays but does not diffuse in order to preserve features in the estimate of the agents’ distribution, which the trace represents.

The input to MCPM is a set of weighted points representing either galaxies, for the SDSS data set, or dark matter halos, for the BP simulation. These points are static and—just like the agents—emit and inject an amount of deposit at each simulation step, but at the same points in space. Unlike the agents however, the data deposit is not constant but proportional to each input object’s mass. Because we typically use 5–50 times more agents than input data points (see “Application of MCPM” below), the deposit emitted by the agents is also significantly smaller (subject to fine-tuning).

MCPM’s main output is the trace field, which we use in all our subsequent analysis as the indicator of cosmic web density. The trace effectively superimposes the trajectories of all agents throughout the recent fitting history, therefore representing their time-amortized density in space. Please refer to Figure 5 for a schematic depiction of both deposit and trace, and to Figure 6 for 3D visualizations of the involved quantities.


```

def propagation_step(in params, inout agents[], inout deposit[][][], out trace[][][]):
  in parallel for each agent i:
    ## sensing phase: panel (a)
    dir_s0 = agents[i].dir
    dir_s1 = random_direction(params.sense_angle)
    dist_s = random_distance(params.sense_distance)
    p0 = deposit[agents[i].pos + dist_s * dir_s0] ** params.sharpness
    p1 = deposit[agents[i].pos + dist_s * dir_s1] ** params.sharpness

    ## sampling phase: panel (b)
    if (unit_random() < p0 / (p0+p1)):
      dir_m = dir_s0
    else:
      dir_m = rotate_towards(dir_s0, dir_s1, params.move_angle)
      dist_m = random_distance(params.move_distance)

    ## update phase: panel (c)
    agents[i].dir = dir_m
    agents[i].pos += dist_m * dir_m
    deposit[agents[i].pos] += params.agent_deposit
    trace[agents[i].pos] += 1
    
```

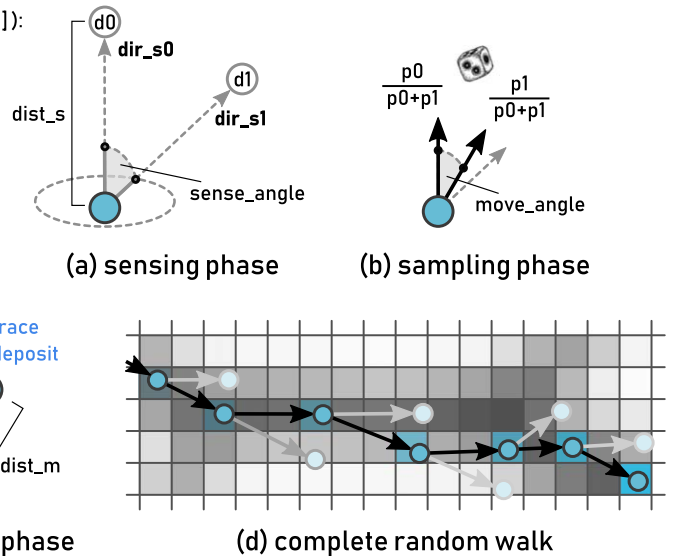


Figure 5. Pseudocode and a schematic depiction of the MCPM’s agent propagation step. In the pseudocode, function calls are highlighted orange, arrays and matrices in blue, structures in green, vectors in boldface. Each propagation phase is illustrated in the corresponding diagrams (a–c). Diagram (d) shows a sample random walk of an agent, coupled with the simulation lattice containing the deposit (grayscale) and the trace values (light blue). We see that the agent’s trajectory is a chain of binary directional decisions. Each propagation step is followed by a relaxation step as explained in the text.




Figure 6. Visualizing the elements of the MCPM model. (a) Deposit “halos” emitted by the input data (red points). (b) Particle visualization of the agents (white points) navigating through the input data (red points). (c) Resulting agents’ trace (blue) superimposed over the input deposit halos (purple). (d) Visualizing only the trace using the same color definitions as in Figure 10.

A.3. Relation of MCPM to Previous Work

Being based on the virtual Physarum machine model of Jones (2010), MCPM shares a number of its features and we encourage the reader to refer to this seminal paper and complement our exposition. Here we examine the extensions to this model we have made to facilitate our cosmic web reconstruction.

1. Generalization from 2D to 3D. The extension of Jones’s model to 3D follows with minor modifications, although the obtained network morphologies are quite different due to the peculiar topological properties of 3D space (as shown in Figure 7). The only notable modification is the directional sampling (see Figure 5(a) and the random_direction function in the pseudocode therein): we now need to sample a cone of directions rather than simply examine two directions at fixed angles from the agent’s current propagation direction.
2. Stochastic Monte-Carlo exploration of the domain. Jones’s agents (Jones 2010) always navigate toward the local maxima of the deposit (thus we will be referring to this model as Max-PM). This yields very detailed solutions (example shown in Figure 7, right), but results

in a highly nonlinear behavior especially when the fitting constraints are complex distributions of many points (we typically fit to 10^4 – 10^5 input points using millions of agents).

Instead, we view the simulation as a *stochastic process*: in the propagation step, the deposit is interpreted as a probability distribution that guides the agents. This has multiple implications illustrated in Figure 5. First, rather than using fixed distances and directions in the propagation step, we generate them probabilistically: drawing uniformly from a cone for directional sampling (function random_direction) and drawing from the 1D Maxwell–Boltzmann distribution for distance sampling (function random_distance). Second, rather than probing the space of directions more densely (Jones 2010) we rely on a single *binary decision*. That is (using the unit_random function, Figure 5(b)), we simply choose between mutating or not mutating the propagation direction. This design is commonly used, e.g., in Markov Chain Monte Carlo methods.

In order to convert the deposit values to sampling probabilities (p_0 and p_1 in the pseudocode) we




Figure 7. Comparison of our MCPM model to Jones’s Max-PM model in a 30 Mpc slice from the BP data set. Our Max-PM adaptation to 3D uses nine directional probes for the deposit sampling (see Jones 2010), with all other parameters kept identical with MCPM. Different *sharpness* values in MCPM trade filament definition for structural complexity (left vs. center); Max-PM (right) being the limit case with sharp (nearly “tubular”) filaments, but lowest geometric complexity and connectivity at filament intersections. That in turn causes about 5% of input halos to be missed by the Max-PM fit, compared to the negligible 0.05% of MCPM.

opted for a geometric proportionality controlled by a single *sharpness* parameter $\in \mathbb{R}^+$. As demonstrated in Figure 7, this grants direct control over the actual sharpness (acuity, or falloff) of the filamentary structures: for *sharpness* $\rightarrow 0$ the agents diffuse uniformly and all anisotropic structures gradually disappear; to the other extreme, the model slowly converges to the Max-PM behavior (Jones 2010) as *sharpness* $\rightarrow \infty$.

3. Extraction of the agent trace. While the deposit field could serve as an indicator function for the cosmic web density reconstruction, this would not yield optimal results because the relaxation step effectively acts as a low-pass filter. Instead, we introduce the trace field, which is a direct recording of the temporally averaged (aggregate) agent trajectories. Thanks to the stochastic formulation of our model, the trace is equivalent to a probability density function (after proper normalization), and serves as a suitable basis for constructing mappings to the cosmic overdensity and other relevant cosmological quantities.

A.4. Application of the MCPM Model

Similarly to the actual *Physarum polycephalum* organism, MCPM and its other virtual realizations produce a set of pathways that—with respect to a relevant measure in the operating domain—form approximate optimal transport networks spanning the input data (“food”). The conjecture that gravitationally relaxed matter has the tendency to form such networks is indirectly supported by our findings.

Our implementation of MCPM is written in C++ and uses GPU compute shaders for executing the parallel propagation and relaxation steps, as well as the volume visualization. The GPU implementation allows us to run the fitting with 2M agents over 960^3 deposit and trace grids at interactive speeds using an NVIDIA TitanX GPU. The model typically converges in fewer than 1000 timesteps (about 1 minute). All of the model’s meta-parameters (see pseudocode in Figure 5, and

Table 1
Parameter Values used in the MCPM Fitting

<i>sense_angle</i>	20 (deg)	<i>sharpness</i>	3–7
<i>sense_distance</i>	2.5 (Mpc)	<i>agent_deposit</i>	0.005–0.05
<i>move_angle</i>	10 (deg)	<i>data_deposit</i>	$M_{\odot}/10^{11}$
<i>move_distance</i>	0.2 (Mpc)		

Table 1 for typical values) can be modified at any point during the simulation, which allows for easy tuning of the obtained 3D structures.

Of critical importance is the question regarding MCPM’s *robustness* when applied to data sets with different spatial densities. This will become particularly important when applying MCPM to galaxy samples at higher redshift in flux-limited surveys, such as the SDSS, which will contain only more massive, rarer galaxies. We therefore examine this robustness in Figure 8, where we fit the MCPM model to three subsamples of the BP halo catalog: including all $M > 10^{11} M_{\odot}$ halos ($\sim 456,000$ objects), including all $M > 10^{12} M_{\odot}$ halos ($\sim 57,000$ objects), and randomly subsampling the $M > 10^{11} M_{\odot}$ catalog by a factor of 1/4 ($\sim 114,000$ objects). For this comparison, we simply adopted the same MCPM parameters for all fits without additional optimization. Figure 8 demonstrates that MCPM finds salient filamentary structures even when the input data is significantly reduced. Crucially, if configured well, the model does not *imagine* nonexistent structures: if the input data is too sparse with regard to the characteristic feature scales (per given geometric configuration), the agents will either be near-uniformly dispersed in space, or form only “fuzzy” tentative (low-probability) filaments. This is another desirable property of MCPM relative to the Max-PM model (Jones 2010), which has no way to assess the “confidence” in the reconstructed structures.

We more quantitatively assess the impact of sparsely sampled data by comparing the trace values from the mass-downsampled fit shown in Figure 8 (right) with the more




Figure 8. Robustness evaluation for the MCPM model. We compare fits starting from the 456 k most massive halos ($\geq 10^{11} M_{\odot}$) from the BP data set in the left column. We apply two decimation schemes: uniform subsampling to a halo *density* comparable to our SDSS data set (middle column), and selective subsampling to a comparable halo *count* using only halos more massive than $10^{12} M_{\odot}$ (right column). In both reduced cases, some fine-level details are lost, but the overall structure is well preserved (see sample insets).

densely sampled fit (left; hereafter referred to as the reference fit). Upon extracting trace volumes from each fit, we standardized each distribution by subtracting its mean and dividing by its standard deviation. We then sampled the reference fit in 110 bins of $\Delta \log_2 \rho_{\text{Phys}} \sim 0.07$, identifying 2000 cells with trace values in each bin, and then logged the trace values at identical locations in the comparison fit. Figure 9 shows the fractional residuals ($[\text{comparison-reference}]/\text{reference}$) as a function of reference fit trace (bottom panel) as well as the global reference trace distribution (top panel). Even without optimizing the fits for the downsampled data set, we note that the reference and comparison fits agree well at the peak of the ρ_{Phys} distribution. Within the lower tail of the distribution, the two fits diverge from one another, suggesting that including only more massive halos will poorly constrain low-density structures. This result is not surprising, although we underscore that the comparison fit was not tuned even given the dramatic decrease (by a factor of 8) in data points used in the fit. A more rigorous performance and optimization analysis of the model under varying conditions (including sampling density) will be presented in a subsequent publication.

A.5. MCPM Trace/BP Matter Density Calibration

The “trace” output of the Physarum model effectively provides density values (ρ_{Phys}) throughout the volume sampled by our SDSS galaxies. To calibrate these values to more

cosmologically meaningful terms of the cosmic matter density ($\rho_m/\langle\rho_m\rangle$), we fitted the MCPM model to the halo catalog for the BP $z = 0$ snapshot produced by the Rockstar algorithm (Behroozi et al. 2013; Figure 10). We then sampled the trace values within this volume in logarithmic bins of $\Delta \log_2 \rho_{\text{Phys}} = 0.05$, crossmatching with the cube of smoothed matter density values, and produced a mapping between the two via the running median, as shown in Figure 11.

We highlight two distinct features of this mapping: (1) the relatively flat relation for $\log_2 \rho_{\text{Phys}} \lesssim 1$ is a result of the larger dynamic range spanned by MCPM trace values than matter density values of the BP data product. (2) The MCPM trace and BP matter densities are well correlated for $\log_2 \rho_{\text{Phys}} \gtrsim 1$. It is in this regime where the mapping shown on the upper horizontal axis of Figure 3 panels is well matched. Consequently, the densities just below and including where our detected absorption signal occurs also has the most robust mapping between ρ_{Phys} and ρ_m . Finally, we standardized the overall distributions of ρ_{Phys} between the MCPM fits of the BP halos and SDSS galaxies. For each distribution, we calculated its mean and standard deviation then subtracted the mean and divided by the standard deviation.

A.6. H I Absorption Signal

As briefly sketched in the main text, our MCPM cosmic web density reconstruction provides a ρ_{Phys} value at every point in




Figure 9. Slicing through the density field reconstructed from the most massive 456 k halos ($\geq 10^{11} M_{\odot}$) of the BP simulation volume. Top: false-color visualization of the mapped overdensity. Bottom: spatial distribution of the trace density values, divided into three intervals.

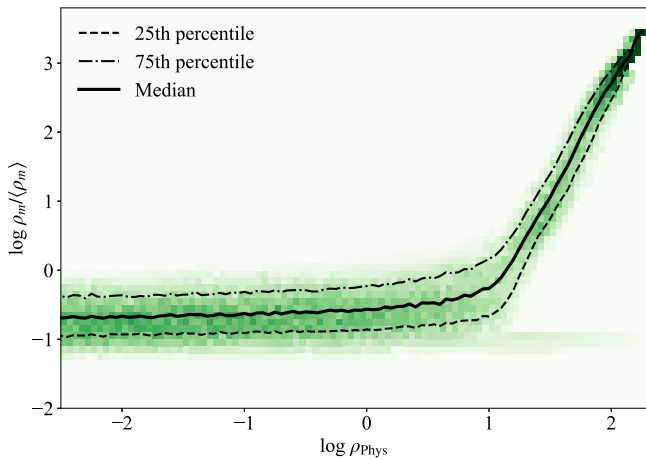


Figure 10. Calibration of the Physarum densities with dark matter densities from the BP simulation. Within bins of MCPM trace values (ρ_{Phys}) from our fit to the BP halo catalog, we matched the spatial locations of points in the MCPM volume to local density values from the BP particle data. This provides a mapping to the cosmic mean matter density ($\rho_m/\langle\rho_m\rangle$). The bold black curve shows the median ρ/ρ_m values in each bin, and the dashed and dashed-dotted black curves denote the 25th and 75th percentiles, respectively, of each ρ_{Phys} bin’s distribution.

the 3D volume, which in turn provides a ρ_{Phys} value for each pixel in the *HST*/*COS* spectra probing the volume. The galaxy survey data we input to MCPM includes the 2D celestial coordinates plus

the luminosity distance from the galaxies’ redshifts. Therefore, we mapped coordinates in the output MCPM density field to spectral pixels using the celestial coordinates of each *HST*/*COS* sightline and adopted as each pixel’s “distance” the luminosity distance at $z_{\text{pix}} = \lambda_i/1215.67 \text{ \AA} - 1$, where λ_i is the observed wavelength of each pixel and 1215.67 Å is the rest wavelength of H I Ly α . We then tabulated the normalized flux ρ_{Phys} values for each pixel in each of the 346 sightlines in our sample.

The observed wavelength range for H I Ly α corresponding to the redshifts of our galaxy sample is 1231.25–1257.96 Å, which includes several transitions observed in the Milky Way’s interstellar medium (ISM): NV $\lambda\lambda 1238, 1242 \text{ \AA}$; Mg II $\lambda\lambda 1239, 1240 \text{ \AA}$; and S II $\lambda\lambda 1250, 1253 \text{ \AA}$. To avoid introducing a spurious signal due to absorption from these transitions, we masked out the following spectral regions from our analysis (all wavelengths in Å): (1238.00, 1240.60), (1242.10, 1243.20), (1250.20, 1250.95), and (1253.30, 1254.30).

We then calculated the median normalized flux of the spectral pixels in bins of $\Delta \log_2 \rho_{\text{Phys}} = 0.4$ (prior to the standardization step described above). We adopted the standard deviation of medians from 1000 pixel bootstrap resamples as the uncertainty on the median. Although we had masked spectral regions likely impacted by the absorption from the Milky Way ISM, a residual, spurious absorption signal certainly remains from interloping absorbers at redshifts $z > 0$. This signal was apparent upon our initial calculation




Figure 11. Comparing the model trace values for MCPM fits to BP catalogs including halos with $M > 10^{11} M_{\odot}$ (reference fit) and with $M > 10^{12} M_{\odot}$ (comparison fit). The bottom panel shows residuals between standardized trace values (ρ_{Phys}) found in the comparison and reference fits ($(\text{comparison-reference})/\text{reference}$) at the same spatial locations. The upper histogram shows the distribution of ρ_{Phys} in the reference fit. Included in the bottom panel are points sampled evenly in bins of ρ_{Phys} to highlight the relationship across the full dynamic range, and both panels are color coded to emphasize the relative frequency of points in each ρ_{Phys} bin. We note agreement between the models near the peak of the distribution, while the downsampled model fit diverges in the lower tail of the distribution.




Figure 12. Density-dependent absorption signal juxtaposed with nearby galaxy statistics As in Figure 3, the black curve and blue shading shows the absorption measure as a function of environmental density. The green curve shows the median of the impact parameters normalized to the virial radii of closest galaxies to the sightline pixels represented in each density bins (values are shown in the right-hand vertical axis). A vertical dashed line marks the cosmic mean matter density and the approximate beginning of the intermediate-density regime.

of the binned median fluxes, particularly at low densities ($\log_2 \rho_{\text{Phys}} < 0.8$), where the median normalized flux remained relatively flat with increasing density but was offset from unity. We verified that the offset was indeed due to interloping absorption by isolating a subset of the sightlines for which we had identified the absorption lines at all redshifts. We then calculated the median flux for the sightlines in this subset with known interloping absorption lines having observed

wavelength within the range analyzed and for those sightlines without such interlopers. Indeed, upon comparing the absorption profiles, the residual signal arose significantly in sightlines with interloping absorbers but did not in those free of interlopers. In the full data set, we corrected for this spurious absorption signal as follows. To establish the baseline flux f_b , we performed a linear least-squares fit to the median flux values at $\log_2 \rho_{\text{Phys}} < 0.8$. The fitted slope was 0.00026 ± 0.00097 , suggesting no statistically significant change in the signal over this range. We adopted the y-intercept ($f = 0.981$) of the fit as f_b in Equation (1). The offset of this baseline from a normalized Ly α flux of 1 (i.e., no H I absorption) is mostly due to interloping lines from other redshifts.

To better understand the galaxy environments probed by our absorption data, we crossmatched each spectral pixel with the SDSS galaxy catalog, searching for galaxies within 5 Mpc of the sightline and having redshifts within $\pm 500 \text{ km s}^{-1}$ of z_{pix} . For each pixel, we found the “nearest” galaxy according to two criteria: (1) that with the smallest impact parameter in proper distance and (2) that with the smallest impact parameter normalized to the galaxy’s virial radius (Burchett et al. 2016). In the lowest ρ_{Phys} bins ($\log_2 \rho_{\text{Phys}} < 1.5$), no galaxies were found within 5 Mpc for a large fraction of the pixels, as expected. Similar to the normalized flux, we then calculated the medians and uncertainties of the smallest impact parameters in each ρ_{Phys} bin. Figure 12 shows, in addition to the absorption signal from Figure 3, the median of the impact parameter to the nearest galaxy within each bin (right vertical axis), where the impact parameters have been normalized by galaxy virial radius (R_{vir}). As we alluded in the main text, the high-density regime begins where the spectra typically probe within $\sim 1-2 R_{\text{vir}}$, whereas the intermediate-density regime begins beyond $\sim 4-5 R_{\text{vir}}$ of any galaxy.

ORCID iDs

Joseph N. Burchett <https://orcid.org/0000-0002-1979-2197>
 Oskar Elek <https://orcid.org/0000-0003-0549-3302>
 Nicolas Tejos <https://orcid.org/0000-0002-1883-4252>
 J. Xavier Prochaska <https://orcid.org/0000-0002-7738-6875>
 Todd M. Tripp <https://orcid.org/0000-0002-1218-640X>
 Rongmon Bordoloi <https://orcid.org/0000-0002-3120-7173>

References

- Abell, G. O. 1965, *ARAA*, 3, 1
 Adamatzky, A. 2010, *Physarum Machines* (Singapore: World Scientific)
 Aragón-Calvo, M. A., Jones, B. J. T., van de Weygaert, R., & van der Hulst, J. M. 2007, *A&A*, 474, 315
 Bahé, Y. M., McCarthy, I. G., Balogh, M. L., & Font, A. S. 2013, *MNRAS*, 430, 3017
 Baldry, I. K., Glazebrook, K., Brinkmann, J., et al. 2004, *ApJ*, 600, 681
 Balogh, M. L., Morris, S. L., Yee, H. K. C., Carlberg, R. G., & Ellingson, E. 1999, *ApJ*, 527, 54
 Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *ApJ*, 762, 109
 Bergeron, J., & Boissé, P. 1991, *A&A*, 243, 344
 Bond, J. R., Kofman, L., & Pogosyan, D. 1996, *Natur*, 380, 603
 Bordoloi, R., Lilly, S. J., Knobel, C., et al. 2011, *ApJ*, 743, 10
 Burchett, J. N., Abramov, D., Otto, J., et al. 2019, *Computer Graphics Forum*, 38, 491
 Burchett, J. N., Tripp, T. M., Bordoloi, R., et al. 2016, *ApJ*, 832, 124
 Burchett, J. N., Tripp, T. M., Wang, Q. D., et al. 2018, *MNRAS*, 475, 2067
 Butsky, I. S., Burchett, J. N., Nagai, D., et al. 2019, *MNRAS*, 490, 4292
 Cautun, M., van de Weygaert, R., & Jones, B. J. T. 2013, *MNRAS*, 429, 1286
 Cen, R., Miralda-Escudé, J., Ostriker, J. P., & Rauch, M. 1994, *ApJL*, 437, L9
 Chen, H.-W., & Mulchaey, J. S. 2009, *ApJ*, 701, 1219
 Chen, Y.-C., Ho, S., Brinkmann, J., et al. 2016, *MNRAS*, 461, 3896

- Chen, Y.-C., Ho, S., Mandelbaum, R., et al. 2017, *MNRAS*, **466**, 1880
- Davé, R., Hernquist, L., Katz, N., & Weinberg, D. H. 1999, *ApJ*, **511**, 521
- Davé, R., Oppenheimer, B. D., Katz, N., Kollmeier, J. A., & Weinberg, D. H. 2010, *MNRAS*, **408**, 2051
- de Lapparent, V., Geller, M. J., & Huchra, J. P. 1986, *ApJL*, **302**, L1
- Flesch, E. W. 2015, *PASA*, **32**, e010
- Fuller, C., Davies, J. I., Smith, M. W. L., et al. 2016, *MNRAS*, **458**, 582
- Goh, T., Primack, J., Lee, C. T., et al. 2019, *MNRAS*, **483**, 2101
- Hong, S., Coutinho, B. C., Dey, A., et al. 2019, *MNRAS*, **459**, 2690
- Johnson, S. D., Chen, H.-W., & Mulchaey, J. S. 2015, *MNRAS*, **449**, 3263
- Jones, J. 2010, *Artificial Life*, **16**, 127
- Kauffmann, G., White, S. D. M., Heckman, T. M., et al. 2004, *MNRAS*, **353**, 713
- Keeney, B. A., Stocke, J. T., Rosenberg, J. L., et al. 2013, *ApJ*, **765**, 27
- Klypin, A., Yepes, G., Gottlöber, S., Prada, F., & Heß, S. 2016, *MNRAS*, **457**, 4340
- Kuutma, T., Tamm, A., & Tempel, E. 2017, *A&A*, **600**, L6
- Lan, T.-W., Ménard, B., & Zhu, G. 2014, *ApJ*, **795**, 31
- Lanzetta, K. M., Bowen, D. V., Tytler, D., & Webb, J. K. 1995, *ApJ*, **442**, 538
- Lee, C. T., Primack, J. R., Behroozi, P., et al. 2017, *MNRAS*, **466**, 3834
- Libeskind, N. I., van de Weygaert, R., Cautun, M., et al. 2018, *MNRAS*, **473**, 1195
- Morris, S. L., Weymann, R. J., Dressler, A., et al. 1993, *ApJ*, **419**, 524
- Nakagaki, T., Yamada, H., & Tóth, Á. 2000, *Natur*, **407**, 470
- Nakagaki, T., Yamada, H., & Tóth, Á. 2001, *Biophysical Chemistry*, **92**, 47
- Nielsen, N. M., Churchill, C. W., & Kacprzak, G. G. 2013, *ApJ*, **776**, 115
- Oppenheimer, B. D., & Davé, R. 2006, *MNRAS*, **373**, 1265
- Peeples, M., Tumlinson, J., & Fox, A. 2017, The Hubble Spectroscopic Legacy Archive, Instrument Science Report COS, 2017-4, STScI
- Peng, Y.-j., Lilly, S. J., Kovač, K., et al. 2010, *ApJ*, **721**, 193
- Penton, S. V., Stocke, J. T., & Shull, J. M. 2002, *ApJ*, **565**, 720
- Pessa, I., Tejos, N., Barrientos, L. F., et al. 2018, *MNRAS*, **477**, 2991
- Prochaska, J. X., Burchett, J. N., Tripp, T. M., et al. 2019, *ApJS*, **243**, 24
- Prochaska, J. X., Kasen, D., & Rubin, K. 2011, *ApJ*, **734**, 24
- Rasmussen, J., Mulchaey, J. S., Bai, L., et al. 2012, *ApJ*, **757**, 122
- Richter, P., Paerels, F. B. S., & Kaastra, J. S. 2008, *SSRv*, **134**, 25
- Richter, P., Savage, B. D., Sembach, K. R., & Tripp, T. M. 2006, *A&A*, **445**, 827
- Roca-Fàbrega, S., Dekel, A., Faerman, Y., et al. 2019, *MNRAS*, **484**, 3625
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521
- Sousbie, T., Pichon, C., & Kawahara, H. 2011, *MNRAS*, **414**, 384
- Stark, D. V., Kannappan, S. J., Eckert, K. D., et al. 2016, *ApJ*, **832**, 126
- Stocke, J. T., Danforth, C. W., Shull, J. M., Penton, S. V., & Giroux, M. L. 2007, *ApJ*, **671**, 146
- Strauss, M. A., Weinberg, D. H., Lupton, R. H., et al. 2002, *AJ*, **124**, 1810
- Tejos, N., Morris, S. L., Crighton, N. H. M., et al. 2012, *MNRAS*, **425**, 245
- Tejos, N., Morris, S. L., Finn, C. W., et al. 2014, *MNRAS*, **437**, 2017
- Tejos, N., Prochaska, J. X., Crighton, N. H., et al. 2016, *MNRAS*, **455**, 2662
- Tempel, E., Stoica, R. S., Martínez, V. J., et al. 2014, *MNRAS*, **438**, 3465
- Tepper-García, T., Richter, P., Schaye, J., et al. 2012, *MNRAS*, **425**, 1640
- Tero, A., Tagaki, S., Saigusa, T., et al. 2010, *Sci*, **327**, 439
- Tonnesen, S., Bryan, G. L., & van Gorkom, J. H. 2007, *ApJ*, **671**, 1434
- Tonnesen, S., Smith, B. D., Kollmeier, J. A., & Cen, R. 2017, *ApJ*, **845**, 47
- Treyer, M., Kraljic, K., Arnouts, S., et al. 2018, *MNRAS*, **477**, 2684
- Tripp, T. M., Lu, L., & Savage, B. D. 1998, *ApJ*, **508**, 200
- Tripp, T. M., Sembach, K. R., Bowen, D. V., et al. 2008, *ApJS*, **177**, 39
- Tumlinson, J., Peeples, M. S., & Werk, J. K. 2017, *ARAA*, **55**, 389
- Tumlinson, J., Thom, C., Werk, J. K., et al. 2011, *Sci*, **334**, 948
- Tumlinson, J., Thom, C., Werk, J. K., et al. 2013, *ApJ*, **777**, 59
- Wakker, B. P., Hernandez, A. K., French, D. M., et al. 2015, *ApJ*, **814**, 40
- Wakker, B. P., & Savage, B. D. 2009, *ApJS*, **182**, 378
- Yoon, J. H., & Putman, M. E. 2013, *ApJL*, **772**, L29
- Yoon, J. H., & Putman, M. E. 2017, *ApJ*, **839**, 117
- Zinger, E., Dekel, A., Kravtsov, A. V., & Nagai, D. 2018, *MNRAS*, **475**, 3654