Laidlaw Research Project

Disentangling Cosmological and Feedback Drivers of the Cosmic Web

Author: Ross McPhee

Supervisor: Professor Rita Tojeiro

School of Physics and Astronomy, University of St Andrews

rm378@st-andrews.ac.uk

Date: 1 September 2025

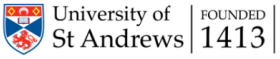

Abstract

This project investigates how variations in both cosmological and astrophysical parameters shape the large-scale structure of the Universe, known as the cosmic web. Using the CAMELS (Cosmology and Astrophysics with MachinE Learning Simulations) suite, specifically the IllustrisTNG 1P hydrodynamical simulations, we explore the impact of matter density (Ω_m), fluctuation amplitude (σ_8), and feedback processes from stars and black holes (ASN and AGN parameters). Two-dimensional slices from the simulations were systematically compared to fiducial cases through residual analysis, quantifying changes in neutral hydrogen (HI), total gas, and cold dark matter (CDM).

Results show that total gas remains stable (<3% variation), HI exhibits moderate but non-monotonic sensitivity (34% to +1%), while CDM responds in a highly symmetric and predictable way, scaling nearly linearly with Ω_m ($\pm$80%). Beyond these results, this project also proposes a new framework for analysing simulations: combining residual map slicing, integrated statistics, and block-based sensitivity mapping. This method provides a reproducible and quantitative way of disentangling cosmological and astrophysical effects, offering a foundation for future studies and observational comparisons.

Acknowledgments

I would like to thank my supervisor, Professor Rita Tojeiro, for her invaluable guidance and encouragement throughout this project. I am also grateful to Finley and Celina for their support, and to Lord Laidlaw for making this research possible through the Laidlaw Scholars program. Finally, I hope this work demonstrates a new way of looking at the cosmic web—one that disentangles cosmological and feedback processes in a reproducible and quantitative manner, and that others can build upon in future research.

Contents

Chapter 1

Introduction

1.1 Cosmological Scales and Structure

When we wake up each morning, our senses are quickly overwhelmed by the world around us. A simple cup of coffee, walking the dog, or chatting with a neighbour all fill our lives without us ever leaving the street. Yet if we step back, the scale of our world expands dramatically. A journey across the globe takes hours by plane, and even then, most of Earth will remain unseen to us. The planet suddenly feels vast, almost unreachably so. That was how I used to think about the world until I discovered cosmology and astrophysics. Our "world" is not just our towns or even our planet. Earth is one of many bodies orbiting the Sun, a star with planets and moons still only partly explored. The Sun itself, which feels enormous to us, is modest compared to Betelgeuse, a red supergiant nearly 700 times its size [1]. All of these belong to the Milky Way galaxy, home to hundreds of billions of stars [2].

If we zoom out again, galaxies themselves are not isolated. They gather in groups and clusters, gravitationally bound collections often containing hundreds or thousands of galaxies. Beyond clusters, galaxies are arranged in vast filaments, thread-like structures spanning hundreds of millions of light years. Between these filaments lie voids, enormous, relatively empty regions where few galaxies exist. Together, filaments, voids and clusters form the cosmic web, the largest-scale structure of the universe, revealed through both observation and simulation [3]. This widening sense of scale, from the mug in our hand to the cosmic web of galaxies, reminds us of both our smallness and our curiosity. The universe is not random chaos, but a structure shaped by fundamental parameters: how much matter it contains, how strongly it clumps, and how galaxies form within it. It is this framework, and how it changes when we vary those parameters, that forms the basis of my exploration in this essay. From this widening perspective of scale, the natural question arises: how can we study and test the processes that give rise to such vast cosmic

structures?

1.2 Simulating the Cosmic Web

To understand how the universe evolved into the intricate structure we observe today, scientists rely not only on telescopes and observations, but also on **cosmological simulations**. These are virtual laboratories that allow researchers to test hypothetical universes by adjusting fundamental parameters such as dark matter density, dark energy, gravity, and baryonic physics. Simulations help answer otherwise untestable questions. For example: *What would the universe look like today if there was less matter to start with? What if star formation feedback were stronger?* By comparing these simulated universes to real data, we can constrain the physical laws that govern our own. The image below shows one such simulation slice, revealing the intricate web of matter that emerges when gravity acts over billions of years. This structure is the backbone of cosmic evolution.

1.3 The Cosmic Web

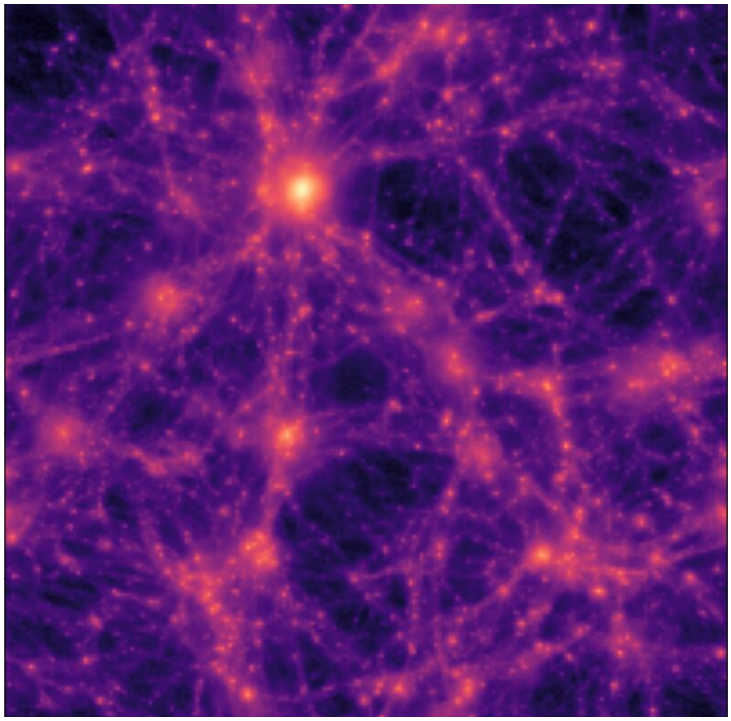

Figure 1.1: Slice from a cosmological simulation showing the large-scale matter distribution. Bright regions are high-density nodes and filaments; darker regions are voids.

As shown in Fig. 1.1, gravity forms a filamentary web. One of the most profound insights from simulations and observations is the emergence of the cosmic web, a vast network of filaments, nodes (clusters), and voids that form the backbone of the large-scale structure of the universe. Galaxies form preferentially along these filaments and nodes,

while vast voids remain nearly empty. This structure was first explained in detail by Bond et al. [3], who showed how anisotropic gravitational collapse gives rise to web-like patterns of matter in the expanding universe. The cosmic web is shaped by gravity acting on initial fluctuations in the early universe and reflects the interplay of dark matter, baryons, and feedback mechanisms. Understanding the cosmic web is crucial because it informs us about:

- Where galaxies form
- How matter flows across cosmic scales
- How the universe evolved from near-uniformity to its current complex state

1.4 The CAMELS Project

To explore how cosmic structures emerge and evolve, we use the **Cosmology and Astrophysics with MachinE Learning Simulations** (CAMELS) suite. CAMELS is designed to systematically vary both cosmological parameters (e.g., matter density Ω_m, fluctuation amplitude σ_8) and astrophysical parameters (e.g., star formation efficiency, AGN feedback strength) across hundreds of simulations.
Our focus is on the IllustrisTNG 1P hydrodynamical subset of CAMELS, which includes gas, stars, and black holes, unlike the N-body-only CAMELS simulations that model dark matter only. The 1P simulations allow independent variation of both cosmological and astrophysical parameters, making them especially suitable for disentangling their respective effects on the cosmic web [4]. The one-parameter variation in the 1P suite suggests that future exploration could extend this analysis using alternative approaches, the most promising being an extrapolation of the algorithm described in Chapter 2. Such an approach would allow adjustments that capture the coupled dynamics of cosmological and astrophysical variations simultaneously.

1.5 Focus: Gas, Dark Matter, and Filaments

We specifically study how gas, dark matter, and galaxy filaments overlap and align, and how these change when we vary physical parameters. Each component traces the cosmic web differently:

- Dark matter is collision-less and forms the gravitational backbone.
- Gas follows gravity but also experiences pressure, heating, cooling, and feedback.

- Galaxies trace dense regions but are redistributed by stellar and AGN feedback.

By comparing the filamentary structure in each component, we aim to answer:

- Which parameters have the strongest impact on filament formation?
- Do certain feedback models erase small-scale filaments?
- What does the alignment (or misalignment) between gas and dark matter tell us about feedback processes?

(See also: [5], [6] for methods on tracing filaments in different components)

1.6 Context and Motivation

This work seeks to determine which physical parameters most significantly shape the form of the cosmic web. By exploring this with simulations like CAMELS, we aim to gain insight into the hidden rules that govern the universe's large-scale evolution. The details of the methodology will be discussed in chapter 2.

Chapter 2

Methodology

2.1 Parameters and Setup

This work varies two cosmological parameters and four astrophysical feedback parameters. These define both the large-scale structure of the universe and the smaller-scale physical processes that shape how gas and galaxies evolve [4].

2.1.1 Cosmological and Astrophysical Parameters

- Cosmological parameters describe the initial conditions and global contents of the universe. They determine the overall "scaffolding" of the cosmic web: where filaments form, how matter clusters, and the rate of structure growth.
- Astrophysical (feedback) parameters describe baryonic processes inside galaxies, such as star formation and black hole feedback. These regulate how energy and momentum are injected back into the surrounding medium, influencing gas content and small-scale web features.

In this study, cosmology sets the skeleton of the cosmic web, while feedback parameters regulate how brightly and densely that skeleton is filled with gas and galaxies.

Table 2.1: *Description of parameters used in this study.*

Parameter	**Type**	**Meaning**	**Range**	**Fiducial**
Ω_m	Cosmological	Total matter density fraction; controls overall structure formation rate and clustering amplitude.	0.1–0.5	0.3
σ_8	Cosmological	RMS amplitude of matter fluctuations at $8h^{-1}$ Mpc; sets the density contrast on intermediate scales.	0.6–1.0	0.8
ASN1	Astrophysical	Energy per unit star formation rate injected into galactic winds (supernova feedback strength).	Varied	3.6
AGN1	Astrophysical	Energy per unit black hole accretion rate injected into feedback-driven outflows (AGN heating efficiency).	Varied	1.0
ASN2	Astrophysical	Characteristic wind launch velocity in stellar feedback models.	Varied	7.4
AGN2	Astrophysical	AGN kinetic-mode ejection velocity; controls feedback-driven gas expulsion.	Varied	20.0

Parameter Summary.

The fiducial parameter set defines the assumed 'true' universe, used as a baseline for all variations:

$$[\Omega_m, \sigma_8, \text{ASN1}, \text{AGN1}, \text{ASN2}, \text{AGN2}] = [0.3, 0.8, 3.6, 1.0, 7.4, 20.0].$$

Having defined the cosmological and astrophysical parameters governing the simulation suite, following sections describe the computational framework and techniques.

2.2 Data Slicing Framework

The 1P simulation set explores 28 parameters, each varied independently across five values. This design yields 140 unique 3D simulations. To analyse them efficiently, each 3D simulation box is subdivided into 15 slices, giving a total of $140 \times 15 = 2100$ two-dimensional maps. Each 2D map is stored on a 256×256 pixel grid.

2.2.1 Rationale for Slicing?

Handling full 3D simulation boxes can be computationally expensive. By slicing each cube into 15 2D cross-sections, we:

- Reduce storage and computational costs.
- Make it easier to visualise structures.
- Enable like-for-like comparisons between simulations.

A simple analogy is slicing a loaf of bread: each 2D slice represents one thin cut through the 3D simulation box. Looking at the slices side by side lets us see internal structure without handling the whole loaf at once.

2.2.2 Data Organisation

- Slices are stored in consecutive blocks of 15, grouped by their parent simulation.
- This simplifies indexing in code and ensures slices can be quickly retrieved for analysis.
- Visual comparisons between simulations are straightforward, since the same slice index corresponds to the same relative depth across all simulations.

2.2.3 Illustrative Example

In Fig. 2.1 the matter density parameter is varied, and the rest are fixed. Group refers to the slice the information is taken from so group 0 represents the first 15 maps, then group 2 represents the 3rd set of 15 maps. Notice the overall shape of the web is retained, but there are slight differences which we will explore in section 3 in a detailed analysis.

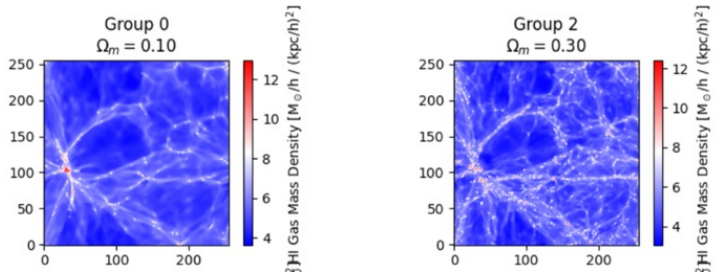

Figure 2.1: Schematic illustration of the slicing procedure. Each 3D simulation cube is subdivided into 15 equally spaced 2D slices (each on a 256×256 grid), providing cross-sectional views of the simulated cosmic web. Grouping slices into blocks of 15 facilitates indexing and cross-simulation comparison.

2.3 Test Case Design

To explore the IllustrisTNG 1P set in detail, the abundance of HI (neutral hydrogen), total gas and Cold Dark Matter was compared across simulations with varying parameters.

- Only simulations where one parameter varies at a time (others fixed) are selected.
- This isolates the effect of each parameter individually.

To quantify differences, the **logarithmic ratio** is computed between test maps and the fiducial map:

$$\log_{10}\left(\frac{\text{map}_{\text{test}}}{\text{map}_{\text{fid}}}\right) = \log_{10}(\text{map}_{\text{test}}) - \log_{10}(\text{map}_{\text{fid}})\,.$$

- **Positive values** $\rightarrow$ test map is denser than fiducial.
- **Negative values** $\rightarrow$ test map is less dense.
- **Zero** $\rightarrow$ no difference.

A value of zero does not imply the maps are identical; rounding and the sensitivity of the log scale can mask differences. We use a tolerance of 10^{-8} for the computations.

Test Case Selection

For each parameter there are **four test cases**:

- Two below the fiducial value.
- Two above the fiducial value.

This provides confidence in detecting consistent trends.

Alongside the direct map comparisons, the absolute log ratio is plotted. This instead tells us about sensitivity rather than change. Thus, implications are made based on the largest changes in areas of the plot. This is done by creating a sorting algorithm, which sorts all the pixels (from the 256x256 grid) into a readable array, then selecting 16x16 sub-grids from start to finish, and finally rearranges the collective value of the 16x16 grids in order from highest to lowest. The code then takes the top 3 sensitive areas and plots a green box around the sub-grid. This can be seen below in Fig. 2.2 Note that this algorithm can be applied to any slice. This is important because it is reproducible, so anyone can use this algorithm for IllustrisTNG.

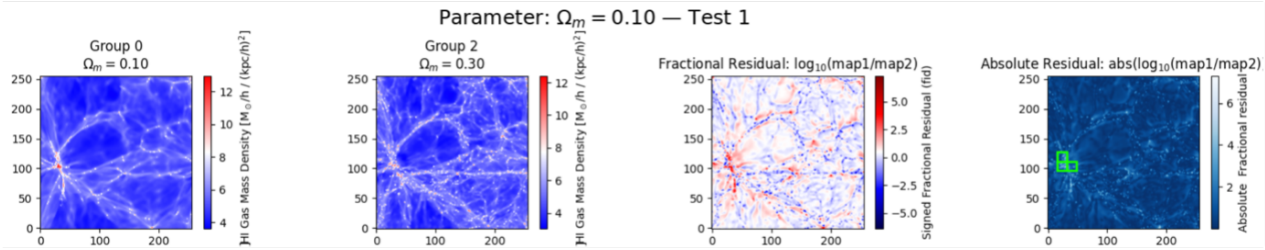

Figure 2.2: Example comparison of fiducial and test cases, this example uses the same parameter case from Figure 2.1, varying matter density and measure HI density maps. The first two panels are as above, and the third panel represents the signed change; the fourth panel represents the absolute log ratio with the green squares showing most sensitive areas. Note that the term 'fid' represents fiducial as explained in section 2.1.

This will be crucial in identifying regions having the most sensitivity to different processes.

2.4 Quantification Framework

In section 3 analysis is performed across neutral hydrogen (HI) and total gas content. These are density maps, and are on a log scale so the procedure is as detailed:

To obtain the total mass content in each map, the projected logarithmic density values were first converted back to linear units by computing $10^{m_{ij}}$. The total mass is then obtained by summing over all pixels:

$$M_{\text{tot}} = \sum_{i,j} 10^{m_{ij}}, \quad \text{where } m_{ij} \text{ represents the logarithmic density value in pixel } (i, j).$$

The framework converts logarithmic map data into quantitative measures of total gas and hydrogen content for statistical comparison.

However, it is important to note that neutral hydrogen exists in different phases. If the parameters are adjusted such that the HI becomes ionised, then the HI density in that

region effectively disappears from the map, since ionised gas is not recorded as neutral hydrogen in this dataset. This is why the total gas file is important for looking at the full picture.

The following chapter builds on this framework to examine how each parameter affects the structure and mass distribution of HI and total gas.

Chapter 3

Analysis of Findings

Although a full sweep across all parameter variations produces 288 residual maps, here we show representative slices for all the parameters, where the impact is most pronounced. For the remaining parameters, we report only the integrated statistics, with all maps available.

3.1 Matter Density

Table 3.1: **Quantitative results for matter density effects.** All values are given as percentage differences relative to the fiducial case $\Omega_m = 0.3$. While total gas remains nearly unchanged ($< 3\%$), HI shows moderate variation, and CDM exhibits large shifts up to $\pm 80\%$.
Note: Red cells indicate decreases, green cells indicate increases relative to fiducial values.

Ω_m	$\Delta_{\mathbf{HI}}$ (%)	$\Delta_{\mathbf{GAS}}$ (%)	$\Delta_{\mathbf{MCDM}}$ (%)
0.1	1.31	-2.26	-80.26
0.2	-16.70	-0.07	-40.13
0.4	-33.71	0.59	40.35
0.5	-20.08	0.04	80.65

3.1.1 Neutral Hydrogen

The variation in neutral hydrogen is most difficult to interpret. Across different values of Ω_m, HI does not follow a simple monotonic trend. Instead, it shows moderate decreases (-16% to -34%) relative to the fiducial case, with one slight positive offset ($+1.3\%$ at $\Omega_m = 0.1$). This suggests that the HI distribution is sensitive to parameter changes in a way that is not straightforwardly symmetric, likely reflecting the competing processes of ionisation and recombination that are not directly encoded in the density maps.

3.1.2 Gas

In contrast, the total gas content remains remarkably stable across all simulations, with changes consistently below 1%. While the variations are small, there is a subtle asymmetry: reductions in Ω_m, tend to produce slight decreases, whereas increases produce slight increases. This indicates that total gas content is more robust to matter density changes than HI, though still responsive at a low level.

Background for M_{CDM}

The CAMELS simulations are constructed under a flat ΛCDM cosmology, consistent with the *Planck 2018* results [7], from which the fiducial parameter set of IllustrisTNG [8] was derived. In this framework, the sum of the density parameters satisfies unity ($\Omega_{\mathrm{total}} = 1$), with spatial curvature fixed to zero ($\Omega_k = 0$). This ensures that

$$\Omega_m + \Omega_\Lambda = 1,$$

and within the matter component,

$$\Omega_m = \Omega_b + \Omega_{\mathrm{cdm}}.$$

Consequently, Ω_{cdm} can be obtained as

$$\Omega_{\mathrm{cdm}} = \Omega_m - \Omega_b,$$

with CAMELS adopting $\Omega_b = 0.049$. The flat-universe assumption and fiducial cosmological parameters are explicitly documented in the CAMELS parameter files, where Ω_k is fixed to 0.0 [4]. For full details of the simulation design and model variations, see Villaescusa-Navarro, Hahn, Delgado, et al. [4].

3.1.3 CDM

The cold dark matter component exhibits the most dramatic variation. Compared to the fiducial $\Omega_m = 0.3$, decreases in Ω_m produce substantial losses in CDM density (-40% to -80%), while increases produce equally large gains ($+40\%$ to $+80\%$). Unlike HI and total gas, this behaviour is highly symmetric around the fiducial case and closely tracks the definition of $\Omega_{\mathrm{CDM}} = \Omega_m - \Omega_b$. This demonstrates that CDM is directly governed by the overall matter density parameter, making it the most sensitive and predictable tracer in the analysis.

Comparison with Theory

Fixing the baryon fraction at $\Omega_b = 0.049$, the fiducial CDM fraction is:

$$\Omega_{\text{cdm}}^{\text{fid}} = 0.3 - 0.049 = 0.251$$

The expected percentage change relative to the fiducial is then:

$$\Delta_{\text{CDM}}(\%) = \left[\frac{\Omega_m - \Omega_b}{\Omega_m^{\text{fid}} - \Omega_b} - 1\right] \times 100$$

Table 3.2: Comparison of simulated and predicted percentage changes in CDM relative to the fiducial case. The predicted values are computed from $\Omega_m - \Omega_b$ with $\Omega_b = 0.049$.

Ω_m	**Simulated $\Delta_{\text{CDM}}(\%)$**	**Predicted $\Delta_{\text{CDM}}(\%)$**	**Difference (%)**
0.1	-80.26	-79.70	-0.56
0.2	-40.13	-39.80	-0.33
0.4	40.35	39.80	+0.55
0.5	80.65	79.70	+0.95

The agreement between simulation and prediction is very close ($< 1\%$ difference), confirming that CDM scales almost linearly with Ω_m as expected.

The rows in Fig. 3.1 (below) represent HI, Gas, and CDM respectively, top to bottom. It is clear that although the total gas content remains nearly constant, the structure is still affected. Investigating the green squares on the last column, HI is mostly changed at the filamentary node, Gas is changed most dramatically in potential areas where formation occurs. CDM is most notably filamentary in extreme cases.

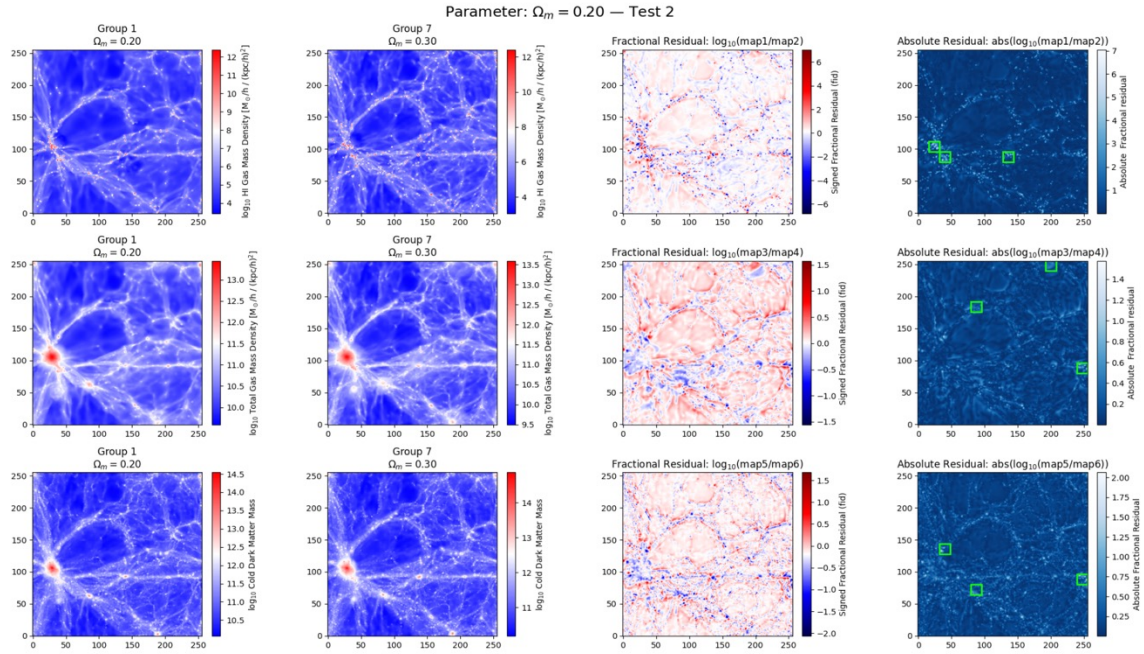

Figure 3.1: Example residual maps for $\Omega_m = 0.2$ compared to the fiducial case. Each row shows neutral hydrogen (HI), total gas, and cold dark matter (CDM), with residuals highlighted in fractional and absolute form. The plots illustrate the pronounced sensitivity of CDM to changes in Ω_m, in contrast to the relative stability of the gas content.

Overall, it seems that knowing more about the full 3D picture would help identify exactly where these types of structures changes are located. Where green boxes are located on the edge of the slice, one could postulate a flow out or into the system from a slice we can't see.

3.2 Matter Fluctuations

The parameter σ_8 quantifies the root-mean-square (RMS) amplitude of matter density fluctuations on scales of $8\,h^{-1}\,\mathrm{Mpc}$. It directly modulates the level of clustering, with higher σ_8 enhancing non-linear growth and low values suppressing structure formation. Table 3.3 summarises the resulting percentage deviations in matter density across the HI, gas, and CDM components, relative to the fiducial case $\sigma_8 = 0.8$.

Table 3.3: Quantitative results for matter density across neutral hydrogen (HI), total gas, and cold dark matter (CDM). All values are given as percentage differences relative to the fiducial case $\sigma_8 = 0.8$.

σ_8	Δ_{HI} (%)	Δ_{GAS} (%)	Δ_{MCDM} (%)
0.60	25.19	-2.22	-2.94
0.70	2.54	-0.97	-1.33
0.90	-21.13	1.98	1.49
1.00	-21.50	4.56	6.62

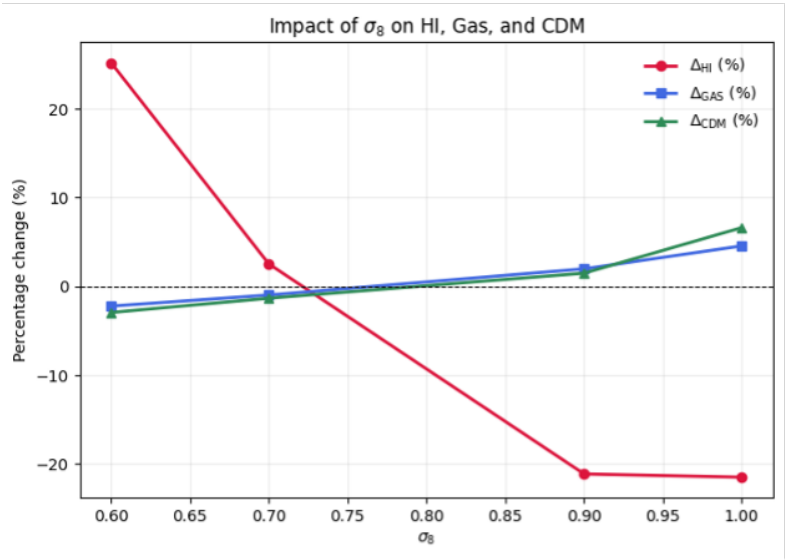

Figure 3.2: Line plots showing the quantitative results in Table 3.3. A clear trend emerges: HI exhibits a monotonic decrease with increasing σ_8, while gas and CDM display modest positive responses at higher fluctuation amplitudes.

The sensitivity of each component to changes in σ_8 reveals distinct physical responses. Neutral hydrogen (HI) decreases sharply with rising σ_8, reflecting enhanced ionisation and gas heating in regions of amplified density contrast. In contrast, CDM and total gas densities rise mildly, consistent with the gravitational amplification of collapsed structures in high-σ_8 cosmologies.

These results are qualitatively consistent with predictions from linear perturbation theory [9, 10], which describe how the amplitude of matter fluctuations drives early structure growth, and with subsequent non-linear refinements [11]. The chosen fiducial value of $\sigma_8 = 0.8$ also aligns with constraints from the *Planck* 2018 cosmological analysis [7], supporting the adopted parameter baseline.

Physically, the decline in HI mass fraction at large σ_8 values indicates a transition from diffuse to ionised states, as stronger clustering induces local heating and feedback. This behaviour reinforces the interpretation that σ_8 modulates not only the total matter distribution, but also the thermodynamic state of the baryonic component.

3.3 Astrophysical Feedback Effects

Astrophysical feedback processes, driven by stars and black holes, play a crucial role in regulating the gas content of the cosmic web. By varying stellar and AGN feedback parameters, we can test how injection and winds reshape the distributions of HI, total gas and dark matter.

3.3.1 Stellar Feedback (ASN1)

Table 3.4: Percentage changes in HI, total gas, and CDM for different stellar feedback strengths (ASN1). HI declines sharply with stronger feedback, reflecting the removal or heating of neutral gas, while gas and CDM remain comparatively stable.

ASN1	Δ_{HI} (%)	Δ_{GAS} (%)	Δ_{MCDM} (%)
0.90	-9.44	-3.62	0.15
1.80	-18.11	-1.38	0.04
7.20	-16.03	1.32	0.03
14.40	-13.89	1.97	0.03

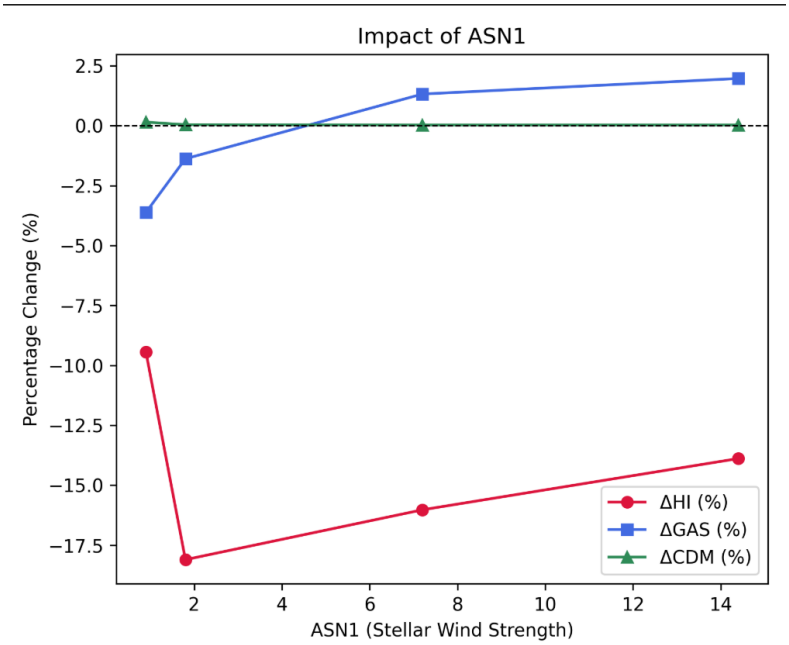

Figure 3.3: Impact of stellar feedback (ASN1) on HI, gas, and CDM. The results highlight HI's strong suppression under increasing stellar wind energy, with gas and CDM remaining comparatively stable (see Table 3.4).

3.3.2 AGN Accretion Feedback (AGN1)

Table 3.5: Percentage differences in HI, total gas, and CDM under varying AGN accretion feedback (AGN1). HI shows maximum suppression at intermediate accretion efficiencies, while gas and CDM exhibit negligible variation, indicating limited large-scale coupling.

AGN1	Δ_{HI} (%)	Δ_{GAS} (%)	Δ_{MCDM} (%)
0.25	-2.21	-0.21	0.00
0.50	-9.56	-0.15	0.08
2.00	-12.96	0.19	0.00
4.00	0.26	0.02	0.00

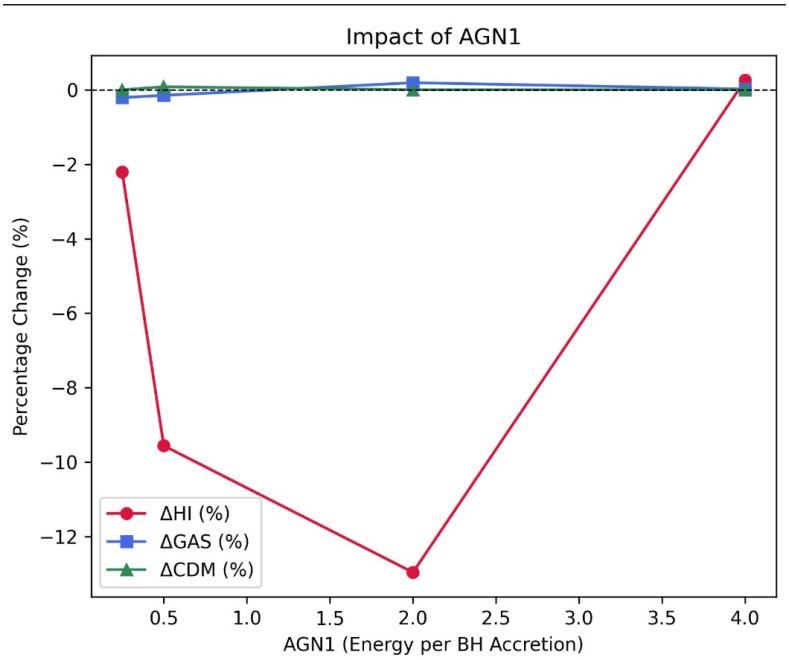

Figure 3.4: Impact of AGN accretion feedback (AGN1) on HI, gas, and CDM. HI shows maximum suppression at intermediate AGN energy injection, while gas and CDM are minimally affected (see Table 3.5).

3.3.3 Stellar Wind Velocity (ASN2)

Table 3.6: Percentage differences in HI, total gas, and CDM as a function of stellar wind launch velocity (ASN2). HI exhibits extreme positive and negative swings, revealing its strong sensitivity to wind dynamics, whereas gas and CDM remain largely unaffected.

ASN2	Δ_{HI} (%)	Δ_{GAS} (%)	Δ_{MCDM} (%)
3.70	62.22	-1.95	0.00
5.23	24.74	-1.09	-0.06
10.47	-50.13	0.98	0.05
14.80	-63.71	1.59	0.05

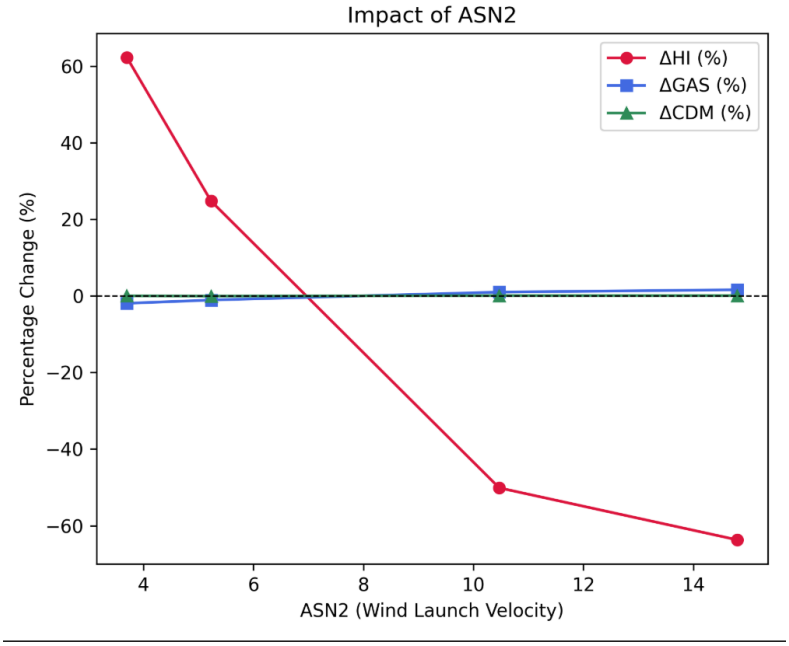

Figure 3.5: Impact of stellar wind velocity (ASN2) on HI, gas, and CDM. HI exhibits extreme swings in abundance, revealing strong sensitivity to wind-launch dynamics, while gas and CDM remain largely unchanged (see Table 3.6).

3.3.4 AGN Kinetic Feedback (AGN2)

Table 3.7: Percentage differences in HI, total gas, and CDM with increasing AGN kinetic-mode ejection velocity (AGN2). HI responses are moderate and non-monotonic, while gas and CDM show only small variations, consistent with weak coupling at large scales.

AGN2	Δ_{HI} (%)	Δ_{GAS} (%)	Δ_{MCDM} (%)
10.00	3.54	-0.65	0.01
14.14	-17.08	-0.33	0.03
28.28	-3.01	0.50	0.01
40.00	6.81	0.48	0.00

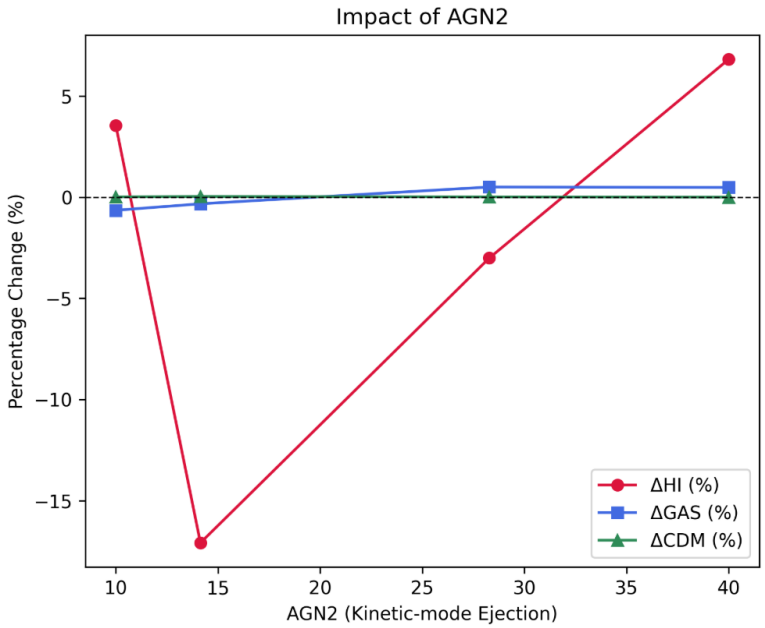

Figure 3.6: Impact of AGN kinetic feedback (AGN2) on HI, gas, and CDM. HI responds moderately to changes in kinetic-mode ejection, whereas gas and CDM show only weak variation (see Table 3.7).

3.4 Conclusions

The analysis across cosmological and astrophysical parameters highlights the contrasting ways in which different components of the cosmic web respond to physical variations.

For cosmological drivers, Ω_m produced the clearest and most predictable trend: CDM scaled nearly linearly with Ω_m ($\pm 80\%$), validating theoretical expectations from flat ΛCDM. In contrast, HI showed moderate but non-monotonic changes, reflecting the complexity of ionisation and recombination processes. Meanwhile, σ_8 revealed that HI is extremely sensitive to the amplitude of matter fluctuations, rising by $\sim 25\%$ at low σ_8 but falling by $\sim 21\%$ at high σ_8, while CDM shifted only modestly, demonstrating its robustness as a tracer of clustering.

The astrophysical feedback parameters introduced greater complexity. Stellar wind strength (ASN1) and AGN accretion feedback (AGN1) both suppressed HI to varying degrees but produced only weak effects on total gas or CDM. Stellar wind velocity (ASN2), however, generated the most dramatic response: HI shifted from large positive enhancements to deep suppression depending on wind speed, highlighting how small-scale baryonic processes strongly reshape the neutral fraction. Finally, AGN kinetic-mode feedback (AGN2) induced subtler changes, with HI fluctuating between modest gains and losses, again with negligible impact on total gas or CDM.

Taken together, these findings demonstrate a key distinction: cosmological parameters set the global scaffold of the cosmic web, governing CDM and overall structure, while astrophysical feedback redistributes baryons, leaving dark matter largely untouched but dramatically altering HI visibility. This interplay emphasises the dual necessity of cosmology and astrophysics. Only by disentangling both can we meaningfully interpret observational tracers of the web.

The patterns revealed here provide not only validation of theoretical expectations but also a foundation for future generalisation. Machine learning models could potentially extrapolate these relationships across the full 3D simulation suite and guide observational comparisons.

Chapter 4

Future Work and Outlook

This project presents a novel framework for disentangling cosmological and astrophysical drivers of the cosmic web, combining residual slicing, integrated statistics, and block-based sensitivity mapping. By systematically varying parameters in the CAMELS IllustrisTNG 1P suite, we established a clear hierarchy of sensitivities:

- *Cold Dark Matter (CDM)* responds in a symmetric, predictable way, scaling almost linearly with Ω_m ($\pm 80\%$). It therefore serves as the most robust and theoretically well-understood tracer of large-scale structure.
- *Neutral Hydrogen (HI)* is highly sensitive to both cosmology and feedback. Its response is non-monotonic under Ω_m, strongly varying under σ_8, and dramatically reshaped by feedback parameters, especially stellar wind velocity. This makes HI the most informative but also the most complex tracer.
- *Total Gas* remains remarkably stable ($< 3\%$ variation), acting as a near-conserved reservoir that reflects baryon redistribution rather than net gain or loss.

This stratified response is a new quantitative result: while the robustness of CDM is well known, the distinct divergence between HI and total gas and their complementary roles has not been captured in this much detail yet.

Future research could extend this work in several directions. First, the methodology can be expanded from *2D slicing* to *full 3D* statistical measures, such as filament-finding algorithms and power spectrum analysis, providing deeper insight into structural evolution. Second, applying this framework to *observational data*, particularly from upcoming HI surveys (e.g. SKA, MeerKAT), could test whether the predicted non-monotonic HI sensitivities are borne out. Finally, the reproducible and quantitative design of this framework makes it well-suited for integration with machine learning models, which

could generalise parameter sensitivities across the entire CAMELS suite and enable rapid inference of cosmological parameters from future survey data.

Overall, this work demonstrates a new, quantitative way of *separating the roles of cosmology and feedback* in shaping the cosmic web. By formalising this distinction into a reproducible framework, it provides both a methodological foundation for future simulation studies and a bridge to upcoming observational efforts, offering the potential to sharpen our interpretation of the largest structures in the universe.

Bibliography

[1] E. M. Levesque et al. “The effective temperature scale of galactic red supergiants: Cool, but not as cool as we thought”. In: *The Astrophysical Journal* 628.2 (2005), pp. 973–985. DOI: 10.1086/430901. URL: https://doi.org/10.1086/430901.

[2] T. C. Licquia and J. A. Newman. “Improved estimates of the Milky Way’s stellar mass and star formation rate from hierarchical Bayesian meta-analysis”. In: *The Astrophysical Journal* 806.1 (2015), p. 96. DOI: 10.1088/0004-637X/806/1/96. URL: https://doi.org/10.1088/0004-637X/806/1/96.

[3] J. R. Bond, L. Kofman, and D. Pogosyan. “How filaments of galaxies are woven into the cosmic web”. In: *Nature* 380.6575 (1996), pp. 603–606. DOI: 10.1038/380603a0. URL: https://doi.org/10.1038/380603a0.

[4] F. Villaescusa-Navarro, C. Hahn, A. Delgado, et al. “The CAMELS project: Cosmology and astrophysics with machine learning simulations”. In: *arXiv preprint* (2021). eprint: 2010.00619. URL: https://arxiv.org/abs/2010.00619.

[5] M. Cautun, R. van de Weygaert, B. J. T. Jones, and C. S. Frenk. “Evolution of the cosmic web”. In: *Monthly Notices of the Royal Astronomical Society* 441.4 (2014), pp. 2923–2973. DOI: 10.1093/mnras/stu768. URL: https://doi.org/10.1093/mnras/stu768.

[6] N. I. Libeskind et al. “Tracing the cosmic web”. In: *Monthly Notices of the Royal Astronomical Society* 473.1 (2018), pp. 1195–1217. DOI: 10.48550/arXiv.1705.03021. URL: https://doi.org/10.48550/arXiv.1705.03021.

[7] P. Collaboration, N. Aghanim, et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910.

[8] A. Pillepich et al. “Simulating galaxy formation with the IllustrisTNG model”. In: *Monthly Notices of the Royal Astronomical Society* 473.3 (2018), pp. 4077–4106. DOI: 10.1093/mnras/stx2656.

[9] K. A. Malik and D. R. Matravers. "A concise introduction to perturbation theory in cosmology". In: *Classical and Quantum Gravity* 25.19 (2008), p. 193001. DOI: 10.1088/0264-9381/25/19/193001. URL: https://doi.org/10.1088/0264-9381/25/19/193001.

[10] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro. "Large-scale structure of the Universe and cosmological perturbation theory". In: *Physics Reports* 367.1 (2002), pp. 1–248. ISSN: 0370-1573. DOI: https://doi.org/10.1016/S0370-1573(02)00135-7. URL: https://www.sciencedirect.com/science/article/pii/S0370157302001357.

[11] C. Uggla and J. Wainwright. "Cosmological perturbation theory revisited". In: *Classical and Quantum Gravity* 28.17 (2011), p. 175017. DOI: 10.1088/0264-9381/28/17/175017. URL: https://doi.org/10.1088/0264-9381/28/17/175017.

Data Acknowledgment

This research made use of the CAMELS IllustrisTNG 1P simulations. We gratefully acknowledge the CAMELS team for making the data publicly available. For data access see https://users.flatironinstitute.org/~camels/CMD/.