

Perturbative and non-perturbative theories of structure formation

DISSERTATION

zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von
Mandar Karandikar

aus
Ratlam – Indien

Bonn, März 2025

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter/Betreuer: Prof. Dr. Cristiano Porciani
Gutachterin: Prof. Dr. Jennifer Schober

Tag der Promotion: 10.07.2025
Erscheinungsjahr: 2025

Abstract

Understanding the formation of structure in the Universe is one of cosmology's primary goals. The material content of the Universe is dominated by an invisible, non-baryonic component which we call dark matter. The gravitational force drives the clustering of dark matter from initially-small density fluctuations to virialised objects. At early times, and on large scales, cosmological perturbation theory (PT) provides a good model for structure formation. As local overdensities approach unity, the assumptions underlying PT start to break down. The EFTofLSS is an extension to PT which fixes some of PT's conceptual issues and make accurate predictions in this regime. At late times, and/or on small scales, dark matter clumps collapse to form virialised structures called halos. In this non-linear regime, analytical models are approximations, and we require N -body simulations to accurately probe these scales.

This thesis is made up of two key projects. The first of these probes the quasi-linear regime and makes conceptual inquiries into the EFTofLSS. The framework assumes that the stress tensor of the dark matter fluid may be expanded in a Maclaurin series about the overdensity and velocity fields. The coefficients of this expansion may be interpreted in two ways. In the top-down sense, they are free parameters calculable by matching the dark matter power spectrum to observations, while in the bottom-up sense they are measurable in simulations. We specialise to a one-dimensional, dark-matter only universe, where these two approaches can be directly compared. We find that the top-down and bottom-up predictions agree with each other, providing evidence for the robustness of the EFTofLSS assumptions. Additionally, we perform controlled tests that shine a light on non-linear dynamics within the EFTofLSS.

The second project explores the formation of halos in the non-linear regime. The conventional analytical approach is to treat gravitational collapse with an idealised spherical model. Within the excursion-set theory framework, spherical collapse (and its ellipsoidal extension) maps the initial overdensity field of a proto-halo to its collapsed descendent. In this work, we provide an alternative to this traditional approach. Namely, we train a neural network to predict the accretion history of halos from the tidal properties of the proto-halo. We find that the neural network predicts the accretion history and derived properties of the halo population with greater accuracy than the analytical models.

As cosmologists, our modelling of structure formation is incomplete. Although we are familiar with the processes that govern it, many details and nuances are yet to be fully understood. In this thesis, we have attempted to answer some open questions in this field. Through this work, we have validated the current theoretical framework to model structure formation on large scales, and introduced a novel approach to understand the formation of non-linear structures.

List of Publications

The following is a list of first-author publications and manuscripts in preparation that are relevant to this thesis.

1. **Karandikar M.**, Porciani C., and Hahn O. (2024), *Testing the assumptions of the Effective Field Theory of Large-Scale Structure*, Journal of Cosmology and Astroparticle Physics, 2024, 051, <https://doi.org/10.1088/1475-7516/2024/01/051>

Author contributions: M.K. ran the simulations, performed the primary analysis, and led the writing of the paper. C.P. supervised the scientific exploration, and contributed in writing parts of the manuscript. O.H. provided the simulation code, which was adapted by M.K. for this use case. C.P. and O.H. helped in interpreting the results, and provided comments to the final manuscript.

2. **Karandikar M.** and Porciani C. (2025), *Predicting halo formation using a neural network*, in prep.

Author contributions: M.K. designed the neural networks and performed the primary analysis. C.P. provided feedback on the analysis, and helped interpret the results. M.K. drafted the manuscript, and C.P. provided their comments.

Acknowledgements

My first note of gratitude is towards Cristiano. Thank you for giving me the opportunity to do a PhD, and thank you for guiding me when I was most lost. Our scientific discussions were truly a pleasure. Thanks to Oliver Hahn for being a supportive mentor, I learned a lot from you. Emilio, thank you for all your scientific and non-scientific support, and that your door was always open. A special thanks to the members of my doctoral committee, Prof. Schober, Prof. Wang, and Prof. Löhner, for the time you took out of your busy schedules to read this thesis.

I could not have asked for a better group to do my research in. Special thanks to Matteo, Alex and Lucas for always being the adults in the JCs, and entertaining all of our complaints at lunch. My most sincere thanks to Anna, Benedetta, Dato, Elena, Prachi, Yvonne, and Bo for your friendship on the good days, the unwavering support on the bad ones, and for the ice-cream on both. This journey would have been far more difficult without you. I owe a large debt of gratitude to Yousry, whose optimism and positivity was always a welcome relief.

I would like to thank my friends, who have been an integral part of my life all these years. To Chandu, for all the tea and coffee in the world, for the endless stream of easy conversation, and for your reassurance at every step of doubt. To Aditya, for matching (and often, exceeding) all the silliness in the forests, for the football, and for your endless brightness. To Estera, for creating a community where I could belong, for the energy you bring to every moment, and for being you. To Kartik, for being an excellent flatmate, and more importantly, the best of friends for nearly a decade. To Aparna and Raphael for your love, hospitality, and friendship. And to Amith, for always having an open ear, and an open heart. Maggie, Rafa, Julia, Isabelle, Federica, Stefan, Carl, Lio, and you have all made life truly enjoyable. Thank you for the improv, the movie nights, and the murder mysteries.

There are no words that could adequately express the debt I owe to my family for their unending and unconditional love and support. Aai, Baba, Aaji and Aaji, I'm

sorry for always being so far away. Mayura, I would not be here if not for you. For sowing the seed of Astronomy in my mind, for tolerating my teenage self, and for being the most wonderful sister one could ask for, a heartfelt thanks. And finally to Mira, for everything.

Contents

1. Introduction	1
1.1. A concise history of the Universe	2
1.2. Cosmic expansion	6
1.3. Dark matter	8
1.4. The CMB, and the shape of the Universe	9
1.5. Notation and terminology	11
2. Theory	13
2.1. Cosmological perturbation theory	16
2.1.1. Standard (Eulerian) perturbation theory	18
2.2. Effective field theory of large-scale structure	23
2.2.1. Equations of the EFTofLSS	24
2.2.2. EFT corrections	26
2.3. Formation and assembly of dark matter halos	28
2.3.1. Gravitational collapse under spherical symmetry	29
2.3.2. Excursion-set theory	32
2.3.3. Halo mass function	37
2.3.4. Predictions for individual halos	39
3. Numerical Methods	41
3.1. N -body simulations	41
3.1.1. Historical note	44
3.1.2. Time integration	46
3.1.3. Poisson solver	49
3.1.4. Mass assignment	53
3.1.5. On finding structure in simulations	54
3.1.6. Convergence and other numerical issues	59

3.2. Simulations used in this work	60
3.2.1. Simulations in one-dimension	60
3.2.2. GADGET-3, and simulations in three-dimensions	61
3.3. Artificial neural networks	62
3.3.1. Backpropagation	64
3.3.2. Learning tasks	65
3.3.3. Activation functions	66
3.3.4. Training a neural network	67
4. Testing the assumptions of the EFTofLSS	69
4.1. Simulation setup	69
4.2. The non-linear scale, and effective stress	71
4.3. EFTofLSS estimators	73
4.3.1. Estimates from spatial cross-correlations	74
4.3.2. Estimates from fitting	75
4.4. Effective speed of sound	75
4.5. The EFTofLSS power spectrum	78
5. Predicting halo formation using a neural network	83
5.1. Network properties	84
5.1.1. Input features	84
5.1.2. Architecture	84
5.1.3. Convergence	85
5.2. Mass function and halo masses	86
5.3. Accretion history and formation time	88
5.4. Effect of tidal anisotropy	92
5.5. Extension to higher masses	95
6. Conclusions	99
A. Statistical description of particles	105
A.1. The Euler-Lagrange equations	105
A.2. Hamilton's equations	106
A.3. Poisson brackets	107
A.4. Liouville's theorem	107
A.5. BBGKY hierarchy and the Vlasov equation	108
B. Predicted accretion histories by population	111
C. Comparison plots from Bo17	113
D. Testing the assumptions of the EFTofLSS	115

Bibliography	151
List of Figures	171
List of Tables	177

CHAPTER 1

Introduction

Cosmology refers to the study of the origin and evolution of the observable Universe. In the sense of a modern science, cosmology is young. Our current concordance theory of gravity – the only fundamental force that operates on cosmological scales – was first proposed a little more than a century ago¹. Evidence for the distribution of the matter-energy budget of the Universe is only a few decades old, and the nature of the two components that dominate this budget is still unclear (see Sec. 1.3, 1.4). Further, pressing issues regarding the evolution of the Universe – and the structure within it – are emerging with the ever-changing observational landscape (see e.g. Freedman, 2021; Abdalla et al., 2022, for recent reviews). Nonetheless, we have made tremendous progress – we have a comprehensive (but perhaps non-exhaustive) list of the physical processes involved, and an estimate of the order in which they played a role in shaping the formation of structure. However, the interplay of these processes is not completely understood. At later times, and on smaller (sub-megaparsec) scales, non-linearity becomes important, and one must supplement the analytical techniques that guide us in establishing a picture of the early Universe. This is achieved using numerical simulations at different scales. Both analytical and numerical methods must explain the Universe as we *observe* it, and it is through the cooperation of the three branches that we can build a consistent story of the evolution of our Universe. The wealth of constantly improving observational data, and the richness and power of new numerical techniques makes

¹By the middle of 1913, Einstein, in collaboration with Grossman had put together nearly all the pieces that make up modern general relativity. However, the equations they published (Einstein & Grossmann, 1913) were not generally covariant. In 1915, Einstein (often with Hilbert) published four papers, the last of which introduced the now-familiar covariant field equations (Eq. 2.1; Einstein, 1915).

this an exciting moment in modern cosmological studies. It is the objective of this thesis to leverage some of these techniques to shine a small flashlight into the dark gaps in our theoretical models of the Universe. The remainder of this section is a primer on our current understanding of the Universe. We sketch its history in brief, summarise its key properties, and build towards a more detailed description of structure formation that follows in Chapter 2.

1.1. A concise history of the Universe

Our knowledge of the very early Universe is highly speculative, as we have no observational data from this time. The best we can do is to extrapolate our understanding of physics to these early times, and look for imprints of these processes on the observations we do possess. The proposed model begins with a *hot big bang* – a dense primordial soup of standard model particles at a pressure and temperature so high as to strip atomic nuclei of their constituents² (e.g. Lemaître, 1931b). At these high temperatures and energy densities, classical physics fails, and to truly understand this epoch, a quantum theory of gravity must be formulated. The first event that cosmologists *can* claim knowledge of, is a short-lived but critical era of inflation, where the primordial soup expanded exponentially, cooling as it grew (e.g. Guth, 1981; Starobinsky, 1982; Albrecht & Steinhardt, 1982; Linde, 1982). In the conditions of the early Universe, the energy density is dominated by relativistic photons and neutrinos (this is called the *radiation-dominated* epoch). All particle species and radiation components can be safely assumed to be in thermal equilibrium with each other³. Thermal equilibrium may be assumed when particles interact at a rate much larger than the expansion rate of the Universe. This allows us to approximate the time at which different species decoupled from equilibrium. Neutrinos⁴ interact via the weak nuclear force, and they are the first species to decouple at a temperature⁵ around a few MeV k_B^{-1} (see e.g. Yanagisawa, 2014).

²This is called the quark-gluon plasma – quarks which go on to construct hadrons (e.g. protons and neutrons), and gluons which mediate the strong nuclear force. The existence of this state of matter was widely discussed in the late 1970s (e.g. Freedman & McLerran, 1977; Chin, 1978; Kapusta, 1979), and it was experimentally verified at CERN at the end of the millennium (Heinz & Jacob, 2000).

³After inflation, there is expected to be a symmetry between matter and antimatter, as opposed to the drastic asymmetry we observe today. This implies a mechanism to create this imbalance, known as *baryogenesis*. Many plausible theories of baryogenesis have been proposed, but no observational signatures have been detected. See, e.g. Riotto & Trodden (1999) for a review.

⁴Within the standard model of particle physics, neutrinos are massless. However, there is compelling evidence suggesting that neutrinos oscillate in three flavours, implying that they have mass. Massive neutrinos are particles of great cosmological interest, from hypothetically making up a small fraction of dark matter, to possibly damping the clustering of matter on small scales. We refer the curious reader to Lesgourgues & Pastor (2006) for a review on the subject.

⁵With one notable exception (see Sec. 1.4), we report the temperatures in this chapter in relation to the Boltzmann constant k_B which translates energy to temperature. Note that in much of

Let us now consider the baryonic components at this time. Neutrons and protons maintain equilibrium by decaying into each other. Once the weak interaction decouples from equilibrium, the neutron to proton ratio *freezes-out*. This occurs when the Universe is about a second old, and has a temperature of $\approx 0.85 \text{ MeV } k_B^{-1}$. Free neutrons have a short half-life, and without another process that preserves them, all neutrons would decay, leading to a universe populated entirely by hydrogen. Although neutrons can combine with protons to create Deuterium, energetic photons destroy any Deuterium that is produced at higher temperatures. This bottleneck is only crossed at $\approx 0.1 \text{ MeV } k_B^{-1}$, at which point, most of the Deuterium that survives is converted into ^4He . Trace amounts of ^3He , ^7Li , and ^7Be are also produced. This process, known as primordial or big-bang nucleosynthesis (BBN) is one of the first physical processes to leave observable imprints, viz. in the present-day abundances of these light elements in metal-poor environments⁶ (e.g. Walker et al., 1991; Fields & Olive, 2006; Cyburt et al., 2016). Thus, at the end of BBN, the Universe consists of decoupled neutrinos, relativistic photons, trace amounts of light nuclei, electrons, and decoupled dark matter particles⁷.

As the energy density of radiation decays faster than that of matter (see Eq. 2.13), the Universe eventually transitions from being radiation-dominated to matter-dominated. The turnover epoch is called (matter-radiation) equality, and occurs when the temperature is $\approx 0.8 \text{ eV } k_B^{-1}$. It is important at this point to consider the relationship between matter and radiation. Photons interact with free electrons (i.e., electrons not bound to a nucleus) via Thomson scattering. Further, photons with enough energy can ionise neutral atoms (photo-ionisation), and free electrons can combine with ions to create neutral atoms once again (radiative recombination). The next epoch in the evolutionary history of the Universe is the interplay of these three processes. The scattering cross-sections for photo-ionisation and radiative recombination are much smaller than that for Thomson scattering. As the temperature falls, fewer photons are energetic enough to cause ionisation, thus lowering the number of free electrons. In this way, the Universe makes a transition from fully-ionised plasma to neutral gas consisting primarily of hydrogen. This is characterised by the *recombination temperature*, estimated to be $\approx 0.32 \text{ eV } k_B^{-1}$. Without free electrons to scatter off of, the photons can then stream freely, rendering the Universe transparent for the first time. This photon-decoupling is estimated to occur at a temperature of $\approx 0.26 \text{ eV } k_B^{-1}$. Photon-decoupling and recombination were

the literature on early-Universe cosmology, temperature is expressed in natural units, where $k_B = 1$, and the unit of temperature is eV.

⁶The other channel through which these elements may be synthesised, is nuclear fusion within stars. Subsequently, stellar winds and supernovae enrich the stellar neighbourhood with these (and heavier) elements. Thus, the evidence for BBN boils down to measuring the abundances in environments that have not been enriched by stellar feedback.

⁷Dark matter is assumed to have decoupled by the end of BBN. Exactly when and how this occurs depends on the nature of the candidate particle.

not instantaneous processes. They were deeply intertwined with each other, and occurred together over a longer period of time (see e.g. Di Bari, 2018 for a review).

Decoupling has a critical observational consequence – the free streaming photons are visible today (albeit redshifted to a much lower temperature), giving us a glimpse at this so-called *layer of last scattering*. They constitute the cosmic microwave background (CMB), and its serendipitous observation⁸ by Penzias & Wilson (1965) was sufficient evidence for favouring the big bang model over the then-popular steady-state model⁹. The CMB has now been widely studied by three generations of probes, viz. COBE (e.g. Mather et al., 1990; Boggess et al., 1992), WMAP (e.g. Spergel et al., 2003, 2007), and *Planck* (e.g. Ade et al., 2014; Aghanim et al., 2020), and is one of the pillars of modern cosmology (see Sec. 1.4).

When the CMB photons start free-streaming, the Universe is about 375,000 years old, corresponding to a redshift (z) of 1090 (Di Bari, 2018). In the cold dark matter (see Sec. 1.3) paradigm, matter is expected to clump hierarchically, with the first objects to virialise being low-mass dark matter halos at $z \approx 100$, which merge to create bigger objects at later times (e.g. Blumenthal et al., 1984; Springel, 2005). These halos may range from earth-mass ($10^{-6} M_{\odot}$; Diemand et al., 2005) to as low as $10^{-13} M_{\odot}$, depending on the nature of the dark matter particle. There are only two sources of observable photons between the epoch of decoupling and the formation of the first stars. The first is the CMB, and the second are photons emitted by neutral hydrogen by way of the spin transition which emits at a wavelength of 21 cm. Since the former are nearly unaltered since recombination, the latter are our only source of information about these *dark ages* (see e.g. Pritchard & Loeb, 2008). The first generation of stars (termed population III) formed in low-metallicity environments within these low-mass halos at $z \approx 20$ –30 (Stark, 2016). These stars enriched the environments in which they lived by radiative feedback, and via supernovae (see Klessen & Glover, 2023 for a recent review). The first dwarf galaxies are expected to have formed within those low-mass halos which contain gas left behind by the population III stars. Galaxies are then assumed to grow by rapidly accreting cold gas (see Stark, 2016 for a review).

There is another transition within the hydrogen atom which is cosmologically relevant. When the sole electron in neutral hydrogen interacts with an energetic photon, it can absorb the photon to make the transition from its ground state to the first excited state. This Lyman-alpha ($\text{Ly-}\alpha$) line appears as an absorption feature in the spectra of background objects. Quasars are highly-luminous objects that reside

⁸It is worth mentioning here that a group led by Robert Dicke was also working on detecting the CMB at this time (Dicke et al., 1965).

⁹In steady-state cosmology, the expansion of the universe is matched by creation of new matter which keeps the average density constant. Championed, for example, by Hoyle, Bondi, and Gold (see in particular Hoyle, 1948), the model was essentially ruled out by the observation of the CMB, although its proponents continued to argue in favour of a modified version (see e.g. Hoyle et al., 1995). Incidentally, it was Fred Hoyle who coined the term ‘big bang’ in 1949, although perhaps not in a pejorative manner as is often reported (Kragh, 2013).

at the centers of galaxies, and can therefore be observed at very high redshifts. If we measure the spectra of distant quasars, the presence of the Ly- α absorption line is an indicator of neutral hydrogen in the intervening medium (Gunn & Peterson, 1965). This offers a clear test for the presence and abundance of ionised material. Remarkably, in spectra of quasars at $z \leq 6$, these absorption lines are absent, while they are present in higher-redshift quasar spectra. This implies the transition of the Universe from neutral to fully-ionised, likely due to energetic radiation emitted by population III stars. This epoch of reionisation (EoR) is an important piece in the evolutionary history of our Universe. For a review of the connection between early galaxies and the EoR in the context of the latest and upcoming observations, see Robertson (2022).

Evolution between $z \approx 10$ and today is moulded, in part, by astrophysical processes. The population III stars left behind gas in the galaxies, which – when cooled – is a fertile breeding ground for a new generation of stars. The efficiency of star formation depends on the bulk properties of this medium – the temperature, pressure, density, and metallicity. Observational estimates indicate that the global star formation rate density increased to a peak at $z \approx 2$ (termed *cosmic noon*), and has declined rapidly since then. The main reason for this decline is the lowered availability of cold gas in galaxies (see Madau & Dickinson, 2014 for a review). Galaxies themselves grow via mergers, and by accreting gas from their surrounding medium. Most massive galaxies contain a supermassive black hole (SMBH) at their centres, and feedback from the SMBH regulates their growth. A combination of these complex dynamical processes has led to a large diversity of galaxies, ranging widely in age, mass, shape, and miscellaneous physio-chemical properties. Although an exploration of these processes is interesting in its own right, it is beyond our scope.

Instead, our focus is on the non-baryonic structure of the Universe. At the epoch of photon decoupling, the Universe was relatively smooth, with small fluctuations in the matter density. Today, we observe a rich, multiscale clustering with a diverse zoo of structures. Galaxy-redshift surveys have revealed that the spatial distribution of galaxies follows a web-like pattern (e.g. De Lapparent et al., 1986; Huchra et al., 2012). The underlying dark matter distribution is thus thought to be similarly organised (e.g. Bond et al., 1996). The components of this *cosmic web* are typically classified based on the number of dimensions along which they have undergone collapse. The dominant component is an elongated structure collapsed along two dimensions called a *filament*. Matter is funnelled through filaments from low-density environments to high-density *clusters* where different filaments meet. The driving force behind this evolution of structure is gravitational amplification of the initially-small fluctuations. Broadly, this process may be divided into three distinct but overlapping regimes, based on the size of the fluctuations. When the dark matter overdensity field is much smaller than unity, we are in the linear regime, where cosmological perturbation theory (Sec. 2.1) can make accurate analytical predictions. As more matter clumps together, the local overdensity field approaches

unity, and we enter the quasi-linear regime. In the non-linear regime, the local overdensity exceeds unity, and analytical predictions become largely untenable. This is the realm of numerical simulations, which can approximate non-linear gravitational collapse. Overall, our physical models of structure formation paint a reasonable, broad-strokes picture. However, as the depth and richness of observational data improves, it is imperative to model the quasi- and non-linear regimes precisely and accurately. In this thesis, we aim to analyse different aspects of gravitational evolution by constructing and scrutinising such models. In the rest of this chapter, we will introduce some important ingredients that make up our concordance cosmology, and then raise specific questions that outline our objectives.

1.2. Cosmic expansion

In an expanding universe, the velocity of a galaxy due to expansion is proportional to its distance from the observer

$$v = H_0 d, \quad (1.1)$$

where the proportionality constant is named after Hubble. The discovery that our Universe is expanding depended on two key observations – the recession velocity of the Andromeda Galaxy, and its distance to us. In 1912, Vesto Slipher made the first observations of galactic redshift by measuring the spectrum of the Andromeda Galaxy. From these observations, he inferred the radial velocity of the galaxy with respect to the observer to be about -300 km s^{-1} (Slipher, 1913). In 1925, Hubble used observations of Cepheid variables in Andromeda to estimate its distance from us to be¹⁰ 285 kpc (Hubble, 1925). Lemaître (1927) derived a dynamic solution to Einstein’s equations, along with Eq. (1.1), thus predicting an expanding universe. Further, Lemaître linked the above observations¹¹, to estimate $H_0 = 625 \text{ km s}^{-1} \text{ Mpc}^{-1}$. In the English-language literature, Eq. (1.1) appeared for the first time in Hubble (1929), where Hubble used his own distances and Slipher’s redshifts to estimate $H_0 = 500 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (see Fig. 1.1). Although both Lemaître and Hubble’s estimates are highly inaccurate compared to modern measurements (chiefly due to errors in calibrating the distance), they were qualitatively correct. Eq. (1.1) is jointly credited to both, and is called the Hubble-Lemaître law. Subsequent estimates of H_0

¹⁰Hubble was not the first to estimate this distance (see e.g. Öpik, 1922 and references therein).

While Hubble’s estimate was inaccurate (the modern estimate is 778 kpc), it answered whether spiral nebulae were intra- or extra-galactic. This issue was subject to a lively discussion in the community that culminated in ‘The Great Debate’ between Shapley and Curtis at the Smithsonian in 1920. Curtis argued that the nebulae were extragalactic, and was proven right as distances to them were finally measured to be too large to belong to the Milky Way.

¹¹The English translation of Lemaître’s paper appeared in 1931 without mention of Eq. (1.1). It is now known (Livio, 2011) that Lemaître himself made the omission, deciding not to repeat what Hubble had reported in 1929. An historical account of discovery of cosmic expansion may be found in Nussbaumer & Bieri (2009).




Figure 1.1. – The distance-velocity plot from Hubble (1929). The filled and empty circles represent – respectively – measurements from individual ‘nebulae’ (galaxies), and from combining them into groups. The solid and dashed lines are linear fits to the two sets of data, and the cross is the estimated mean of 22 nebulae for which individual distances were not found. The slope of the best-fit line is the estimated value of H_0 . *Credit:* PNAS; Hubble (1929).

have consistently shrunk (Trimble, 1996), and the current estimate from Cepheids and Type Ia supernovae¹² is $H_0 = 73.2 \pm 1.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Riess et al., 2021). The story of this parameter is, however, still unsettled. The measurement of H_0 from the CMB yields a lower estimate of $H_0 = 67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$. This disagreement between the early- and late-time measurements has been christened the ‘Hubble tension’ (see e.g. Verde et al., 2019 for a review, and Di Valentino et al., 2021 for possible resolutions).

The constant H_0 merely refers to the expansion rate of the Universe today. The complete expansion history is captured by the time-dependent Hubble function $H(t)$. This function is heavily dependent on the choice of cosmological model. Observational investigations compared the apparent magnitudes of high-redshift Type Ia supernovae to their low-redshift counterparts to measure $H(t)$, and found that the expansion of the Universe is accelerating (Riess et al., 1998; Perlmutter et al., 1999). The hypothesised cause for this acceleration is called ‘dark energy’. Its nature is unknown, and models range from Einstein’s cosmological constant, to scalar fields (see Copeland et al., 2006 for a review). The phenomenological impact is that the energy density of the present-day Universe is dominated by dark energy. For the rest of this thesis, we assume dark energy to be described by the cosmological constant

¹²Cepheid variable stars and Type Ia supernovae are what astronomers call *standard candles*. Their intrinsic brightness (absolute magnitude) is known, and thus their apparent magnitude is a measure of their distance from us.

Λ . Physically, it can be treated as a perfect fluid with negative pressure, and it is thought to be the energy of empty spacetime. With this definition, the energy density of dark energy is constant in time. Dynamic models of dark energy have also been proposed (e.g. Zhao et al., 2017), and one model in particular – early dark energy – has been suggested as a solution to the Hubble tension (e.g. Poulin et al., 2019; Kamionkowski & Riess, 2023).

1.3. Dark matter

The existence of invisible material in the Milky Way was discussed in the late 19th century, and was subject to debate in the early 20th century. For instance, Kapteyn (1922) and Oort (1932) had already used the term ‘dark matter’ to explain the motion of stars in the galaxy (see Bertone & Hooper, 2018 for a historical review). By the 1930s, the Coma galaxy cluster was known to have an unusually high velocity dispersion. Zwicky (1933, 1937) applied the virial theorem to estimate the mass of the cluster. Comparing this dynamical mass estimate with the observed luminosity of galaxies within the cluster, led Zwicky to a mass-to-light ratio of ≈ 500 – there seemed to be 500 times more mass in the cluster than the luminosity would indicate. Zwicky was using Hubble’s estimate of H_0 , and rescaling the mass-to-light ratio to the more accurate H_0 value leads to a value of ≈ 8.3 , which is much closer to the modern estimates. Smith (1936), using a similar approach to Zwicky, measured the mass of the Virgo cluster, and found the average mass of a galaxy two orders of magnitude larger than expected. In the following decades, many ideas were put forth to resolve this discrepancy, e.g. modifications to Newtonian gravity (Finzi & Pirani, 1963), or suggestions that clusters were unstable, making the assumption of virialisation incorrect (Karachentsev, 1966). Around the same time, galaxy rotation curves – plots of the orbital velocity of visible galactic material against distance from the galactic center – were being increasingly discussed. In the 1970s, it became clear that the rotation curves of spiral galaxies were flatter, and extended to larger radii than the prediction from Newtonian gravity (Freeman, 1970; Rubin & Ford Jr, 1970).

In 1974, two landmark papers (Einasto et al., 1974; Ostriker et al., 1974) brought the two discrepant observations together to point out that both problems may be resolved by hidden, invisible matter. The existence of this non-baryonic, dark matter is now largely accepted. Historically, many particles (within and beyond the standard model of particle physics) have been suggested as candidates – often characterised by their mass, and particle dark matter still remains the leading hypothesis – see Bertone et al. (2005) for a review. Although detection efforts over the last decades have increasingly narrowed the search space (Undagoitia & Rauch, 2015; Boveia & Doglioni, 2018), smoking-gun evidence is still elusive. In this thesis, we will concern ourselves with the clustering of dark matter, but a detailed discussion of its nature




Figure 1.2. – The temperature map of the CMB obtained from *Planck*’s observations using the SMICA algorithm (Adam et al., 2016). Note the remarkable isotropy – the temperature deviation is of the order 10^{-4} K, compared to the CMB temperature of 2.7 K. *Credit:* ESA and the *Planck* collaboration.

is beyond our scope. We assume that a majority of dark matter is made of identical particles, and that the particles are

- *cold*, i.e., their velocity dispersion is (at least at early times) negligible. This is supported by the fact that structure formation can only have happened in a bottom-up manner – smaller structures forming first, then accreting and merging to form larger ones (Blumenthal et al., 1984), and
- *collisionless*, i.e. they move under the influence of their aggregate potential, and two-body interactions are rare.

With these assumptions, the system of particles can be treated statistically, as we do in Sec. 2.1 (see also App. A).

1.4. The CMB, and the shape of the Universe

The CMB is observed to have a thermal blackbody spectrum with a temperature¹³ of ≈ 2.7 K (Fixsen, 2009), as expected from the conditions predicted by a hot big

¹³It is customary to report the CMB temperature in Kelvin. This is equivalent to $2.35 \times 10^{-4} \text{ eV } k_{\text{B}}^{-1}$.

bang. The temperature map of the CMB is highly isotropic with nearly-Gaussian statistics (e.g. Akrami et al., 2020), lending credence to an inflationary origin (see Fig. 1.2). The CMB can also be used to probe the geometry of the Universe. This can be quantified by the dimensionless energy-density parameter Ω_0 (the subscript indicates that it is measured at $z = 0$; see Sec. 2.1 for more details)

- $\Omega_0 < 1$ implies an *open* universe (negative spatial curvature),
- $\Omega_0 > 1$ implies a *closed* universe (positive spatial curvature), and
- $\Omega_0 = 1$ implies a *flat* universe (zero spatial curvature).

CMB observations from *Planck* put strong constraints on Ω_0 to be almost exactly unity. Thus, our best estimate is that the Universe is flat.

Taking all of these observational results together, the Universe is flat, and dominated ($\approx 69\%$ of the budget) by dark energy (Λ) which accelerates its expansion. The matter content ($\approx 31\%$) is dominated by a non-baryonic component – cold dark matter (CDM; $\approx 26\%$ of the total budget), and a (relatively) small amount of baryonic matter makes up the stars and galaxies that we can directly observe. Theoretical approaches set against the backdrop of this simple Λ CDM model make accurate predictions, and can explain nearly all observational results, which is a remarkable feat. However, there are lingering questions about the discrepancies between early- and late-type measurements (e.g. the Hubble tension), and the nature of the dark sector remains unresolved. In this chapter, we have narrated a broad story of the evolution of the Universe. Now, we narrow our focus to comment on several open questions regarding the clustering of dark matter, which we summarise below

- How do different analytical theories of structure formation compare with each other in the quasi-linear and non-linear regimes?
- Are the assumptions that these theories rely on justified?
- Gravitational instability leads to the formation of virialised structures called halos. How do their statistical and individual properties depend on the environments in which they are born?
- More generally, what is the best theoretical approach to map the initial dark matter overdensity field to collapsed structures?

In Chapters 2 and 3, we introduce the tools we exploit to answer these questions in depth. We attempt to answer the first two questions in Chapter 4 and the other two in Chapter 5. Chapter 6 summarises our results, and puts them into a broader context.

1.5. Notation and terminology

- Whenever possible, we use bold-faced (Latin) letters to denote vectors and co-vectors, and bold-faced Greek letters to denote higher-rank tensors. The exception to this is the text before Sec. 2.1, where we denote the tensors in the Einstein field equations with their conventional Latin symbols. The components of a tensor (of any rank) are written in italic font with indices. A tensor may thus be expressed in a component-basis form as $\mathbf{\Gamma} = \Gamma_{ij} \mathbf{e}_i \otimes \mathbf{e}_j$, where Γ_{ij} are the components of the tensor, and the $\mathbf{e}_i$ are the standard basis vectors.
- The ‘ $\otimes$ ’ above is a dyadic (tensor) product. For two vectors $\mathbf{u}$ and $\mathbf{v}$, their dyadic product is a rank-2 tensor with components $A_{ij} = u_i v_j$. The scalar product between two vectors is indicated by a dot (see below).
- We use a tilde over a quantity to indicate its Fourier transform (so that f and $\tilde{f}$ are duals). In our convention, the forward and inverse transforms are defined as follows

$$\tilde{f}(\mathbf{k}) = \frac{1}{(2\pi)^3} \int d^3\mathbf{x} \, f(\mathbf{x}) \exp(-i\mathbf{k} \cdot \mathbf{x}) , \quad (1.2a)$$

$$f(\mathbf{x}) = \int d^3\mathbf{k} \, \tilde{f}(\mathbf{k}) \exp(i\mathbf{k} \cdot \mathbf{x}) . \quad (1.2b)$$

- We use δ at different points in the text to indicate different quantities. The distinction is either obvious from the context, or explicitly stated. First, δ is used to denote the dark matter overdensity field. Second, δ_{ij} is used to denote the Kronecker delta, defined as

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j , \\ 0 & \text{if } i \neq j . \end{cases} \quad (1.3)$$

Finally, δ^D represents the three-dimensional Dirac delta (the continuous form of the Kronecker delta), defined as the inverse Fourier transform of unity

$$\delta^D(\mathbf{x}) = \int d^3\mathbf{k} \, \exp(i\mathbf{k} \cdot \mathbf{x}) . \quad (1.4)$$

- Operators, i.e. linear maps which act on a function to give a real number, are denoted by cursive letters (e.g., $\mathcal{H}, \mathcal{D}, \mathcal{L}, \dots$).
- We use ‘ e ’ and ‘ $\exp$ ’ interchangeably to indicate exponentiation.
- We do *not* opt for natural units, so k_B and c are explicitly written out every time they should appear.

CHAPTER 2

Theory

We infer our Universe using a wide range of observables. Electromagnetic radiation, the primary guide, comes in many wavelengths spanning a vast spectrum – from highly energetic X-rays that inform us about the dynamics of galaxy clusters, to radio waves that reveal the interiors of compact neutron stars. Neutrinos – through the cosmic neutrino background – may permit us a glimpse into the infancy of the Universe, shortly after the big bang. Gravitational waves offer a window to the mechanics of black hole mergers. Any theoretical framework must explain all of these observations, which is a monumental task. In this thesis, we are mainly concerned with the evolution of structure in the Universe – how we arrived at the rich, multiscale clustering that we observe today from the homogeneous, isotropic starting point immediately after inflation. The most important physical process that governs this evolution is gravitation. Thus, a natural starting point is a review of our theory of gravity. General relativity (GR) proposes that our Universe can be mathematically modelled by a four-dimensional, pseudo-Riemannian manifold (called spacetime) endowed with an intrinsic curvature. Then, gravity arises as a natural consequence of motion in curved spacetime. The theory is described by the Einstein field equations which characterise how matter and energy affect the curvature of spacetime,

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} . \quad (2.1)$$

The equation is solved by specifying the metric tensor $g_{\mu\nu}$ (or, equivalently the spacetime interval $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$) which defines the notion of distance on the manifold. The Ricci curvature tensor $R_{\mu\nu}$ describes deviations of the metric from Euclidean geometry, R is the scalar curvature (trace of $R_{\mu\nu}$), and $T_{\mu\nu}$ is the energy-momentum tensor describing the matter density of the Universe. The equation

is coupled to the geodesic equation which encodes the response of matter to the curvature of spacetime. The simplest solution assumes that the energy-momentum tensor vanishes identically, as in the absence of matter and all non-gravitational fields (i.e., a vacuum solution). In such a universe, the Ricci tensor is also zero, and we have a trivial solution – the Minkowski metric – defined as

$$ds_{\text{Minkowski}}^2 = c^2 dt^2 - (dx^2 + dy^2 + dz^2) . \quad (2.2)$$

Schwarzschild (1916) (see also Droste, 1917) proposed the first non-trivial solution in vacuum – describing the gravitational field outside a spherical mass. In fact, this is the most general spherically-symmetric vacuum solution, and no time-dependent solutions of this form exist (Jebsen, 1921; Birkhoff & Langer, 1923).

Friedmann (1924) and, independently, Lemaître (1927, 1931a) derived a class of solutions for a homogeneous, isotropic, expanding spacetime. Robertson (1935) and Walker (1937) proved the uniqueness of this solution. This is called the Friedmann-Lemaître-Robertson-Walker (FLRW) metric and is defined as

$$ds_{\text{FLRW}}^2 = c^2 dt^2 - a^2(t) \left\{ dr^2 + f_K^2(r) [d\theta^2 + \sin^2 \theta d\phi^2] \right\} , \quad (2.3)$$

where f_K defines a curvature-dependent comoving angular diameter distance. For flat spacetime, $f_K(r) = r$, and we recover the Minkowski metric in spherical coordinates, with an additional time-dependent cosmic expansion factor (hereafter, *scale factor*) that encodes the expansion of spacetime. A scale factor of a denotes the time when the Universe had a volume a^{-3} times the volume today (naturally, the scale factor at the present epoch is 1).

The FLRW metric proves to be an excellent approximation to the real Universe on large scales. This is a consequence of the fact that the Universe appears to be spatially homogeneous and isotropic on large scales. Although the veracity of this *cosmological principle* has come into increased scrutiny over the last decade (e.g. Aluri et al., 2023), it remains a useful guide. Given the FLRW metric, we can construct the other pieces of Eq. (2.1). First, consider the matter-energy content of the universe. As a first approximation, this can be modelled as a perfect fluid, i.e. a fluid without shear stress, anisotropic pressure, and viscosity. Operating in the fluid rest-frame, we have the following expression for the energy-momentum tensor

$$T_{\mu\nu} = \left(\rho + \frac{p}{c^2} \right) U_\mu U_\nu - p g_{\mu\nu} , \quad (2.4)$$

where ρ is the energy density of the fluid, p is its isotropic pressure, and $U_\mu = (c, 0, 0, 0)$ is the velocity four-vector. The non-vanishing components of $T_{\mu\nu}$ are

$$T_{00} = \rho c^2 , \quad T_{ii} = p a^2(t) . \quad (2.5)$$

The Ricci tensor is diagonal in the FLRW metric, and its non-zero components are

$$\begin{aligned} R_{00} &= 3\frac{\ddot{a}}{ac^2}, & R_{11} &= -\left(\frac{a\ddot{a} + 2\dot{a}^2}{c^2}\right), \\ R_{22} &= -\left(\frac{a\ddot{a} + 2\dot{a}^2}{c^2}\right)r^2, & R_{33} &= -\left(\frac{a\ddot{a} + 2\dot{a}^2}{c^2}\right)r^2\sin^2\theta, \end{aligned} \quad (2.6)$$

where the dot indicates a derivative with respect to proper time t . This leads to the following expression for the Ricci scalar

$$R = \frac{6}{c} \left[\frac{\ddot{a}}{a} + \left(\frac{\dot{a}}{a}\right)^2 \right]. \quad (2.7)$$

Substituting the relevant quantities in Eq. (2.1), we get

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3}, \quad (2.8)$$

$$\frac{2\ddot{a}}{a} + \left(\frac{\dot{a}}{a}\right)^2 = -\frac{8\pi G\rho}{c^2}. \quad (2.9)$$

The two equations are often combined to derive

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} \left(\rho + \frac{3p}{c^2} \right). \quad (2.10)$$

Eq. (2.8) and Eq. (2.10), taken together, determine the evolution of $a(t)$ and are called the Friedmann equations. To solve the Friedmann equations, we need a model for the time evolution of the fluid density and pressure. To that end, we can examine the conservation laws associated with $T_{\mu\nu}$. The conservation of momentum implies the translational invariance of the metric, and reveals no new information. The conservation of energy ($\nabla_\nu T^{0\nu}$) leads to the following equation

$$\frac{\partial\rho}{\partial t} + \frac{3\dot{a}}{a} \left(\frac{p}{c^2} + \rho \right) = 0. \quad (2.11)$$

This admits a solution of the form $p = w\rho$. For a constant w , this leads to an equation of state

$$\rho \propto a^{-3(1+w)}. \quad (2.12)$$

Different components of matter and energy can be identified with different values of w . For non-relativistic baryonic matter, $w = 0$, while for relativistic photons, $w = 1/3$. When modelled as a perfect fluid with negative pressure, dark energy has an equation of state with $w = -1$. The total matter-energy density of the Universe is then given as the sum of these three components

$$\rho = \frac{\rho_{\text{m},0}}{a^3} + \frac{\rho_{\text{r},0}}{a^4} + \rho_\Lambda, \quad (2.13)$$

where the subscript zero indicates that the values are taken to be at the present epoch. For the Universe to be flat, the present-day total matter-energy density must equal a critical value, defined as

$$\rho_{\text{cr}} = \frac{3H_0^2}{8\pi G}, \quad (2.14)$$

where H_0 is the Hubble constant. The densities in Eq. (2.13) are typically scaled by ρ_{cr} to obtain dimensionless quantities related as

$$\Omega_0 = \Omega_{\text{m}} + \Omega_{\text{r}} + \Omega_{\Lambda}. \quad (2.15)$$

We have discussed the effect of Ω_0 on the geometry of the Universe in Sec. 1.4. Now, we turn our attention to the process of structure formation – the gravitational interaction of cold, collisionless dark matter particles to form the multiscale clustering we observe today. Broadly, structure formation can be studied in the linear and quasi-linear regimes using perturbative techniques which we describe in Sec. 2.1 and Sec. 2.2. In the non-linear regime, analytical methods are limited to approximations. We introduce some such techniques in the context of halo formation in Sec. 2.3. In this regime, N -body simulations are the primary computational tool, and they are discussed in Sec. 3.1.

2.1. Cosmological perturbation theory

As discussed in Sec. 1.3, we assume a majority of dark matter to be made of identical, classical particles. Although there is a wide range of particle masses that have not been observationally ruled out, all of them would imply a very large particle number. Such a collection can be modelled statistically by a distribution function. In full generality, this is the N -particle distribution function that resides in the $6N$ -dimensional Γ -space (six numbers for position and momentum per particle; see App. A), and whose evolution is governed by the BBGKY hierarchy of equations. The approximate treatment that is sufficient for our purposes, models the system using the one-particle distribution function (since collisions are assumed to be rare). It is defined such that its integral over a six-dimensional phase-space volume is the number of particles contained in that volume. Such a system follows Liouville's theorem – *the evolution of the distribution function along trajectories of the system in phase space is constant* (see App. A). This reflects the conservation of phase space density, and can be expressed mathematically as the collisionless-Boltzmann (Vlasov) equation

$$\frac{\partial f}{\partial t} + \frac{\partial f}{\partial \mathbf{x}} \cdot \frac{\partial \mathbf{x}}{\partial t} + \frac{\partial f}{\partial \mathbf{p}} \cdot \frac{\partial \mathbf{p}}{\partial t} = 0, \quad (2.16)$$

where $\mathbf{x}$ is the comoving position and $\mathbf{p} = ma^2\dot{\mathbf{x}}$ is its conjugate momentum. To incorporate the effect of self-gravity, this is coupled with the Poisson equation which

relates the matter overdensity $\delta = \rho/\bar{\rho} - 1$ (where ρ is the matter density, and $\bar{\rho}$ is the mean background density) to the gravitational potential ϕ

$$\nabla^2 \phi(\mathbf{x}) = 4\pi G \bar{\rho} a^2(t) \delta(\mathbf{x}) . \quad (2.17)$$

We have implicitly made an important assumption here. Eq. (2.17) is the Poisson equation for *Newtonian* gravity. We are thus operating in the weak-field ($|\phi| \ll c^2$), small-velocity ($v \ll c$) limit. This is well-justified – at the scales that are of interest to us, general relativistic effects are subdominant. For one, our scales of interest are much smaller than the horizon scale. Also, in the late Universe, radiation plays no role in structure formation, and the only source of gravity is non-relativistic matter. Taken together, these facts allow us to work within the Newtonian approximation.

Solving the Vlasov-Poisson system would provide us with a full picture of how the dark matter particles evolve. Analytically, the full six-dimensional equation can only be solved under special, highly-symmetric conditions. Numerical solutions are also intractably expensive. The typical approach in cosmology is to take momentum moments of the Vlasov equation and solve them perturbatively. First, we summarise the n^{th} moment of the distribution as

$$\boldsymbol{\mu}^{(n)}(\mathbf{x}, t) = \frac{m}{a(t)^3} \int d^3 \mathbf{p} \left[\frac{\mathbf{p}}{ma(t)} \right]^n f(\mathbf{x}, \mathbf{p}, t) , \quad (2.18)$$

where the scaling m/a^3 ensures that the moments are mass-weighted, and in physical (as opposed to comoving) coordinates. Let us name the first few moments to simplify the nomenclature later on – $\boldsymbol{\mu}(\mathbf{x}, t) = \{\rho, \mathbf{P}, \boldsymbol{\Xi}, \dots\}$. Here ρ is a mass density, $\mathbf{P}$ is a momentum density with $\mathbf{V} = \mathbf{P}/\rho$ being the average particle velocity, and $\boldsymbol{\Xi}$ is a tensor which can be decomposed as

$$\boldsymbol{\Xi}(\mathbf{x}, t) = \rho(\mathbf{x}, t) \mathbf{V}(\mathbf{x}, t) \mathbf{V}(\mathbf{x}, t) + \boldsymbol{\Gamma}(\mathbf{x}, t) . \quad (2.19)$$

Here, $\boldsymbol{\Gamma}$ is the velocity dispersion tensor, and describes how particle velocities are spread about their mean. The set of moment equations is generated by taking successive moments of Eq. (2.16). Although each individual moment equation is simpler to solve than the Vlasov equation, the full hierarchy is infinite. The first two moment equations are given as follows

$$\dot{\delta} + \frac{1}{a} \boldsymbol{\nabla} \cdot [(1 + \delta) \mathbf{V}] = 0 , \quad (2.20)$$

$$\dot{\mathbf{V}} + H \mathbf{V} + \frac{1}{a} (\mathbf{V} \cdot \boldsymbol{\nabla}) \mathbf{V} + \frac{1}{a} \boldsymbol{\nabla} \phi + \frac{1}{a\rho} \boldsymbol{\nabla} \cdot (\rho \boldsymbol{\Gamma}) = 0 , \quad (2.21)$$

and treat the system as a fluid. The equations express the conservation of mass and momentum of the fluid respectively. The n^{th} moment equation couples the n^{th} and $(n + 1)^{\text{th}}$ moments of f . A non-zero velocity dispersion generates successively

higher-order moments with time, making it impossible to truncate the hierarchy. The only truncation that does not generate higher-order moments is at second-order by setting $\mathbf{\Gamma}$ to zero. In this truncated form, the Vlasov hierarchy admits a solution of the type

$$f(\mathbf{x}, \mathbf{p}, t) = \rho(\mathbf{x}) \delta^D[\mathbf{p} - \nabla \phi(\mathbf{x}, t)] , \quad (2.22)$$

where δ^D is the Dirac delta. This *dust* model breaks down in the presence of velocity dispersion, leading to formation of caustics on small scales. Nonetheless, it is a good approximation in the linear regime of structure formation, and remains valid into the non-linear regime at large scales. With this assumption, the two moment equations coupled with the Poisson equation can be solved perturbatively. In what follows, we describe the Eulerian frame solutions, which is referred to as standard perturbation theory (SPT, e.g., Vishniac, 1983; Fry, 1984; Jain & Bertschinger, 1994; Scoccimarro, 1997; Bernardeau et al., 2002). An equivalent Lagrangian frame solution has also been widely studied (e.g. Zeldovich, 1970; Shandarin & Zeldovich, 1989; Buchert, 1989; Bouchet et al., 1994; Ehlers & Buchert, 1997).

2.1.1. Standard (Eulerian) perturbation theory

In the Eulerian frame of reference, the fluid elements are positioned on a fixed grid, and they are time-dependent. The comoving coordinates $\mathbf{x}$ are defined in the Eulerian frame. Under the assumption of zero velocity dispersion, the vorticity (curl of the velocity) of the fluid vanishes. Even if vorticity is initially non-zero, it is damped by a factor of a , and decays as the Universe expands. Given this, it is reasonable to assume that the velocity field is irrotational, i.e., that it is completely specified by its divergence $\theta(\mathbf{x}, t) = \nabla \cdot \mathbf{V}(\mathbf{x}, t)$. We can cast equations Eq. (2.20) and Eq. (2.21) in terms of δ and θ as

$$\frac{\partial \delta}{\partial t} + \frac{1}{a} \nabla \cdot [(1 + \delta) \mathbf{V}] = 0 , \quad (2.23)$$

$$\dot{\theta} + H\theta + 4\pi G \bar{\rho} a \delta + \frac{1}{a} \nabla \cdot (\mathbf{V} \cdot \nabla \mathbf{V}) = 0 , \quad (2.24)$$

where we have taken the divergence of Eq. (2.21) and used Poisson's equation (Eq. 2.17) to substitute ϕ . SPT assumes that the equations can be solved order-by-order by expanding δ and θ in a power series. First, the linearised equations are solved. Then, the linear solution is substituted in Eq. (2.23) and Eq. (2.24), the resulting equations are solved, and so on. In practice, the non-linear solutions are best obtained in Fourier space, where the equations take the following form

$$a \frac{\partial \tilde{\delta}(\mathbf{k}, t)}{\partial t} + \tilde{\theta}(\mathbf{k}, t) = - \int d^3 \mathbf{k}_1 d^3 \mathbf{k}_2 \delta^D(\mathbf{k} - \mathbf{k}_{12}) \alpha(\mathbf{k}_1, \mathbf{k}_2) \tilde{\theta}(\mathbf{k}_1, t) \tilde{\delta}(\mathbf{k}_2, t) , \quad (2.25)$$

$$a \frac{\partial \tilde{\theta}(\mathbf{k}, t)}{\partial t} + \dot{a}(t) \tilde{\theta}(\mathbf{k}, t) + 4\pi G \bar{\rho} a^2 \tilde{\delta} = - \int d^3 \mathbf{k}_1 d^3 \mathbf{k}_2 \delta^D(\mathbf{k} - \mathbf{k}_{12}) \beta(\mathbf{k}_1, \mathbf{k}_2) \tilde{\theta}(\mathbf{k}_1, t) \tilde{\theta}(\mathbf{k}_2, t) , \quad (2.26)$$

where $\mathbf{k}_{12} = \mathbf{k}_1 + \mathbf{k}_2$, and

$$\alpha(\mathbf{k}_1, \mathbf{k}_2) = \frac{\mathbf{k}_{12} \cdot \mathbf{k}_1}{\mathbf{k}_1} , \quad \beta(\mathbf{k}_1, \mathbf{k}_2) = \frac{k_{12}^2 (\mathbf{k}_1 \cdot \mathbf{k}_2)}{2k_1^2 k_2^2} \quad (2.27)$$

are functions which quantify the coupling of Fourier modes. In the following section, we will use the SPT solutions in an Einstein-de Sitter (EdS) cosmology. The EdS solution for δ and θ decomposes the position- and time-dependence of the fields (e.g. Goroff et al., 1986; Makino et al., 1992; Bernardeau et al., 2002)

$$\tilde{\delta}(\mathbf{k}, t) = \sum_{n=1}^{\infty} a^n(t) \tilde{\delta}_n(\mathbf{k}) , \quad \tilde{\theta}(\mathbf{k}, t) = -\dot{a}(t) \sum_{n=1}^{\infty} a^n(t) \tilde{\theta}_n(\mathbf{k}) , \quad (2.28)$$

where $\tilde{\delta}_n$ and $\tilde{\theta}_n$ are time-independent fields expressible in terms of the linear solution as

$$\tilde{\delta}_n(\mathbf{k}) = \int d^3 \mathbf{q}_1 \dots \int d^3 \mathbf{q}_n \delta^D(\mathbf{k} - \mathbf{q}_{1\dots n}) F_n(\mathbf{q}_1, \dots, \mathbf{q}_n) \tilde{\delta}_1(\mathbf{q}_1) \dots \tilde{\delta}_1(\mathbf{q}_n) , \quad (2.29)$$

$$\tilde{\theta}_n(\mathbf{k}) = \int d^3 \mathbf{q}_1 \dots \int d^3 \mathbf{q}_n \delta^D(\mathbf{k} - \mathbf{q}_{1\dots n}) G_n(\mathbf{q}_1, \dots, \mathbf{q}_n) \tilde{\delta}_1(\mathbf{q}_1) \dots \tilde{\delta}_1(\mathbf{q}_n) . \quad (2.30)$$

Here, $\mathbf{q}_{1\dots n} = \mathbf{q}_1 + \dots + \mathbf{q}_n$, and F_n and G_n are kernels that obey the following recurrence relations (see e.g. Bernardeau et al., 2002). The first-order kernels $F_1 = G_1 = 1$, and for $n > 1$

$$F_n(\mathbf{q}_1 \dots \mathbf{q}_n) = \sum_{m=1}^{n-1} \frac{G_m(\mathbf{q}_1 \dots \mathbf{q}_m)}{(2n+3)(n-1)} [(2n+1)\alpha(\mathbf{k}_1, \mathbf{k}_2) F_{n-m}(\mathbf{q}_{m+1} \dots \mathbf{q}_n)] \\ + 2\beta(\mathbf{k}_1, \mathbf{k}_2) G_{n-m}(\mathbf{q}_{m+1} \dots \mathbf{q}_n) , \quad (2.31)$$

$$G_n(\mathbf{q}_1 \dots \mathbf{q}_n) = \sum_{m=1}^{n-1} \frac{G_m(\mathbf{q}_1 \dots \mathbf{q}_m)}{(2n+3)(n-1)} [3\alpha(\mathbf{k}_1, \mathbf{k}_2) F_{n-m}(\mathbf{q}_{m+1} \dots \mathbf{q}_n)] \\ + 2n\beta(\mathbf{k}_1, \mathbf{k}_2) G_{n-m}(\mathbf{q}_{m+1} \dots \mathbf{q}_n) . \quad (2.32)$$

The kernels are often symmetrised by averaging over permutations of their arguments (e.g. Goroff et al., 1986). We denote the symmetric kernels by a superscript as $F_n^{(s)}$ and $G_n^{(s)}$.

Note that higher-order perturbative solutions are proportional to powers of the linear overdensity. In practice, the full overdensity field is not observable. What we *can* observe is statistical summaries of galaxy clustering. Galaxies themselves are biased tracers, and do not follow the dark matter potential wells perfectly. This *galaxy bias* is a crucial component of cosmological inference, but it is beyond the scope of this thesis. Instead, we turn to the summaries of dark matter clustering.

Let us try to understand how the overdensities at two nearby points $\mathbf{x}_1$ and $\mathbf{x}_2$ in configuration space are related to each other at any given time t . Treating $\delta(\mathbf{x}_1)$ and $\delta(\mathbf{x}_2)$ as random variables, we can measure their *correlation*. A high positive

(negative) correlation implies that if $\delta(\mathbf{x}_1)$ is larger than its mean value, then $\delta(\mathbf{x}_2)$ is likely to be larger (smaller) than *its* mean value. A zero correlation points to two random variables being independent. A commonly-employed measure of correlation is the *two-point correlation function* (2PCF) defined as

$$\xi(\mathbf{x}_1, \mathbf{x}_2) = \langle \delta(\mathbf{x}_1) \delta(\mathbf{x}_2) \rangle , \quad (2.33)$$

where $\langle \cdot \rangle$ indicates a spatial averaging. This is the *excess-over-random* probability of finding a dark matter particle at $\mathbf{x}_2$ given another particle at $\mathbf{x}_1$. Assuming translational invariance, the 2PCF does not depend on where the points are located, but only on their separation vector (this is equivalent to assuming spatial homogeneity)

$$\xi(\mathbf{x}_1 - \mathbf{x}_2) = \langle \delta(\mathbf{x}_1) \delta(\mathbf{x}_2) \rangle . \quad (2.34)$$

As we did before, it is often useful to work in Fourier-space

$$\begin{aligned} \langle \tilde{\delta}(\mathbf{k}_1) \tilde{\delta}(\mathbf{k}_2)^* \rangle &= \int \frac{d^3 \mathbf{x}_1}{(2\pi)^3} \frac{d^3 \mathbf{x}_2}{(2\pi)^3} e^{-i\mathbf{k}_1 \cdot \mathbf{x}_1} e^{i\mathbf{k}_2 \cdot \mathbf{x}_2} \langle \delta(\mathbf{x}_1) \delta(\mathbf{x}_2)^* \rangle \\ &= \int \frac{d^3 \mathbf{x}_1}{(2\pi)^3} \frac{d^3 \mathbf{x}_2}{(2\pi)^3} e^{-i\mathbf{k}_1 \cdot \mathbf{x}_1} e^{i\mathbf{k}_2 \cdot \mathbf{x}_2} \xi(\mathbf{x}_1 - \mathbf{x}_2) . \end{aligned} \quad (2.35)$$

To make further progress, we introduce the following coordinate transformation

$$\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{r} + \mathbf{R} \\ \mathbf{R} \end{bmatrix} . \quad (2.36)$$

The determinant of the Jacobian of this transformation is unity, and the volume elements transform in a simple way. Then Eq. (2.35) takes the following form

$$\langle \tilde{\delta}(\mathbf{k}_1) \tilde{\delta}(\mathbf{k}_2)^* \rangle = \int \frac{d^3 \mathbf{r}}{(2\pi)^3} \frac{d^3 \mathbf{R}}{(2\pi)^3} e^{-i\mathbf{k}_1 \cdot (\mathbf{R} + \mathbf{r})} e^{i\mathbf{k}_2 \cdot \mathbf{R}} \xi(\mathbf{r}) . \quad (2.37)$$

This can be written as a product of two independent integrals

$$\langle \tilde{\delta}(\mathbf{k}_1) \tilde{\delta}(\mathbf{k}_2)^* \rangle = \int \frac{d^3 \mathbf{r}}{(2\pi)^3} e^{-i\mathbf{k}_1 \cdot \mathbf{r}} \xi(\mathbf{r}) \cdot \int \frac{d^3 \mathbf{R}}{(2\pi)^3} e^{-i(\mathbf{k}_1 - \mathbf{k}_2) \cdot \mathbf{R}} . \quad (2.38)$$

Re-labeling $\mathbf{k}_1 \rightarrow \mathbf{k}$ and $\mathbf{k}_2 \rightarrow \mathbf{k}'$,

$$\langle \tilde{\delta}(\mathbf{k}) \tilde{\delta}(\mathbf{k}')^* \rangle = P(\mathbf{k}) \delta^D(\mathbf{k} - \mathbf{k}') , \quad (2.39)$$

where we have used the definition of the Dirac delta (Eq. 1.4) for the second integral, and

$$P(\mathbf{k}) = \int \frac{d^3 \mathbf{r}}{(2\pi)^3} e^{-i\mathbf{k} \cdot \mathbf{r}} \xi(\mathbf{r}) \quad (2.40)$$

is the Fourier transform of the 2PCF. It is called the *power spectrum* of the overdensity field, and measures the amplitude of fluctuations at different scales. Examining Eq. (2.39) reveals that if $\mathbf{k} \neq \mathbf{k}'$, their correlation is zero – different Fourier modes are uncorrelated with each other. Solving the PT equations, the leading-order contribution to the power spectrum is second-order in the overdensity. This is the correlation of two linear-order fields, and its Fourier-space counterpart is thus called the *linear* power spectrum

$$P_L(\mathbf{k}, t) = a^2(t) \langle \delta_1(\mathbf{k}) \delta_1(\mathbf{k}')^* \rangle, \quad (2.41)$$

where we re-instate the time-dependence explicitly. This is also called the ‘tree-level’ (or ‘0-loop’) power spectrum. Importantly, if the linear overdensity is Gaussian, the only allowed correlators are those with pairs of fields. This is guaranteed by Isserlis’ theorem¹, which proves that the odd N -point correlators are zero for a Gaussian field, and the even ones are expressible as a sum of the product of different two-point correlators (Isserlis, 1918). Then, the next-to-leading order contribution to the power spectrum must be fourth-order in δ_1 , and the power spectrum that includes this contribution is termed ‘1-loop’

$$\begin{aligned} P_{1\text{-loop}}(\mathbf{k}, t) &= a^2(t) \langle \delta_1(\mathbf{k}) \delta_1(\mathbf{k}')^* \rangle + 2a^4(t) [\langle \delta_1(\mathbf{k}) \delta_3(\mathbf{k}')^* \rangle + \langle \delta_2(\mathbf{k}) \delta_2(\mathbf{k}')^* \rangle] \\ &= P_L(\mathbf{k}, t) + 2P_{13}(\mathbf{k}, t) + P_{22}(\mathbf{k}, t). \end{aligned} \quad (2.42)$$

An important simplification here comes from the assumption of *statistical isotropy*, which implies that there is no preferred direction in space, and the correlators depend only on the magnitude of separation between the two points. From now on, we will indicate this isotropy by writing $P(k)$ instead of $P(\mathbf{k})$. The solutions for the 1-loop corrections are expressed as the following integrals

$$P_{13}(k, t) = 3P_L(k) \int d^3\mathbf{q} P_L(q) F_3^{(s)}(\mathbf{q}, -\mathbf{q}, \mathbf{k}), \quad (2.43)$$

$$P_{22}(k, t) = 2 \int d^3\mathbf{q} P_L(q) P_L(|\mathbf{k} - \mathbf{q}|) \left[F_2^{(s)}(\mathbf{q}, -\mathbf{q}, \mathbf{k}) \right]^2 \quad (2.44)$$

As a consequence of Isserlis’ theorem, the 2PCF (or equivalently, the power spectrum) fully characterises a Gaussian field. For non-Gaussian fields, the correlation hierarchy extends ad infinitum. Cosmological fields in the early Universe are Gaussian or nearly-Gaussian. However, gravitational evolution and astrophysical phenomena couple different spatial scales and render the late-time fields highly non-Gaussian. The power spectrum is, in these cases, a limited measure of correlation, and information leaks into the higher-order correlators (Samushia et al., 2021). In practice, these correlators are increasingly difficult to model – their utility and their alternatives are subject to further research.

¹A version of this result for operators in quantum field theory was proved by Wick (1950).

At lowest-order in the correlation hierarchy, PT predicts the power spectrum of the dark matter overdensity field. To close the gap to observational data, the canonical first step is to define a unit of clustering called a dark matter *halo*, which then hosts observable galaxies. We discuss the modelling of halos in detail in Sec. 2.3. PT and its extensions have been successful in predicting the clustering of structure on large scales (e.g. Crocce & Scoccimarro, 2006a; Taruya et al., 2009; Carlson et al., 2009). However, PT suffers from several conceptual and practical problems, some of which we highlight below

- In the non-linear regime, loop corrections can grow large enough to be of the same order as the linear power spectrum. Further, higher-order loop corrections may grow larger than their lower-order counterparts (e.g. Carlson et al., 2009). This is a consequence of the perturbation variable δ no longer being small. One way to overcome this is by leveraging field-theoretic methods to reorganise the perturbative series. The key ingredient is a propagator that captures the evolution of density (or velocity) fluctuations. The propagator is renormalised and re-summed, leading to a well-defined PT (Crocce & Scoccimarro, 2006a,b). When cast in the language of the renormalisation group, a UV cutoff naturally arises, which enhances the convergence of the loop integrals on small scales (McDonald, 2006; Matarrese & Pietroni, 2007). However, the underlying issue persists – PT treats non-linear modes as perturbative.
- PT assumes the dark matter fluid to be pressureless and perfect. The assumption of non-zero velocity dispersion breaks on successively larger scales due to non-linear gravitational evolution, and PT stops making meaningful predictions beyond shell crossing (Pueblas & Scoccimarro, 2009).
- The loop integrals of SPT suffer from divergences on large- and small-scales. For scale-free power spectra – $P(k) \propto k^n$ – with spectral index $n < -1$, the leading-order large-scale (IR) divergences are expected to cancel at each order in the PT, though subdominant divergences may remain (Jain & Bertschinger, 1996). Physically, this cancellation is a consequence of the fact that the equations of motion are invariant under Galilean transformations. The UV-divergences that arise from small-scale effects are, however, more terminal. Their cancellation requires counterterms to include the effect of small-scale physics, which is beyond the scope of SPT.

Due to these fundamental issues that arise from small-scale, non-linear effects, the focus of much of the recent theoretical work has been in using an effective theory to predict structure formation. Such a theory would retain the structure of PT, with a coarse-graining scale that explicitly integrates out the small-scale physics. Invoking the renormalisation techniques often used in field theory, this approach alleviates a majority of the problems with PT. We discuss this in detail the following section.

2.2. Effective field theory of large-scale structure

Physical processes in the Universe tend to unfold on a wide range of scales. An effective theory is a tool that allows us to extract meaningful physics from specific scales, while remaining largely agnostic to all other scales. A remarkable observation about our Universe is that the physical processes on large spatial (low-energy, IR) scales *decouple* from the processes on small (high-energy, UV) scales. This allows us to treat the two scales independently as a first-order approximation. Higher-order corrections can be made to model the coupling of scales to an arbitrary degree of precision (at least, in principle). Effective field theories (EFTs) find their utility in cases where

- the underlying UV physics is either unknown (e.g., GR can be thought of as a low-energy EFT of some yet-unknown theory of quantum gravity), or
- calculations within the UV-complete theory are intractable. For example, the theory of the strong interaction – quantum chromodynamics – is non-perturbative at low energies, and an EFT called chiral perturbation theory is constructed.

In cosmology, EFTs were first introduced in the context of inflation (Creminelli et al., 2006; Cheung et al., 2008). An EFT has also been formulated to model dark energy (Gubitosi et al., 2013; Frusciante & Perenon, 2020). In what follows, we formally introduce the effective field theory of large-scale structure (EFTofLSS; e.g., Baumann et al., 2012; Carrasco et al., 2012, 2014b) – an EFT for the growth of density fluctuations. In Karandikar et al. (2024) (hereafter, K24), we performed controlled numerical experiments to test the validity of the assumptions that the EFTofLSS makes. The results of our work are detailed in Chapter 4. This section is intended to serve as the necessary and sufficient background to understand the contents of Chapter 4, and consequently shares some material with K24.

Let us return to the shortcomings of SPT we discussed in the previous section, and see how they may be addressed. First, consider the fact that SPT treats non-linear scales as perturbative. Define the non-linear scale k_{NL} to be the scale at which density fluctuations are typically larger than unity. In our Universe, this is roughly the scale of galaxy clusters – of the order of a few Mpc. To remove spatial scales larger than k_{NL} (or wavenumbers smaller than k_{NL}) from the theory, we introduce a momentum cutoff $\Lambda < k_{\text{NL}}$, and coarse-grain our equations of motion (a similar approach was discussed in Pietroni et al., 2012). Next, let us relax the assumption of zero velocity dispersion, and include an *effective* stress tensor in the equations of motion that is explicitly sourced by modes with $k > k_{\text{NL}}$. Finally, the presence of this effective stress tensor leads to counterterms in the expressions for the correlators, which cancel the UV-divergences in their loop integrals. The EFTofLSS is a synthesis of all these ingredients, and thus alleviates the conceptual problems of PT.

2.2.1. Equations of the EFTofLSS

To start the construction, we define the coarse-graining operation as a convolution with a general window function as follows

$$g_\ell(\mathbf{x}) = \int d^3\mathbf{x}' W_\Lambda(\mathbf{x} - \mathbf{x}') g(\mathbf{x}) , \quad (2.45)$$

where the subscript ℓ indicates a long-wavelength quantity, and Λ is the coarse-graining scale. In principle, the choice of the window function should not change the final results. However, a Gaussian kernel greatly simplifies calculations (e.g. Carrasco et al., 2012), and we will define the window as

$$W_\Lambda(\mathbf{x}) = \left(\frac{\Lambda}{\sqrt{2\pi}} \right)^3 \int d^3\mathbf{x} \exp\left(-\frac{x^2}{2\Lambda^2}\right) . \quad (2.46)$$

We return to the Vlasov equation (Eq. 2.16) and apply the Gaussian window

$$\frac{\partial f_\ell}{\partial t} + \frac{\partial f_\ell}{\partial \mathbf{x}} \cdot \frac{\mathbf{p}}{ma^2} + m \int d^3\mathbf{x}' W_\Lambda(\mathbf{x} - \mathbf{x}') \frac{\partial \phi}{\partial \mathbf{x}'} \cdot \frac{\partial f}{\partial \mathbf{p}} = 0 , \quad (2.47)$$

where ϕ , as before, is the gravitational potential. We proceed (as in SPT) by taking successive moments of Eq. (2.47). The first two-moment equations are the coarse-grained analogues of Eq. (2.23) and Eq. (2.24)

$$\frac{\partial \delta_\ell}{\partial t} + \frac{1}{a} \nabla \cdot [(1 + \delta_\ell) \mathbf{U}] = 0 , \quad (2.48)$$

$$\frac{\partial \Theta}{\partial t} + H\Theta + 4\pi G \bar{\rho} a \delta_\ell + \frac{1}{a} \nabla \cdot (\mathbf{U} \cdot \nabla \mathbf{U}) + \frac{1}{a\rho_\ell} \nabla^2 \boldsymbol{\tau} = 0 , \quad (2.49)$$

where² $\mathbf{U} = \mathbf{P}_\ell / \rho_\ell$ ($\mathbf{P}$ as defined in Sec. 2.1), $\Theta = \nabla \cdot \mathbf{U}$, and $\boldsymbol{\tau}$ is the effective stress tensor. This is the coarse-grained analogue of $\boldsymbol{\Gamma}$, which was set to zero in SPT. It has been shown in the literature (e.g. Buchert & Dominguez, 2005; Carrasco et al., 2012) that $\boldsymbol{\tau}$ decomposes into a kinetic piece that describes the velocity dispersion within a coarse-grained volume, and a gravitational piece that arises from the coarse-grained gravitational potential

$$\begin{aligned} \boldsymbol{\tau}(\mathbf{x}, t) &= \boldsymbol{\Xi}_\ell(\mathbf{x}, t) - \rho_\ell(\mathbf{x}, t) \mathbf{U}(\mathbf{x}, t) \mathbf{U}(\mathbf{x}, t) \\ &+ \frac{1}{4\pi G a(t)^2} [\mathbf{w}_\ell(\mathbf{x}, t) - \nabla \phi_\ell(\mathbf{x}, t) \nabla \phi_\ell(\mathbf{x}, t)] , \end{aligned} \quad (2.50)$$

where³

$$\mathbf{w}_\ell(\mathbf{x}, t) = \int d^3\mathbf{x}' W_\Lambda(\mathbf{x} - \mathbf{x}') [\nabla \phi(\mathbf{x}', t) \nabla \phi(\mathbf{x}', t)] . \quad (2.51)$$

²Strictly speaking, the coarse-graining of a product is generically different from the product of the coarse-grained quantities. Thus, we do not use the subscript ℓ for such quantities.

³A subtlety we have ignored so far is that as one passes from a continuum description to discrete particles, the contribution of the self-interaction between particles must be subtracted out. This implies a modification to Eq. (2.47) as well as Eq. (2.51) (see e.g. Carrasco et al., 2012). However, in case of the planar-symmetry we are about to enforce, the equations as written above are consistent with McQuinn & White (2016) and K24.

The effective stress tensor as defined above contains explicit dependence on modes with $k > \Lambda$. To make further progress, it should be modelled in terms of the long-wavelength variables (viz. δ_ℓ and Θ) and the high- k modes must be integrated out. The EFTofLSS addresses this by assuming that the stress tensor can be written as the sum of a deterministic piece that depends on these variables, and a stochastic piece that is uncorrelated with them. The deterministic piece can then be expressed as an expectation value over the $k > \Lambda$ modes conditioned on fixing the modes $k < \Lambda$. Constructing this ensemble in practice is a formidable task, details of which may be found in Chapter 4 and in K24. Assuming the existence of this expectation value, we can express it as a Maclaurin series in the long-wavelength variables, i.e.

$$\begin{aligned} \langle \tau \rangle = & \langle \tau \rangle_{(0,0)} + \left. \frac{\partial \langle \tau \rangle}{\partial \delta_\ell} \right|_{(0,\Theta)} \delta_\ell + \left. \frac{\partial \langle \tau \rangle}{\partial \delta_\ell} \right|_{(\delta_\ell,0)} \Theta \\ & + \frac{1}{2!} \left. \frac{\partial^2 \langle \tau \rangle}{\partial \delta_\ell \partial \Theta} \right|_{(0,0)} \delta_\ell \Theta + \dots + J. \end{aligned} \quad (2.52)$$

The coefficients of this expansion are the free parameters of the theory to be measured/estimated from numerical/observational data. In the language of EFTs, they are called *Wilson* coefficients (Wilson, 1983). We can interpret these coefficients as properties of the dark matter fluid that couple the short- and long-wavelengths. We summarise below each part of the above equation

- $p_b = \langle \tau \rangle_{(0,0)}$ is an effective background pressure. It is spatially homogeneous, and thus does not affect the dynamics.
- $c_s^2 = \left. \frac{\partial \langle \tau \rangle}{\partial \delta_\ell} \right|_{(0,\Theta)}$ is a speed of sound that measures the strength of perturbations in δ_ℓ .
- $\dot{a}c_v^2 = \left. \frac{\partial \langle \tau \rangle}{\partial \Theta} \right|_{(\delta_\ell,0)}$ is a viscosity coefficient, often expressed in units of speed to align more clearly with c_s^2 . The factor $\dot{a}$ is necessary to keep the units consistent.
- $\frac{1}{2!} \left. \frac{\partial^2 \langle \tau \rangle}{\partial \delta_\ell \partial \Theta} \right|_{(0,0)}$ is the coefficient of the second-order term, and its amplitude is subdominant compared to the leading-order coefficients. The ellipses collect all other higher-order terms.
- J is the stochastic term that is uncorrelated with long-wavelength fields.

With this interpretation in mind, we can finally make the promised jump to the one-dimensional universe that forms the backdrop of K24. Practically, this is accomplished by considering perturbations with planar symmetry such that all tensors reduce to scalar quantities. This leads to a simplification of the full three-dimensional

case, which nonetheless retains some of its important features (see also McQuinn & White, 2016). To emphasise this, consider the overdensity and velocity divergence fields. In linear theory, $\delta_1 = \theta_1$ (this can be inferred from Eqs 2.29 and 2.30). Non-linear evolution breaks this degeneracy in the three-dimensional case with the origin of vorticity. However, this degeneracy persists into the quasi-linear regime in one-dimension. Effectively, this means that we have only one dynamical variable for the expansion of τ , and c_s^2 and c_v^2 are nearly degenerate (see Sec. 4 in K24 for more details). Thus, we define the combination of the two⁴, $c_{\text{tot}}^2 = c_s^2 - \dot{a}c_v^2$, as the sole Wilson coefficient of interest in our setup. We call this quantity the effective speed of sound. We leave the details of how c_{tot}^2 is estimated to Chapter 4, and now discuss how it impacts the power spectrum.

2.2.2. EFT corrections

Let us begin by reproducing the two moment equations in one-dimension

$$\partial_t \delta_\ell + \frac{1}{a} \partial_x [(1 + \delta_\ell)U] = 0, \quad (2.53)$$

$$\partial_t \Theta + H\Theta + 4\pi G \bar{\rho} a \delta_\ell + \frac{1}{a} \partial_x (U \partial_x U) + \frac{c_s^2}{a} \partial_x^2 \delta_\ell + \frac{\dot{a}c_v^2}{a} \partial_x^2 \Theta + \frac{1}{a\rho_\ell} \partial_x^2 J = 0, \quad (2.54)$$

where we have used the expression for the scalar form of τ (hereafter, effective stress). To obtain the 1-loop solution, we can combine the two equations, noting that $\delta_\ell \approx -\Theta$ (we denote this coarse-grained, linear-order overdensity as $\delta_{1\ell}$)

$$(\partial_t^2 + 2H\partial_t - 4\pi G\bar{\rho}) \delta_{1\ell} = \frac{1}{a^2} c_{\text{tot}}^2 \partial_x^2 \delta_{1\ell} + \frac{1}{a^2 \rho_\ell} \partial_x^2 J. \quad (2.55)$$

This equation is then solved perturbatively up to third-order in δ_ℓ , and the solutions are combined to obtain the 1-loop power spectrum. For the coarse-grained version of SPT (hereafter, tSPT for *truncated* SPT), the right-hand side of Eq. (2.55) vanishes. In the EFTofLSS, the right-hand side acts as a source term which captures small-scale physics, chiefly through c_{tot}^2 . Compared to tSPT, the source term introduces a correction at the level of the overdensity as follows

$$\tilde{\delta}_\ell^{\text{corr}}(k, a) = k^2 \int_0^\infty da' G(a, a') \left[a' c_{\text{tot}}^2(a') \tilde{\delta}_{1\ell}(k) - \frac{J(k, a')}{\bar{\rho}} \right], \quad (2.56)$$

where G is the Green's function for the time-dependence of Eq. (2.55). In an EdS universe, it can be expressed as (Hertzberg, 2014; Pajer & Zaldarriaga, 2013)

$$G(a, a') = \mathcal{S}(a - a') \frac{2}{5H_0^2} \left[\left(\frac{a'}{a} \right)^{3/2} - \frac{a}{a'} \right], \quad (2.57)$$

⁴The minus sign in the expression is purely a choice of convention that stems from the definition of θ on page 18. Our sign convention here is the same as in K24, and differs from that of McQuinn & White (2016) who define $c_{\text{tot}}^2 = c_s^2 + \dot{a}c_v^2$.

where $\mathcal{S}$ is the Heaviside step function, which is unity for $a > a'$ and zero otherwise. At any time a , the overdensity field is affected by $c_{\text{tot}}^2(a')$ for all $a' \leq a$. Ignoring – for the moment – the effect of the stochastic term, the EFT power spectrum at next-to-leading order is

$$P_{\text{EFT}}(k, a) = P_{\text{tSPT}}(k, a) + 2\alpha_c(a)k^2\widetilde{W}_\Lambda^2 P_{\text{L}}(k, a), \quad (2.58)$$

where the spatial- and time-dependence of the EFT correction has been separated, with the latter defined as

$$\alpha_c(a) = \frac{1}{a} \int_0^\infty da' G(a, a') a' c_{\text{tot}}^2(a'). \quad (2.59)$$

Let us examine the time-dependence of the correction to the power spectrum. In the linear regime⁵, $c_{\text{tot}}^2 \propto a$, and we can use this to evaluate the scaling of α_c with a . The Green's function introduces an additional factor of a through the integral. Specifically (Pajer & Zaldarriaga, 2013),

$$\int_0^\infty da' G(a, a') (a')^n \propto a^{n+1}, \quad (2.60)$$

and thus $\widetilde{\delta}_\ell^{\text{corr}} \propto a^3$ (ignoring the effect of J). This is a third-order correction in δ_ℓ , and thus enters the power spectrum via its correlation with $\delta_{1\ell}$ to give the second term on the right of Eq. (2.58). As a counterterm, this fixes the UV-divergence of P_{13} in SPT (Baumann et al., 2012; Carrasco et al., 2012; Hertzberg, 2014). Similarly, the stochastic correction cancels the UV-divergence of P_{22} . The final step of the EFTofLSS is to renormalise the loop integrals. This is done by choosing the counterterms at fixed cutoff such that the EFT results agree with observations at a chosen renormalisation scale k_{ren} , then formally setting $\Lambda \rightarrow \infty$ (Carrasco et al., 2012). This removes all cutoff-dependence, as the results of a physical theory should not depend on an arbitrarily-chosen scale. The EFTofLSS has extended the range of applicability, and the accuracy of existing perturbative methods. In the Eulerian setting, it has been applied to predict the matter power spectrum (Carrasco et al., 2012, 2014a,b) and bispectrum (Baldauf et al., 2015; Angulo et al., 2015b). In the former case, the EFTofLSS has extended the k -reach of perturbative models by a factor of at least two (Alkhanishvili et al., 2022). Additionally, inclusion of biased tracers of the matter overdensity has been widely studied (Assassi et al., 2014; Senatore, 2015; Angulo et al., 2015a; Mirbabayi et al., 2015; Donath & Senatore, 2020). Finally, the Lagrangian space equivalent has also seen development in recent years (Porto et al., 2014; Vlah et al., 2015). Despite the theory passing many observational checks, its assumptions have not been fully scrutinised. Specifically, the validity of

⁵This can be inferred from the definition of τ in the linear regime. This time-dependence has often been assumed in the literature (Carrasco et al., 2012, 2014b; Hertzberg, 2014; McQuinn & White, 2016), and we show it to be true in the linear regime in Chapter 4 (and K24).

the effective stress expansion Eq. (2.52), and the time-evolution of the coefficients and the counterterm has not been studied in detail. Our main motivation in K24 was to perform these tests in a simple setting. To lay the final groundwork, note that EFTs can be broadly split into two categories

- If the underlying UV-complete theory is known, one can integrate out the UV degrees of freedom to obtain a ‘top-down’ EFT. Then, one can calculate the Wilson coefficients from the UV-complete theory, without the need for free parameters.
- On the other hand, if the UV-complete theory is unknown (or calculations within it are intractable), an EFT can be constructed in a ‘bottom-up’ way, by expressing the quantity of interest in an operator-coefficient expansion (as in Eq. 2.52) consistent with the symmetries of the physical system. In this case, the Wilson coefficients are free parameters.

In this section, we have described a bottom-up construction of the EFTofLSS, following Baumann et al. (2012); Carrasco et al. (2012). However, the two constructions share the same dynamical equations, and the main point of difference is the way in which the Wilson coefficient (c_{tot}^2) is found. Particularly, N -body simulations may be thought of as the UV-complete theory describing the gravitational evolution of a dark matter sheet in one-dimension. Comparing the EFT power spectrum defined in Eq. (2.58) to the N -body power spectrum offers a top-down estimation of the Wilson coefficient. In practice, only α_c can be estimated in this way, as we cannot undo the causal integral of Eq. (2.59) to get c_{tot}^2 . Additionally, by calculating the effective stress (Eq. 2.52) explicitly from the simulations, we can measure c_{tot}^2 in a bottom-up way, and then numerically perform the integration of Eq. (2.59) to obtain α_c . Thus, we measure α_c from both the bottom-up and the top-down approaches. If the assumptions that the EFTofLSS makes are to be true, these two estimates of α_c must agree with each other. We perform this test in K24, and Chapter 4 offers an overview of our results. In the following section, we move beyond the dark matter field, and consider the theoretical modelling of one of its biased tracers – the dark matter halo.

2.3. Formation and assembly of dark matter halos

In the halo-based approach to gravitational clustering, dark matter is assumed to clump in discrete, gravitationally-bound *halos* that host galaxies and galaxy clusters (Neyman & Scott, 1952; White & Rees, 1978; Scherrer & Bertschinger, 1991; Cooray & Sheth, 2002). If dark matter is cold, they are expected to form hierarchically, with less massive halos accreting and merging to form more massive ones. The distribution and physical properties of galaxies depend on the halos in which they form. In turn, the clustering of the halos is driven primarily by their mass. However,

their environment and formation history also contributes to clustering – an effect known as *halo assembly bias* (Sheth & Tormen, 2004; Gao et al., 2005; Wechsler et al., 2006; Croton et al., 2007). Understanding the mechanisms by which halos gain and lose mass is thus important in understanding the evolution of the galaxies that reside within them. In the context of halos, the clustering of dark matter can be understood in two steps: on large scales by studying the spatial distribution of halos, and on small scales by studying the internal structure and distribution of matter within halos. The former regime can be described well by perturbation theory, and on scales where perturbation theory fails, halos can be assumed to be virialised. Therefore, the key aspects of this approach include recipes for the formation of halos, for their density profiles, and for their spatial distribution. Each of these require simplifying assumptions in order to make analytical predictions. Our focus in this work is the first of the three – an exploration of how collapsed objects are born from the smooth initial conditions of the Universe.

The peaks formalism (e.g. Doroshkevich, 1970; Peebles, 1980; Peacock & Heavens, 1985; Bardeen et al., 1986) associates proto-halo patches which later virialise, to local maxima of the initial overdensity field⁶. On the other hand, the excursion set approach (e.g. Press & Schechter, 1974; Bond et al., 1991) considers regions of the smoothed overdensity field which exceed a certain density threshold, as seeds of halo formation. The simplest prescription for the collapse of density peaks or excursion set regions into bound, virialised structures assumes that the collapse is spherically symmetric (e.g. Lemaître, 1958; Lifshitz & Khalatnikov, 1963; Van Albada, 1960; Gunn & Gott III, 1972; Fosalba & Gaztanaga, 1998).

2.3.1. Gravitational collapse under spherical symmetry

As a simple illustration of spherical collapse, consider an EdS universe. Assuming that the initial fluctuations in the Universe were Gaussian, take a spherical patch of radius r , enclosing matter with mass $M(r)$. The overdensity of this patch is defined as

$$\delta_{\text{SC}}(r, t) = \frac{3M(r)}{4\pi r^3 \bar{\rho}(t)} - 1, \quad (2.61)$$

We would like to trace the evolution of this patch due to the gravitational potential of the mass within it. The acceleration of the patch is given as

$$\frac{d^2 r}{dt^2} = -\frac{GM(r)}{r^2}. \quad (2.62)$$

We may integrate Eq. (2.62) to obtain

$$\left(\frac{dr}{dt}\right)^2 = \frac{2GM(r)}{r} + C, \quad (2.63)$$

⁶Note that N -body simulations suggest the existence of a significant population of *peakless* halos (Ludlow & Porciani, 2011).

which expresses the velocity of the expansion. If the universe has no overdense regions, the integration constant $C = 0$, and the only contribution is from Hubble expansion. A positive value of C introduces an additional contribution from underdensities that speeds up expansion, while a negative value points to overdensities that slow it down. First, let us consider the case of pure Hubble expansion with $C = 0$. We can further integrate Eq. (2.63) to get the evolution of the radius

$$r_{\text{Hubble}}(t) = \frac{1}{2} (GM)^{1/3} (6t)^{2/3} , \quad (2.64)$$

assuming that $r = 0$ at $t = 0$ (we drop the argument of M for brevity). Thus, in the absence of an enclosed overdensity, the patch expands indefinitely, with $r \rightarrow \infty$ as $t \rightarrow \infty$, as expected. Let us now consider the effect of an overdensity via a negative value of C . In this case we may re-express Eq. (2.63) in a parametric form as follows

$$r(\theta) = \frac{GM}{C} (1 - \cos \theta) , \quad (2.65)$$

$$t(\theta) = \frac{GM}{C^{3/2}} (\theta - \sin \theta) . \quad (2.66)$$

The general behaviour of $r(t)$ can now be examined. As θ varies between 0 and 2π , $t(\theta)$ is a monotonically-increasing function. Then $r(\theta)$, and therefore $r(t)$, increases until $\theta = \pi$. This is the maximum radius of the patch – $r_{\text{max}} = 2GM/C$, at $t_{\text{max}} = \pi GM/C^{3/2}$. Thus, $C = 2GM/r_{\text{max}}$. This allows us to express the spherical collapse overdensity (δ_{SC}) in terms of the parameter θ , by noting that the mean background density in an EdS universe is given as $\bar{\rho} = (6\pi G t^2)^{-1}$ (e.g. Sahni & Coles, 1995). Then, substituting the parametric relations in Eq. (2.61), we get

$$\delta_{\text{SC}}(\theta) = \frac{9}{2} \frac{(\theta - \sin \theta)^2}{(1 - \cos \theta)^3} - 1 . \quad (2.67)$$

Thus, the overdensity of the patch at maximum radius is

$$\delta_{\text{SC}}(\theta = \pi) = \frac{9\pi^2}{16} - 1 . \quad (2.68)$$

At this epoch, the overdensity of the region overcomes Hubble expansion, and the patch *turns around* and starts to contract. At $\theta = 2\pi$, the patch would collapse to $r = 0$, at which point $\delta \rightarrow \infty$ as Eq. (2.67) becomes singular. This is unphysical, and in reality, the patch virialises before collapsing. Let us label the radius of the patch at virialisation as r_{vir} . The kinetic energy at the time of virialisation is given as

$$T(t_{\text{vir}}) = -\frac{GM}{2r_{\text{vir}}} . \quad (2.69)$$

At r_{max} , the kinetic and potential energies match. Thus, we can relate the maximum and virial radii as $r_{\text{max}} = 2r_{\text{vir}}$. The overdensity may also be calculated by taking

the ratio of the density at virialisation to the density of the sphere expanding purely due to Hubble expansion. As the mass in the patch is conserved, this is simply the inverse of the ratio of the respective volumes

$$\delta_{\text{SC}}(\theta = 2\pi) = \frac{r_{\text{Hubble}}^3(t_{\text{vir}})}{r_{\text{vir}}^3} - 1 . \quad (2.70)$$

Using Eq. (2.66), we find that $t(\theta = 2\pi) = 2\pi GM/C^{3/2}$. Substituting this in Eq. (2.64), we get

$$r_{\text{Hubble}}(t_{\text{vir}}) = \frac{GM}{C} (18\pi^2)^{1/3} . \quad (2.71)$$

But, $GM/C = r_{\text{max}}/2 = r_{\text{vir}}$. We can now cube both sides and move r_{vir} to the left side, such that

$$\frac{r_{\text{Hubble}}^3(t_{\text{vir}})}{r_{\text{vir}}^3} = 18\pi^2 . \quad (2.72)$$

Then, Eq. (2.70) can be rewritten as

$$\delta_{\text{SC,vir}} = 18\pi^2 - 1 . \quad (2.73)$$

Thus, following spherical collapse, the overdensity of the volume at virialisation is ≈ 178 times its mean density. It is instructive at this point to get a similar estimate using linear perturbation theory. At early times, i.e. for small values of θ , we can approximate Eq. (2.66) by taking the first two terms of the Maclaurin series for $\cos \theta$

$$t \approx \frac{GM(r)}{C^{3/2}} \left(\frac{\theta^3}{6} - \frac{\theta^5}{120} \right) . \quad (2.74)$$

Thus,

$$\begin{aligned} \theta^3 &\approx \frac{6tC^{3/2}}{GM} \left(1 - \frac{\theta^2}{20} \right)^{-1} \\ &\approx \frac{6tC^{3/2}}{GM} , \end{aligned} \quad (2.75)$$

where we used $\theta^2 \ll 1$ in the last step. We can similarly approximate r as follows

$$\begin{aligned} r &\approx \frac{GM}{C} \theta^2 \left(1 - \frac{\theta^2}{12} \right) \\ &\approx \frac{GM}{C} \left(\frac{6tC^{3/2}}{GM} \right)^{2/3} \left[\left(1 - \frac{\theta^2}{20} \right)^{-2/3} \left(1 - \frac{\theta^2}{12} \right) \right] \\ &\approx (GM)^{1/3} (6t)^{2/3} \left[1 - \frac{1}{20} \left(\frac{6tC^{3/2}}{GM} \right)^{2/3} \right] , \end{aligned} \quad (2.76)$$

using the binomial theorem and substituting Eq. (2.75) in the last step. Note that this result agrees with Eq. (2.64) in the limit of $C \rightarrow 0$. Thus, the first term accounts for Hubble expansion, and the second term is the effect of the overdensity using linear perturbation theory. To calculate the overdensity at virialisation, we must evaluate the following expression

$$\delta_{\text{SC}}(t_{\text{vir}}) = \left[\frac{r(t=0)}{r(t_{\text{vir}})} - 1 \right]^3. \quad (2.77)$$

Linearising the equation, and using the relation between r and t (Eq. 2.76),

$$\begin{aligned} \delta_{\text{SC,L}}(t_{\text{vir}}) &\approx 3 \left[\frac{r(t=0)}{r(t_{\text{vir}})} - 1 \right] \\ &\approx 3 \left\{ \left[1 - \frac{1}{20} (12\pi)^{2/3} \right]^{-1} - 1 \right\} \\ &\approx 1.686. \end{aligned} \quad (2.78)$$

where we used the expression $t_{\text{vir}} = 2\pi GM/C^{3/2}$ in the second step. Thus, the linear overdensity of a spherically symmetric region extrapolated to the time of virialisation is ≈ 1.686 . Although this prediction is inaccurate compared to the non-linear estimate from Eq. (2.73), it provides us with a tool to determine approximately when a spherically-symmetric region will undergo collapse and virialise. Given some initial overdensity of a spherical region, collapse occurs at the scale factor when its linearly extrapolated overdensity first exceeds 1.686. Since the linear overdensity scales as $\delta_{\text{L}} \propto a$, we can express this collapse time as follows

$$a_{\text{coll}} = \frac{1.686 \cdot a_{\text{in}}}{\delta_{\text{L}}(a_{\text{in}})}. \quad (2.79)$$

This result is the basis of the excursion-set approach we describe in the following section.

2.3.2. Excursion-set theory

Given the analytical prediction of the collapse time, we can construct the accretion histories of dark matter halos. The principle behind this construction is the following – regions of the linear overdensity field smoothed on a mass scale M denser than a threshold τ contain a collapsed structure of mass M . These regions are called excursion sets, and may be formally defined as follows

$$\mathbb{S}_\tau = \{\mathbf{x} \mid \delta_{\text{L}}(\mathbf{x}, M) \geq \tau\}. \quad (2.80)$$

The theory of excursion sets defines rules for assigning mass to virialised objects at certain scales. To extract spatial information from the overdensity field, we will

consider the field filtered at a number of different scales. This smoothing works exactly as in Eq. (2.45), with many possible choices for the window function. If we are to be completely rigorous, the following derivations in excursion-set theory are only valid for a sharp filter in k -space. On the other hand, the conversion of filter size into volume, has ambiguity for all filters except the sharp filter in x -space. The literature on the subject uses the k -space filter for the derivation, and the x -space filter for the conversion. Although this inconsistency is untenable for a formal treatment and extension of excursion-set theory (as in Maggiore & Riotto, 2010a), it may be glossed over for our purposes.

The choice of the window function determines the volume of the smoothed region. For a spherical top-hat window of radius r , the volume is $V_{\text{top-hat}} = 4\pi r^3/3$, while for a Gaussian window, $V_{\text{Gaussian}} = (2\pi)^{3/2} r^3$. Using this volume, we may relate r to a mass scale as $M(r) = \bar{\rho} V(r)$. To proceed concretely, we will focus on the overdensity field at an arbitrary point in space, and as such, suppress the corresponding argument in δ .

As in the original application of spherical collapse to dark matter halos (Press & Schechter, 1974), let us start by identifying the probability distribution of the overdensity field as a Gaussian,

$$\mathcal{P}(\delta_L, S) = \frac{1}{\sqrt{2\pi S}} e^{-\delta_L^2/(2S)} . \quad (2.81)$$

where $S(r) = \sigma^2(r)$ is the variance of the overdensity defined as follows

$$S(r) = \frac{1}{2\pi^2} \int dk \, k^2 P(k) |\widetilde{W}(k, r)|^2 , \quad (2.82)$$

where $P(k)$ is the linear power spectrum and $\widetilde{W}$ is the Fourier-space window. Note that S inherits a redshift-dependence from the power spectrum, which we have suppressed here. Then, the volume fraction occupied by collapsed objects with size larger than r – equivalently, with mass greater than $M(r)$ – is given by integrating over the tail of the above probability distribution. Denoting the collapse threshold by δ_c , this fraction can be calculated as

$$\begin{aligned} F_{\text{PS}}(\delta_c, S) &= \int_{\delta_c}^{\infty} d\delta_L \, \mathcal{P}(\delta_L, S) \\ &= \frac{1}{2} \text{erfc} \left[\frac{\delta_c}{\sqrt{2S}} \right] , \end{aligned} \quad (2.83)$$

where ‘erfc’ is the complementary error function. However, this result is incorrect. As $M \rightarrow 0$, the variance diverges, and we expect the fraction to be unity, but Eq. (2.83) implies $F(\delta_c, 0) = 1/2$. This problem is pointed out in Press & Schechter (1974), who subsequently add a factor of two by hand to get a consistent volume fraction. The underlying reason for this discrepancy is the so-called ‘cloud-in-cloud’

problem. The Press-Schechter (PS) formalism defines all sufficiently overdense regions as halos. At a given smoothing scale, an overdense patch may contain nested regions which are also overdense. Similarly, the patch may itself be nested within a larger region. The PS formalism has no mechanism to properly count these nested regions, leading to an underestimation of high-mass halos, and an overestimation of low-mass halos.

To solve this, Bond et al. (1991) proposed a stochastic approach which leverages excursion sets. Consider the behaviour of the variance as a function of the smoothing scale. For very large smoothing ($r \rightarrow \infty$), the variance (and δ_L) tends to zero. The probability that the corresponding excursion set lies above the threshold is low. As we reduce r , $S(r)$ monotonically increases, and $\delta_L(r)$ is likelier to cross δ_c . The evolution of $\delta_L(r)$ is stochastic, and is governed by the Langevin equation where r acts as a time variable (see e.g. Maggiore & Riotto, 2010a). The solution is an uncorrelated⁷ random walk with respect to r . A halo forms when this trajectory crosses the threshold (or barrier) $\delta_L = \delta_c$.

Within this framework, the cloud-in-cloud problem is represented by trajectories of $\delta_L(r)$ that cross the barrier multiple times. To fix the problem, Bond et al. (1991) count only the first crossing – the one at largest r . For a sharp- k filter, the probability distribution obeys the Fokker-Planck equation (Bond et al., 1991; Maggiore & Riotto, 2010a)

$$\frac{\partial \mathcal{P}}{\partial S} = \frac{1}{2} \frac{\partial^2 \mathcal{P}}{\partial \delta_L^2} . \quad (2.84)$$

The PS formalism can be viewed as solving Eq. (2.84) with boundary conditions such that the probability vanishes at $\delta_L = \pm\infty$, recovering Eq. (2.81). To include the halos neglected due to the cloud-in-cloud problem, Bond et al. (1991) set the following boundary condition

$$\mathcal{P}(\delta_L, S)|_{\delta_L=\delta_c} = 0 . \quad (2.85)$$

This leads to the following solution to Eq. (2.84) (see e.g. Zentner, 2007; Maggiore & Riotto, 2010a for details)

$$\mathcal{P}(\delta_L, S) = \frac{1}{\sqrt{2\pi S}} \left[e^{-\delta_L^2/(2S)} - e^{-2(\delta_c - \delta_L)^2/(2S)} \right] . \quad (2.86)$$

Then, the volume fraction of collapsed objects is given as⁸

$$F_{\text{SC}}(\delta_c, S) = 1 - \int_{-\infty}^{\delta_c} d\delta_L \mathcal{P}(\delta_L, S) . \quad (2.87)$$

⁷Here, it is instructive to note the effect of the choice of smoothing, and put our earlier discussion into context: the random walk trajectories are uncorrelated *only* for a sharp k -space filter. For any other filter choice, we must solve the much harder problem of correlated random walks (Maggiore & Riotto, 2010a,b).

⁸We cannot simply integrate the probability from δ_c to ∞ as before, because $\mathcal{P}(\delta_L, S)$ is only defined for $\delta_L < \delta_c$. Instead, we find the fraction of objects that never cross the barrier, and subtract it from unity (Maggiore & Riotto, 2010a).

Substituting the solution for $\mathcal{P}$ from Eq. (2.86) into Eq. (2.87), we get

$$F_{\text{SC}}(\delta_c, S) = \text{erfc} \left[\frac{\delta_c}{\sqrt{2S}} \right], \quad (2.88)$$

which is the expected result. The probability that a random-walk first crosses the collapse barrier between S and $S + dS$ is given by the first-crossing rate, defined as

$$\frac{dF_{\text{SC}}}{dS} = \frac{\delta_c}{\sqrt{2\pi S^3}} e^{-\delta_c^2/(2S)}. \quad (2.89)$$

We can use the first-crossing rate to find the halo mass function, which is a statistic that counts the distribution of halos in mass bins. It has been shown (e.g. in Jenkins et al., 2001; Zentner, 2007) that the mass function derived from spherical collapse is qualitatively similar to that measured from N -body simulations. However, there is a quantitative mismatch between the two. We will demonstrate this discrepancy in Sec. 2.3.3, where the spherical collapse predictions are labelled as ‘SC’. The reasons for this discrepancy are twofold. First, the physical model of spherical collapse is an oversimplification, and second, the mathematical framework of excursion-set theory has inconsistencies. For the purposes of this work, we will gloss over the subtleties of the latter, which are addressed in Maggiore & Riotto (2010b), and instead turn our attention to the first problem. In the real Universe, halo formation is not spherically symmetric. Indeed, simulations show that the collapse is tri-axial, and is accompanied by a retinue of complex dynamical processes (i.e., mergers, fragmentation, relaxation; see e.g. Springel et al., 2021) which the simple spherical collapse model fails to account for. To make the analytical prediction more realistic, the above framework may be generalised to the collapse of a tri-axial or ellipsoidal patch, which we discuss next.

Ellipsoidal collapse

Viewed within the framework of excursion-set theory, ellipsoidal collapse can be posed as the problem of random walks crossing a barrier of non-constant height (Sheth et al., 2001), or a diffusing barrier (Maggiore & Riotto, 2010b). The shape of this barrier depends on the variance σ^2 of the overdensity field. Sheth et al. (2001) proposed the following expression for the barrier shape

$$B_{\text{EC}}(\sigma) = \delta_c \left[1 + \beta \left(\frac{\delta_c}{\sigma} \right)^{-2\gamma} \right]. \quad (2.90)$$

where β and γ are parameters originating from ellipsoidal collapse dynamics, and may be calibrated against numerical data. Note that the barrier shape carries a cosmology dependence through σ , which we have suppressed with the implicit assumption of $a = 1$. To calculate the barrier at any other time, we must multiply σ^2

by $D_+^2(a)$, where D_+ is the linear growth factor. For this work, we find that the parameter values from Sheth & Tormen (1999) to work well with our simulation data, and thus pick $\beta = 0.47$, and $\alpha = 0.615$. Given this barrier shape, we can use the δ_L trajectory of halos to compute their mass accretion history. We will enumerate an algorithm for this in Sec. 2.3.4, where the ellipsoidal collapse predictions are labelled as ‘EC’.

For a given proto-halo, the departure from spherical symmetry may be expressed in terms of the tidal field, characterised by the linear tidal deformation tensor

$$\omega_{ij} = \frac{\partial^2 \phi}{\partial x_i \partial x_j} , \quad (2.91)$$

where ϕ is the peculiar gravitational potential which solves the Poisson equation for δ_L . The deformation tensor at a proto-halo’s center of mass characterises its neighbouring tidal field. Then, we can consider the eigenvalues of the deformation tensor, whose relative sizes reflect the degree to which a structure has collapsed along a specific direction. Let us label the eigenvalues as λ_i , organised (without loss of generality) such that, $\lambda_1 \geq \lambda_2 \geq \lambda_3$. It is straightforward to see that the trace of the deformation tensor is $\lambda_1 + \lambda_2 + \lambda_3 = \delta_L$. Further, we can define two quantities that capture the shape of the halo viz., the ellipticity (e) and prolateness (p)

$$e = \frac{\lambda_1 - \lambda_3}{2} , \quad p = \frac{\lambda_1 - 2\lambda_2 + \lambda_3}{2\delta_L} . \quad (2.92)$$

The set of variables (e, p, δ_L) contains the same information as the three eigenvalues, but carry more physical intuition. Following this thread, Borzyszkowski et al. (2017a) defined a criterion for halo accretion that considers the ellipticity as an input. Specifically, they observed that halos with $e < 0.5$ were more likely to accrete mass. This suggested the following modification to the spherical collapse mass – if e (measured at the mass scale of the halo) is ≥ 0.5 , the halo stops accreting, even if the spherical collapse mass grows. To understand the reasoning behind this criterion, let us rephrase this in terms of the eigenvalues.

$$e < 0.5 \quad \Rightarrow \quad \lambda_2 > -2\lambda_3 . \quad (2.93)$$

Thus, whenever $\lambda_2 > 0$ and $\lambda_3 < 0$, accretion takes place only when the magnitude of the latter is not too large. Further, a negative λ_2 always suppresses accretion. Let us remark that there is no a priori reason for the threshold to be 0.5. For our primary simulation (see Sec. 3.2.2), we lower this threshold to 0.4. We label the resulting predictions as ‘Bo17’, and we shall discuss them in Chapter 5.

At this point, it is worth enquiring which property best quantifies the impact of halo environment on its growth. We will attempt to answer this question in Chapter 5, and it is useful to define a few quantities that capture certain aspects

of the halo environment. First, we may define the three rotational invariants of the deformation tensor

$$I_1 = \lambda_1 + \lambda_2 + \lambda_3 , \quad (2.94)$$

$$I_2 = \frac{1}{2}(\lambda_1\lambda_2 + \lambda_2\lambda_3 + \lambda_1\lambda_3) , \quad (2.95)$$

$$I_3 = \lambda_1\lambda_2\lambda_3 . \quad (2.96)$$

The first invariant is simply δ_L , while the third is the determinant of the deformation tensor. This set of invariants is degenerate with the three eigenvalues. However, the combination (Catelan & Theuns, 1996; Paranjape et al., 2018)

$$q = I_1^2 - 3I_2 \quad (2.97)$$

quantifies the degree of anisotropy at the smoothing scale and has a relevant, desirable property. Namely, for a Gaussian random field, the distribution of q is independent of δ_L . In the non-Gaussian case, however, q correlates strongly with δ_L . To mitigate this, Paranjape et al. (2018) defined a new quantity that contains information complementary to δ_L even for a non-Gaussian field. This tidal anisotropy may be defined as

$$\alpha = \frac{\sqrt{q}}{1 + \delta_L} . \quad (2.98)$$

In Chapter 5, we use α as an input for our neural network in order to study its impact on halo formation.

2.3.3. Halo mass function

In order to compute statistics of the halo distribution within the excursion-set, ellipsoidal collapse framework, Sheth et al. (2001) suggest the following expression for the first-crossing rate

$$\frac{dF_{\text{EC}}}{dS} = A \left[1 + \left(\frac{S}{\delta_c^2} \right)^q \right] \left(\frac{\delta_c}{\sqrt{2\pi S^2}} \right) e^{-\delta_c^2/(2S)} , \quad (2.99)$$

where q and A may be set by calibrating with simulations. This is calculated using Monte-Carlo simulations of the Langevin equation for δ_L . We will use this result in this section to predict the halo mass function (HMF). In its differential form, the HMF denotes the average number density of halos in a given mass bin. Following from the above discussion, we assume that the initial density field obeys Gaussian statistics. Then, the differential mass function is strongly related to the probability of uncorrelated, Brownian random walks to cross a barrier δ_c . To arrive at an expression for this quantity, let us consider the first-crossing rate as a function of the mass M , instead of S . The fraction of volume occupied by collapsed objects

between mass M and $M + \Delta M$ is given as $|dF/dM|dM$. Then, the average number density of halos is

$$\begin{aligned} \frac{dn}{dM} dM &= \frac{\bar{\rho}}{M} \left| \frac{dF}{dM} \right| dM \\ &= \frac{\bar{\rho}}{M} \frac{dF}{dS} \left| \frac{dS}{dM} \right| dM . \end{aligned} \quad (2.100)$$

Now, we can substitute the first-crossing rates from our two models – spherical and ellipsoidal collapse – to obtain their analytical HMF predictions. Restoring our original variable σ , we find

$$\left(\frac{dn}{dM} \right)_{\text{SC}} = \sqrt{\frac{2}{\pi}} \frac{\bar{\rho}}{M} \frac{\delta_c}{\sigma^2} e^{-\delta_c^2/(2\sigma^2)} \left| \frac{d\sigma}{dM} \right| , \quad (2.101)$$

$$\left(\frac{dn}{dM} \right)_{\text{EC}} = A \left[1 + \left(\frac{\sigma}{\delta_c} \right)^{2q} \right] \sqrt{\frac{2}{\pi}} \frac{\bar{\rho}}{M} \frac{\delta_c}{\sigma^2} e^{-\delta_c^2/(2\sigma^2)} \left| \frac{d\sigma}{dM} \right| . \quad (2.102)$$

At the level of the HMF, the two models differ by the prefactor that is defined by the parameters A and q . Now, we can compare the theoretical mass functions to that obtained from simulations, and this is shown in Fig. 2.1. We define logarithmic mass bins in the range $10^{10} - 10^{16} h^{-1} M_\odot$, which is covered by two simulations with boxes of side length 100 and 300 h^{-1} Mpc (**box100** and **box300** respectively). At the lower end of the mass range, **box100** is limited by its mass resolution, and the measured mass function starts to drop off around $10^{10} h^{-1} M_\odot$. Similarly, **box300** contains no halos more massive than a few times $10^{15} h^{-1} M_\odot$, and the HMF sharply drops off in that limit. In the intermediate range, the spherical collapse prediction is qualitatively close to the N -body measurement, but there is a clear quantitative disagreement. Spherical collapse overpredicts the mass function for masses $< 10^{13} h^{-1} M_\odot$, and slightly underpredicts it for masses $> 10^{14} h^{-1} M_\odot$ (note that only about 5% of **box300** halos have masses $> 10^{14} h^{-1} M_\odot$). This behaviour is well-known in the literature, often plotted as the first-crossing rate against the peak height $-\delta_c/\sigma$ (e.g. Sheth & Tormen, 1999; Zentner, 2007; Robertson et al., 2009). In the $10^{13} - 10^{14} h^{-1} M_\odot$ mass range, SC predictions agree well with the simulation. We shall have more to say on this matter in Sec. 5.5. The ellipsoidal collapse predictions are a good match to the simulations over most of the mass range, with the parameter values $A \approx 0.27$ and $q \approx -0.25$. We note that for masses $> 5 \times 10^{14} h^{-1} M_\odot$, ellipsoidal collapse also slightly underpredicts the mass function, though there are only a handful of halos with this mass. It is important to point out that we have used functions with two distinct sets of parameters for the barrier shape and the HMF. In principle, it is possible to find the HMF for a given barrier shape without running Monte-Carlo simulations. This was demonstrated in Zhang & Hui (2006), who find that the first-crossing distribution of a given barrier shape satisfies an integral




Figure 2.1. – Comparison of excursion-set predictions of the HMF with that obtained from N -body simulations. The spherical collapse prediction (SC; dashed teal line) differs from the simulation HMF through most of the mass range, while the ellipsoidal collapse prediction (EC; dot-dashed dark blue line) is a much better match throughout, except for the highest masses.

equation that may be solved by matrix inversion. Robertson et al. (2009) show that an HMF thus derived, is inconsistent with that measured from the simulation. This points to a failure of the excursion set model that, nonetheless, furnishes reasonable predictions. Maggiore & Riotto (2010a,b) attempt to reformulate excursion-set theory using path integrals. The collapse barrier is treated as a stochastic variable (see also Corasaniti & Aчитouv, 2011a,b), and inconsistencies in filter choice are dealt with by introducing non-Markovian corrections to the random walk. However, a discussion of this and other such extensions are beyond our scope, and we leave the curious reader to refer to the literature for these details.

2.3.4. Predictions for individual halos

Let us now sketch an algorithm for generating our analytical predictions of halo accretion histories via the three models – SC, EC, and Bo17. To discuss this concretely, we pick an arbitrary halo from **box100**. We start with the overdensity field filtered at a set of spatial scales, and proceed as follows

1. Given the positions of the halo particles, we find the proto-halo patch that will collapse to form the halo at $a = 1$, and locate its center of mass.
2. We obtain the trajectory of δ_L vs M for a given scale factor a . Here, δ_L is the overdensity at the center of mass of the proto-halo patch, extrapolated to scale factor a using linear theory. For our example halo, three selected trajectories are shown in the left panel of Fig. 2.2 (solid lines).
3. We compare this trajectory against the two barriers – the constant SC barrier at $\delta_c = 1.686$ (dashed teal line in Fig. 2.2), and the EC barrier defined by




Figure 2.2. – The analytical prediction for the accretion history of an example halo. *Left:* δ_L vs M trajectories at three different times. We show the SC and EC barriers. The SC barrier is constant (dashed grey line), while its EC counterpart changes with time (dot-dashed lines; colours correspond to different scale factors). The SC and EC up-crossings are marked by discs and stars respectively. *Right:* The accretion histories predicted by the three models, and the true history obtained from the simulation (box100).

Eq. (2.90) (dot-dashed lines in Fig. 2.2, with colour corresponding to a). Note that as $a \rightarrow 0$, the EC barrier converges to δ_c . At each a , we identify the first up-crossing (in order of decreasing mass) of δ_L over the barriers. These are marked by discs (stars) for the SC (EC) model, where the colour refers to a as before. The x -intercept of this up-crossing gives a mass scale. This is the largest mass at which δ_L exceeds the barrier height. This becomes the SC (or EC) mass of the proto-halo at that time – $M_p(a)$. If there are no up-crossings, the mass remains unchanged from its value at the previous scale factor (or is zero if no up-crossings have taken place thus far).

4. Given the SC mass, we find the ellipticity of the halo at that mass scale. If $e < 0.4$, the ellipsoidal collapse mass is same as the spherical collapse mass. Instead, if $e \geq 0.4$, the halo is assumed to stop accreting, and its Bo17 mass is set equal to the mass at the previous scale factor.
5. We repeat steps 2–4 for an array of scale factors, and obtain the mass of the proto-halo at each a . The compiled set of masses constitutes the mass accretion history of the halo predicted by excursion-set theory. We compare these predictions with the true accretion history for our example halo on the right side of Fig. 2.2.

Calculating the accretion histories of all halos in the population allows us to evaluate derived properties, e.g. the halo formation time. We use these results to set benchmarks for our neural network predictions in Chapter 5.

CHAPTER 3

Numerical Methods

3.1. N -body simulations

The theoretical methods we described in Chapter 2 to provide an accurate description of structure formation in the linear and quasi-linear regimes, and an approximate model for halo formation. However, to obtain an accurate picture of structure formation in the non-linear regime, we must turn to numerical simulations. At their heart, these simulations solve the gravitational N -body problem – ‘how does a set of N particles move under the influence of its collective gravitational potential?’. As discussed in Chapter 2, we work under the assumption of Newtonian gravity. The issue of relativistic effects in cosmological simulations has been widely discussed (e.g. Chisari & Zaldarriaga, 2011; Green & Wald, 2012; Flender & Schwarz, 2012; Adamek et al., 2013). Particularly, Fidler et al. (2015) show that Newtonian simulations can be interpreted consistently in a so-called N -body gauge (see also Hahn & Paranjape, 2016 for a discussion on this issue). As a result of this, and due to their numerical complexities, relativistic N -body simulations are rarely carried out. A slight exception to this is **GEVOLUTION** (Adamek et al., 2016), which works in the weak-field expansion of GR, and its more modern descendent **GRGADGET** (Quintana-Miranda et al., 2023). There have also been attempts to expand the FLRW metric to obtain approximate solutions, dubbed *post-Friedmann* corrections (Bruni et al., 2014; Milillo et al., 2015; Giblin Jr et al., 2016), and to add relativistic effects in post-processing (e.g. Borzyszkowski et al., 2017b).

The distribution of collisionless DM particles is too large to run a 1:1 simulation. We make the assumption that many physical DM particles make up a *macroparticle*. An ensemble of these macroparticles of mass m is then evolved under Newtonian

gravity. This constitutes an approximate solution to the Vlasov-Poisson system of the form^{1,2}

$$f_{N\text{-body}}(\mathbf{x}, \mathbf{p}, t) = A \sum_{i=1}^N \delta^D[\mathbf{x} - \mathbf{x}_i(t)] \delta^D[\mathbf{p} - \mathbf{p}_i(t)] , \quad (3.1)$$

where the normalisation A determines if f is a number or mass density in phase-space. Chiefly, we are concerned with the evolution of three quantities – the gravitational potential ϕ , the particle positions $\mathbf{x}$ (from which the overdensity δ is calculated), and the conjugate momenta $\mathbf{p}$. The governing equations are

$$\dot{\mathbf{p}}(t) = m a^2(t) \dot{\mathbf{x}}(t) , \quad (3.2)$$

$$\dot{\mathbf{p}}(t) = -m \nabla \phi(\mathbf{x}, t) , \quad (3.3)$$

$$\nabla^2 \phi(\mathbf{x}, t) = 4\pi G \bar{\rho} a^2(t) \delta(\mathbf{x}, t) . \quad (3.4)$$

First, the three fields are calculated at some initial time³ t_i . Then, an iterative algorithm updates them as follows

- Eq. (3.2) is integrated to update the positions,
- the momenta are updated using Eq. (3.3),
- a mass-assignment scheme uses the updated positions to obtain an updated overdensity, and
- the Poisson equation (Eq. 3.4) is solved to update the potential.

This process translates to three key ingredients – (i) a time integrator for updating positions and momenta, (ii) a mass-assignment scheme, and (iii) a Poisson solver. A Newtonian N -body simulation is thus a synthesis of these pieces that moves the ensemble of macroparticles forward in time (typically to $z = 0$). The results of a simulation are stored in *snapshots* spaced out evenly in logarithmic redshifts (this yields more snapshots at late times as high-density regions form). Each snapshot contains the full distribution of particles in phase-space, and is thus expensive to store, especially for large simulations. From these snapshots, derived data products such as summary statistics and halo catalogues (see Sec. 3.1.5) can be calculated

¹This is the discrete *one-particle* Klimontovich distribution (Klimontovich, 1967), as required for collisionless particles. Multi-body interactions are captured by higher-orders in the BBGKY hierarchy, and are ignored here. See e.g. Binney & Tremaine (1987), and App. A.

²One subtlety that we have ignored here is that the N -body distribution must also involve a sum over all periodic copies of the macroparticles. Further, we have assumed the macroparticle mass to be identical. For the exact expression, refer to Angulo & Hahn (2022).

³Typically, this is achieved using LPT truncated to some order. This introduces a truncation error, which can be alleviated by setting an early initial time. However, that makes the simulation prone to numerical errors at late times. A discussion on the balance between these errors, and on the art of picking the initial conditions is found in Michaux et al. (2021).

at relatively low computational cost. Simulations are generally carried out with periodic boundary conditions, to mimic the fact that the Universe contains matter on scales larger than the box size. The mean density of the simulation volume is the same as the mean density of the Universe. This naturally gives rise to a few key parameters that depend on the physical system one wants to study. The box size and number of particles in a simulation set limits on the size of structures that will form in it at late times. Together, these two parameters also set the *mass resolution* of a simulation. For example, a simulation with a cubic box of side length $100 h^{-1} \text{ Mpc}$ with 512^3 particles has a mass resolution of $\approx 6 \times 10^8 h^{-1} \text{ M}_\odot$. In this simulation, dwarf halos with mass $10^8 - 10^9 h^{-1} \text{ M}_\odot$ will not be well-resolved, but Milky Way sized halos will be resolved with $\approx 10^2 - 10^3$ particles. Periodic boundaries are also a natural choice for the Fourier-space methods for evaluating the fields that we will discuss in the following sections.

An important numerical consideration concerns the fact that the gravitational potential between two particles scales inversely with their separation r . This expression diverges as $r \rightarrow 0$. For real particles, dissipative forces prevent any singular behaviour, but two-body interactions may survive as numerical artefacts. Given that we are dealing with collisionless particles (see footnote 1 on page 42), this is undesirable, and motivates the use of a *softening* procedure to remove these artefacts. The key step is the introduction of a softening scale ϵ , such that the gravitational potential is redefined to $\phi \propto (r + \epsilon)^{-1}$. For small separations ($r \ll \epsilon$), the potential is constant, while the large-separation ($r \gg \epsilon$) potential is unaffected. Softening treats a point-mass as an extended object (i.e., the Dirac delta is replaced by a continuous mass distribution), meaning that the small-scale density is smoothed. This works well in concert with the macroparticle assumption, as we are ignorant to these small-scales in any case. However, softening does introduce a bias in the potential calculation – a trade-off against inaccurate small-scale accelerations. We shall have more to say on softening in Sec. 3.1.3, but the curious reader may also refer to, e.g. Dehnen (2001); Barnes (2012).

In this section, we set up the necessary language for discussing numerical simulations. We will briefly describe the development of N -body codes through the years (Sec. 3.1.1), explain how Newtonian simulations broadly work by detailing the three key ingredients (Sec. 3.1.2–3.1.4), summarise common structure finding algorithms (Sec. 3.1.5), and review some miscellaneous numerical issues (Sec. 3.1.6). In Sec. 3.2, we will describe the simulations that were used in this work. Before we begin, we point the reader to some review literature on the subject – an early review by Sellwood (1987), the textbook by Hockney & Eastwood (1988), and the modern reviews by Kuhlen et al. (2012) and Angulo & Hahn (2022).

3.1.1. Historical note

The first numerical simulation of gravitational bodies was performed using light-bulbs as proxies for galaxies, and light as a proxy for mass (Holmberg, 1941). The first simulations on a digital computer were performed in the 1960s, and were limited to around 100 particles (Von Hoerner, 1960, 1963; Aarseth & Hoyle, 1963). One of the first implementations in a realistic cosmological setting was by Peebles (1970), who simulated 300 collisionless particles to model dynamics in the Coma cluster and found the results to be consistent with observations. In the half-century since, the progress of N -body simulations can be tracked through the evolution of gravity (Poisson) solvers and their computational complexity. In the early iterations of the code, force calculations were performed for all pairs of particles, regardless of their distance from each other. This method – *direct summation* – set a limit on the number of particles that could be efficiently simulated. For example, Fall (1978); Aarseth et al. (1979) simulated around 10^3 particles, but found inconsistent results, chiefly due to having too few particles to resolve galaxy clustering. This problem was widely discussed in the contemporary literature (e.g., Efstathiou & Eastwood, 1981; Klypin & Shandarin, 1983; Efstathiou et al., 1985). In particular, Efstathiou & Eastwood (1981) introduced the particle-mesh (PM; see also Doroshkevich et al., 1980) and the particle-particle/particle-mesh (P³M) schemes in cosmological simulations⁴. In a PM code, particles are represented by points on a fixed mesh, allowing for faster estimations of the potential. This is fairly accurate on large-scales, but loses accuracy when the mesh resolution becomes low compared to the particle separation. The discrete gravitational potential at a particle position is a sum over the contributions due to all other particles, and over the periodic copies of the finite simulation box. This is a double sum which converges slowly. To overcome this, the P³M scheme splits the discrete potential into a periodic, long-range (lr) contribution and a non-periodic, short short-range (sr) contribution⁵

$$\phi(\mathbf{x}) = \phi_{\text{lr}}(\mathbf{x}) + \phi_{\text{sr}}(\mathbf{x}) \quad (3.5)$$

where ϕ_{lr} is calculated in Fourier space where it converges quickly using a PM code, and ϕ_{sr} is evaluated in real space by direct summation. This compromise allowed these early simulations to use up to 10^5 particles. However, as time evolution creates high density regions, the direct summation calculations in P³M become much slower.

This problem was addressed with the introduction of the **tree** algorithm (Appel, 1985; Jernigan, 1985), improved upon by Barnes & Hut (1986). The underlying idea is to approximate the force due a group of distant particles by treating them as a single particle located at the group’s center of mass. This is achieved by dividing the simulation volume into a fixed number of subvolumes and creating a

⁴This had already been applied to collisionless plasma by Hockney et al. (1974).

⁵This is called *Ewald summation*, first introduced to calculate the electrostatic energy of ions (Ewald, 1921). It was first applied in cosmology by Bouchet & Hernquist (1988).

hierarchical tree. In three (two) dimensions, each subdivision partitions the volume in eight (four) subvolumes, and the resulting tree is called an *octree* (*quadtree*). The *root* node of the tree contains the full volume and all simulation particles. Following the branches down the tree (see, e.g., the right side of Fig. 3.1), subsequent *internal* nodes contain smaller groups of particles. This recursive process stops when a node contains a single particle (this is called an *external* or *leaf* node).

On average, the algorithm has a complexity of $\mathcal{O}(N \log N)$, which is faster than the $\mathcal{O}(N^2)$ complexity of direct summation. Further, the structure of the tree lends itself naturally to parallelisation, leading to additional speed-ups. On the other hand, in contrast to grid-based methods, the speed of the algorithm depends on the distribution of particles, with high-density regions requiring more direct summations which can slow it down. Modern *N*-body codes tend to overcome this by allowing a node to have more particles (typically $10^1 - 10^3$ depending on the total number of particles in the simulation), and using a hybrid method where the long-range forces are calculated via **PM**, and short-range forces use the **tree** algorithm (Xu, 1995; Bode et al., 2000; Bagla, 2002). This **tree-PM** method was successfully implemented in⁶ **GADGET-1** (Springel et al., 2001) and its successor **GADGET-2** (Springel, 2005).

Although **tree-PM** provides an improvement in computational complexity, it has some limitations. Mainly, the gravitational interactions in the tree are not symmetric, which breaks momentum conservation. Fast multiple methods (**FMM**) are a solution to this, and they come with another improvement in complexity. The basic idea follows the construction of an adaptive hierarchical tree. The interactions between nodes are expressed as a multipole expansion, while the interactions within nodes are expressed as a *local* expansion (usually a polynomial expansion or a Taylor series). The key idea is that particles within the same node feel roughly the same force due to other nodes, and thus the node-node forces can be translated into particle-node forces. This rids us of many redundant calculations from the **tree** algorithm, leading to a remarkable $\mathcal{O}(N)$ complexity. The **FMM** algorithm was first proposed by Greengard & Rokhlin (1987), with the expansions in spherical harmonics. The Cartesian expression was described in Dehnen (2000, 2002). **FMM** is algorithmically more complex, but it increasingly being used in *N*-body simulations, e.g. **GADGET-4** (Springel et al., 2021), **SWIFT** (Schaller et al., 2024), **PKDGRAV-3** (which uses a binary tree instead of an octree; Potter et al., 2017). **GADGET-4**, in particular, implements the **tree**, **tree-PM**, **FMM**, and **FMM-PM** algorithms as alternatives for the gravity solver. We briefly sketch the **FMM** algorithm in Sec. 3.1.3, but a detailed description is beyond our scope. These details may be found in Kurzak & Pettitt (2006); Dehnen & Read (2011); Gnedin (2019).

An alternative way to address the slowness of **P³M** codes in high-density regions is to start with a coarse grid, and successively refine it in the short-range regime,

⁶To be precise, although **GADGET** and its successors can be run in dark matter only mode, they are *hydrodynamical* codes that implicitly include the effects of baryons. A discussion of simulations that include baryonic physics may be found in Vogelsberger et al. (2020).

creating a series of *refinements*. Each refinement is then solved with P^3M scheme as an independent set of particles. This results in an *adaptive* P^3M (AP^3M) code, first introduced by Couchman (1991). Analogously, an octree can also be generated for the simulation box adaptively, with higher levels of refinement for higher-density regions. An N -body code that utilises this method is called an adaptive refinement tree (**ART**), and was first introduced in cosmology by Kravtsov et al. (1997). **ART** is the backbone of the hydrodynamical simulation **RAMSES** (Teyssier, 2002).

In this thesis, we will use simulations run using a development version of **GADGET-3**, first discussed in (Springel et al., 2009). The documentation on **GADGET-3** is limited, and we point the curious reader to Springel et al. (2021) for details on the state-of-the-art **GADGET** methodology. We will describe the setup of the three-dimensional simulations used in this work in Sec. 3.2.2.

3.1.2. Time integration

Time integration refers to updating the positions and momenta of the system of macroparticles. For a fixed step size Δt , a simple update scheme Taylor-expands the position and momentum about time t as follows

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \frac{\mathbf{p}(t)}{ma^2(t)} \Delta t, \quad (3.6)$$

$$\mathbf{p}(t + \Delta t) = \mathbf{p}(t) - m \nabla \phi(\mathbf{x}, t) \Delta t. \quad (3.7)$$

This is the forward Euler method. While it is simple to implement, it is a first-order method, meaning that the local truncation error is $\mathcal{O}(\Delta t^2)$, and the global error over a fixed interval of integration is $\mathcal{O}(\Delta t)$ (see, e.g. Iserles, 2009), which is too large for our purposes.

To make progress, recall that the system of DM particles follow Liouville’s theorem (see Sec. 2.1 and App. A.4). Thus, the system is described by a Hamiltonian, a property that is inherited by the macroparticle N -body ensemble. This Hamiltonian is defined as

$$H(\mathbf{x}, \mathbf{p}, t) = \sum_{i=1}^N \left[\frac{\mathbf{p}_i^2(t)}{2ma(t)^2} + \frac{m\phi_i(\mathbf{x}, t)}{2} \right]. \quad (3.8)$$

We can cast Eq. (3.2) and (3.3) in the Hamiltonian formalism as (see also App. A)

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}}, \quad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}}. \quad (3.9)$$

The equations admit an exact solution of the form (see e.g. Hairer et al., 2006; Angulo & Hahn, 2022)

$$\zeta_i(t) = \mathcal{T} \exp \left[\int_0^t dt \mathcal{H}(\mathbf{x}, \mathbf{p}, t) \right] \zeta_i(0), \quad (3.10)$$

where $\zeta = (\mathbf{x}, \mathbf{p})$, $\mathcal{H}$ is the Hamiltonian *operator* defined by the action of the Poisson bracket (App. A.3) as $\mathcal{H}(t) = \{\cdot, H(t)\}$, and $\mathcal{T}$ is a time-ordering operator necessary because H contains explicit time-dependence. Given any products of terms evaluated at different times, $\mathcal{T}$ orders the products such that terms evaluated at earlier times appear earlier⁷. Let us make progress towards a numerical implementation of Eq. (3.10). First, the integral from some initial time 0 to time t may be broken down into a series of small, finite steps Δt , such that

$$\int_0^t = \int_0^{\Delta t} + \dots + \int_{t-\Delta t}^t . \quad (3.11)$$

Next, observe that the Hamiltonian of Eq. (3.8) is naturally split into a kinetic part that contains only a momentum dependence, and a potential part that depends only on the position. This allows the definition of two operators called *drift* and *kick*

$$\mathcal{D}(t_2 - t_1) = \sum_{i=1}^N \exp \left(\int_{t_1}^{t_2} dt \frac{\mathbf{p}_i^2(t)}{2ma(t)^2} \right) , \quad (3.12)$$

$$\mathcal{K}(t_2 - t_1) = \sum_{i=1}^N \exp \left(\int_{t_1}^{t_2} dt \frac{m\phi_i(\mathbf{x}, t)}{2} \right) , \quad (3.13)$$

where we exchanged the integral and the sum⁸.

The drift operator keeps the momentum fixed, and evolves the position (in effect, solving Eq. 3.2), while the kick operator updates the momentum at fixed particle positions (solving Eq. 3.3). Then, Eq. (3.10) may be discretised by a series of successive drift and kick operations as follows

$$\begin{aligned} \mathcal{T} \exp (\mathcal{H} \Delta t) &= \mathcal{T} \exp [(\mathcal{K} + \mathcal{D}) \Delta t] \\ &\approx \exp (\mathcal{K} \Delta t) \exp (\mathcal{D} \Delta t) . \end{aligned} \quad (3.14)$$

Here, the exponent of the sum of the two operators is approximated by the product of their exponents. This introduces an error that is proportional, at leading-order, to their commutator. As the commutator does not appear in the expression, the time-ordering has no effect. The following result (Hairer et al., 2006; Saha & Tremaine, 1992),

$$\exp A \exp B = \exp \left(A + B + \frac{1}{2}[A, B] + \frac{1}{12}[A - B[A, B]] + \dots \right) , \quad (3.15)$$

⁷To elaborate on this, the Hamiltonian $H(t)$ does not commute with itself at a different time t' .

If any such commutators appear, the exponentiation becomes intractable. This is the purpose of $\mathcal{T}$, as the commutators will arise as a consequence of the exponential theorem (Eq. 3.15).

The operator is important in quantum field theory, and was first introduced by Dyson (1949).

⁸The interchange of the finite sum and the integral is permitted when the individual integrals are positive. In measure-theoretic language, the finite sum is a measure on a finite subset of natural numbers, and the integral is a Lebesgue measure on the positive real numbers. Then, the interchangeability is a special case of the Fubini-Tonelli theorem (see, e.g. Tao, 2011).

is called the exponential theorem, and often named after⁹ Campbell, Baker, Hausdorff, and (sometimes) Dynkin (CBHD). An improvement may be introduced by combining the drift and kick operators with different weights. Generically (Dehnen & Read, 2011),

$$\exp(H\Delta t) \approx \prod_{i=1}^N \exp(a_i \mathcal{K}\Delta t) \exp(b_i \mathcal{D}\Delta t) , \quad (3.16)$$

where the coefficients a_i and b_i can be chosen to the desired accuracy. For example, a second-order scheme sets $a_1 = a_2 = 1/2, b_1 = 1, b_2 = 0$, and a single time-step can be implemented as follows

$$\zeta_i(t + \Delta t) = \mathcal{T} \left[\mathcal{K} \left(\frac{\Delta t}{2} \right) \mathcal{D}(\Delta t) \mathcal{K} \left(\frac{\Delta t}{2} \right) \right] \zeta_i(t) , \quad (3.17)$$

where the kick may be further split into short- and long-range parts for hybrid methods. This is called the *leapfrog* integrator and has a global error of $\mathcal{O}(\Delta t^3)$. Importantly, the leapfrog integrator is *symplectic*. Such an integrator is exact (the error comes from an approximate Hamiltonian, not an approximate integration), and preserves phase-space volume. This is also called a *kick-drift-kick* (KDK) scheme for the order in which the operators appear. Exchanging a and b defines a *drift-kick-drift* (DKD) scheme which is conceptually identical, but computationally more involved (see e.g. Dehnen & Read, 2011, specifically their footnote 7). N -body simulations, therefore, prefer using KDK for the time integration.

So far, we have assumed that the time-step Δt is constant. Simulations have a vast temporal range – high-density regions evolve at much faster time-scales than their low-density counterparts. With this in mind, it is practical to relax this assumption. Allowing the time-step to vary does, however, destroy symplecticity. This can be partially compensated for by making the scheme time-reversible¹⁰ (Dehnen, 2017). Different N -body simulations have different time-stepping criteria. For example, **GADGET-2** defines the time step for a particle based on its acceleration $\mathbf{a}$ and the gravitational softening such that (Springel, 2005)

$$\Delta t = \min \left[\Delta t_{\max}, \left(\frac{2\eta\epsilon}{|\mathbf{a}|} \right)^{1/2} \right] , \quad (3.18)$$

where $\Delta t_{\max}$ is a user-defined maximum time-step, and η is an accuracy-controlling parameter. For the **tree** part of its **tree-PM** implementation, **GADGET-2** uses the

⁹This theorem is deeply embedded in the study of Lie groups, with a history as rich as its applications. As is often the case, none of CBHD were the first to work on the theorem. Pascal, Poincaré, and Schur all made important early contributions, although these have largely been forgotten. A detailed account of these contributions, and those of the authors to which the theorem is credited can be found in Achilles & Bonfiglioli (2012).

¹⁰Time-reversibility means that the simulation can be run in reverse to recover the initial configuration from any final state by simply changing the sign of the particle momenta.

above criterion, while for the long-range forces, the time-step is fixed to $\Delta t_{\max}$. A more modern time-stepping criterion splits the particles into *fast* and *slow*, and split the kick operator in two corresponding pieces. This is implemented hierarchically, such that each set of fast particles is further split into fast and slow. This *block* (hierarchical) time-stepping was first discussed by Hayli (1967, 1974), and is the scheme of choice for **GADGET-4** (Springel et al., 2021) and **PKDGRAV-3** (Potter et al., 2017).

3.1.3. Poisson solver

To move the particles of a simulation in accordance with Eq. (3.2), the gravitational potential at the location of each particle due to all other particles must be evaluated. This constitutes solving the Poisson equation (Eq. 3.4). As pointed out in Sec. 3.1.1, using direct summation is too expensive for a large set of particles, and most modern simulations use the **tree-PM** method. For the **PM** part, the Poisson equation may be solved using the Green’s function of the Laplacian which may be written in real-space as

$$G(\mathbf{x}) = -\frac{1}{4\pi|\mathbf{x}|} . \quad (3.19)$$

Then, the formal solution of Eq. (3.4) is a convolution with the Green’s function

$$\phi(\mathbf{x}, t) = G(\mathbf{x}) * 4\pi G\bar{\rho}a^2(t)\delta(\mathbf{x}, t) . \quad (3.20)$$

However, the Green’s function is singular, viz., $G \rightarrow \infty$ as $\mathbf{x} \rightarrow 0$. This necessitates the softening procedure hinted at earlier in the parent section. There are broadly two approaches to implement gravitational softening. The first idea is to apply a cutoff to the potential (or the Green’s function – the mathematical procedure is identical) to prevent the small-scale singularity. Physically, this is equivalent to replacing the point-masses with Plummer spheres à la Aarseth & Hoyle (1963) (see also Hernquist & Barnes, 1990). The softened Green’s function is given as

$$G_\epsilon(\mathbf{x}) = -\frac{1}{4\pi\sqrt{|\mathbf{x}|^2 + \epsilon^2}} \quad (3.21)$$

for a softening length ϵ . This *Plummer softening* biases the force calculation as it recovers the Newtonian potential only as $\mathbf{x} \rightarrow \infty$. Alternatively, softening may also be viewed as smoothing the point-masses with a kernel (Barnes, 2012). Depending on the chosen kernel, the Newtonian potential is recovered within a distance of few times the softening length. The **GADGET** codes adopt a spline kernel first defined in Monaghan & Lattanzio (1985), and this factor is 2.8 (Springel et al., 2021). The exact choice of the kernel and that of the softening length is subject to ongoing debate (see, e.g. Athanassoula et al., 2000; Dehnen, 2001; Power et al., 2003; Iannuzzi & Dolag, 2011; van den Bosch & Ogiya, 2018; Zhang et al., 2019).

Thus, a PM code solves the Poisson equation on a fixed mesh by transforming Eq. (3.20) to Fourier space, where the convolution becomes a product, and using a softened Green's function. The transformation is achieved efficiently using discrete Fourier transforms, employing the Fast-Fourier-Transform (FFT) algorithm of Cooley & Tukey (1965). The algorithm has a complexity of $\mathcal{O}(N_g^3 \log N_g^3)$ for N_g grid points. N -body simulations have a large spatial range. At late times, many high-density regions need to be resolved, and a coarse global grid is not sufficient. Using additional grids with higher-resolution is then desirable, at the cost of computational and algorithmic complexity. These additional grids are added adaptively, with high-density regions covered by high-resolution grids. The ART algorithm utilises such an adaptive multigrid method. On the coarsest level, the discrete Poisson equation is solved using FFTs. This solution is then interpolated onto the finer grid, yielding boundary values and a first guess for the potential. The solution at that level is found using a process called *relaxation*, which starts with an initial guess that iteratively arrives at better and better approximations by solving a finite-difference equation. Since the initial guess is already approximately correct, the method converges quickly.

On small scales, the heart of the **tree-PM** code is the Barnes-Hut **tree** algorithm, which consists of three steps

1. First, the simulation volume is divided into a hierarchical structure called a tree. This is done recursively, until each subvolume is either empty, has a pre-defined minimum number particles, or the maximum number of iterations is exceeded. The last two are user-defined parameters. The maximum number of iterations is called the *depth* of the tree.
2. Next, for each node of the tree along all the subvolumes, the center of mass (COM) is calculated and stored.
3. Finally, the force on each particle is calculated by traversing the tree, such that nodes which are further away are grouped together into one macroparticle.

We demonstrate these steps for ten randomly-distributed points in two-dimensional space. The result is visualised in Fig. 3.1, where we show both the quadtree (left), and the graph view of the structure. As an example, if we have to calculate the force on particle 5, we start at root of the tree (the node marked as 'R'). Since R has multiple children, we iterate over them. The node containing particle 8 is an external node, and the force exerted by it on particle 5 is calculated by direct summation. The two children C1 and C2 are internal nodes, and we iterate over them. For C2, we check whether the node is sufficiently far from particle 5 to be considered a macroparticle. This is done by calculating the ratio of the node size to the distance between particle 5 and the node COM. This is called the *opening angle*, and if it is less than a user-defined threshold (typically 0.5), the macroparticle assumption is made. The force exerted by the node COM on particle 5 is then calculated, and added to the total force. For C1, the ratio is greater than




Figure 3.1. – *Left:* Barnes-Hut quadtree for ten random points. *Right:* The corresponding tree graph.

this threshold, and we must iterate over its children. This process continues until the full tree is resolved. These steps are then repeated for each particle.

An important approximation is to then write the gravitational potential in a multipole expansion. Let us illustrate this with an example for an arbitrary Green's function (loosely following Salmon, 1991). First, assume a partition of the simulation volume $\mathcal{V}$ into disjoint subvolumes (nodes) as follows

$$\mathcal{V} = \sum_S \mathcal{V}_S . \quad (3.22)$$

Given a node $\mathcal{V}_A$ with COM $\boldsymbol{\lambda}_A$, and a particle $\mathbf{x}_A \in \mathcal{V}_A$, we would like to compute the gravitational potential this particle feels due to all other nodes. This can be expressed as a sum over the contributions from each node

$$\phi(\mathbf{x}_A, t) = \sum_S \phi_S(\mathbf{x}_A, t) , \quad (3.23)$$

Now, consider a node $\mathcal{V}_B$ well-separated from $\mathcal{V}_A$ with COM $\boldsymbol{\lambda}_B$. The contribution to the potential at $\mathbf{x}_A$ due to $\mathcal{V}_B$ is given as (following Eq. 3.20)

$$\phi_B(\mathbf{x}_A, t) = 4\pi G \bar{\rho} a^2 \int_{\mathcal{V}_B} d^3\mathbf{y} \, G_\epsilon(\mathbf{x}_A - \mathbf{y}) \delta(\mathbf{y}, t) . \quad (3.24)$$

Thus, for a given point $\mathbf{y} \in \mathcal{V}_B$, we would like to evaluate the Green's function $G_\epsilon(\mathbf{x}_A - \mathbf{y})$. We can rewrite the latter as

$$\begin{aligned} G_\epsilon(\mathbf{x}_A - \mathbf{y}) &= G_\epsilon[(\mathbf{x}_A - \boldsymbol{\lambda}_A) - (\mathbf{y} - \boldsymbol{\lambda}_A)] \\ &\approx G_\epsilon[(\mathbf{x}_A - \boldsymbol{\lambda}_A) - \mathbf{R}] , \end{aligned} \quad (3.25)$$

where $\mathbf{R} = \boldsymbol{\lambda}_B - \boldsymbol{\lambda}_A$, and we have used the fact that $\mathbf{y} - \boldsymbol{\lambda}_A \approx \mathbf{R}$. Let us define the separation between the particle at $\mathbf{x}_A$ and the COM of $\mathcal{V}_A$ as

$$\mathbf{r} = \mathbf{x}_A - \boldsymbol{\lambda}_A . \quad (3.26)$$

This is small compared to the separation $\mathbf{R}$ between the subvolumes. Then,

$$G_\epsilon(\mathbf{r} - \mathbf{R}) \approx G_\epsilon(-\mathbf{R}). \quad (3.27)$$

The Taylor series for this quantity about $\mathbf{r}$ is

$$G_\epsilon(-\mathbf{R}) = G_\epsilon(\mathbf{r}) - (\mathbf{R} \cdot \boldsymbol{\nabla}) G_\epsilon(\mathbf{r}) + \frac{1}{2!} (\mathbf{R}^{\otimes 2} \cdot \boldsymbol{\nabla}_2) G_\epsilon(\mathbf{r}) - \frac{1}{6!} (\mathbf{R}^{\otimes 3} \cdot \boldsymbol{\nabla}_3) G_\epsilon(\mathbf{r}) + \dots , \quad (3.28)$$

where, $\mathbf{R}^{\otimes i} = \underbrace{\mathbf{R} \otimes \dots \otimes \mathbf{R}}_{i \text{ times}}$ is a rank i tensor, and contracts with $\boldsymbol{\nabla}_i$ to yield a scalar.

Going back to Eq. (3.25), we can write

$$G_\epsilon(\mathbf{x}_A - \mathbf{y}) = \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} [(\mathbf{y} - \boldsymbol{\lambda}_A)^{\otimes n} \cdot \boldsymbol{\nabla}_n] G_\epsilon(\mathbf{x}_A - \boldsymbol{\lambda}_A) , \quad (3.29)$$

where we re-substituted the definition of $\mathbf{r}$, and the approximation $\mathbf{R} \approx \mathbf{y} - \boldsymbol{\lambda}_A$. Plugging this expression into Eq. (3.24), we get

$$\phi_B(\mathbf{x}_A, t) = 4\pi G \bar{\rho} a^2 \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} [\boldsymbol{\mu}^{(n)} \cdot \boldsymbol{\nabla}_n] G_\epsilon(\mathbf{x}_A - \boldsymbol{\lambda}_A) , \quad (3.30)$$

where we exchanged the integral and infinite sum using the Fubini-Tonelli theorem (see footnote 8 on page 47). The n^{th} multipole moment of the mass distribution in $\mathcal{V}_B$ is a rank- n tensor defined as

$$\boldsymbol{\mu}^{(n)} = \int_{\mathcal{V}_B} d^3\mathbf{y} (\mathbf{y} - \boldsymbol{\lambda}_A)^{\otimes n} \delta(\mathbf{y}, t) . \quad (3.31)$$

Although we chose the pivot $\boldsymbol{\lambda}_A$ as the COM of the node, this is not compulsory. The FMM algorithm as implemented by Greengard & Rokhlin (1987), e.g. uses the geometric center of the node. Using the COM is equivalent to moving to a frame of reference where the COM is zero, and then using the origin as the pivot for the moment calculation. However, in this frame of reference, the dipole moment is exactly the same as the COM, which is zero. Thus, the dipole moment always vanishes when using the COM as the pivot. This is a consequence of the conservation of linear momentum.

There are two possibilities for evaluating the moments of Eq. (3.31). One is to consider all bodies in $\mathcal{V}_B$, and calculate the following discrete moments by brute force

$$\boldsymbol{\mu}^{(n)} = \frac{1}{\bar{\rho}} \sum_{\mathbf{y} \in \mathcal{V}_B} m(\mathbf{y} - \boldsymbol{\lambda}_A)^{\otimes n} , \quad (3.32)$$

To simplify the computation for a given node, one can use only its eight children instead of all particles within it. This leads to a recursive procedure starting from the external nodes and traversing up the tree, using the previous computations for its parent node. Comparing this procedure to the integral in Eq. (3.24), we can finally see the power of the multipole expansion – the complicated convolution integral over the subvolume has been converted to a series of multipole moments that can be pre-computed and looked up during the potential calculation. Typically, only the monopole is used for estimation, but higher-order multipoles may be used if higher accuracy is desirable. The opening angle criterion determines whether a node is resolved by going further down the tree and considering its child nodes, or whether a multipole expansion is used to approximate the potential (and thus the force) due to this node. **GADGET-1** used a quadrupole expansion allowing for larger opening angles. This means that fewer computations are required for the desired accuracy. However, the complexity per computation increases (Springel, 2005). **GADGET-2** and its successors use the monopole instead, which significantly lowers the required memory usage. In the example tree that we constructed earlier, we used a geometric definition of the opening angle – for a particle at a distance l from a node of size r , the opening angle must be smaller than a threshold θ_t for the multipole expansion to be applied

$$\frac{l}{r} \leq \theta_t, \quad (3.33)$$

where θ_t controls the accuracy and the computational cost of the approximation. The above definition does not depend on the mass of the node. In both **GADGET** codes, a dynamic criterion using the node mass is employed instead. Specifically, in **GADGET-2**,

$$\frac{GM}{r^2} \left(\frac{l}{r} \right)^2 \leq \alpha |\mathbf{a}|, \quad (3.34)$$

where $\mathbf{a}$ is the total acceleration of the last time-step, and α defines a tolerance parameter.

3.1.4. Mass assignment

The mass-assignment is generally done using either of the nearest-grid-point (NGP; Hockney, 1966), cloud-in-cell (CIC; Birdsall & Fuss, 1969), or triangular-shaped-cloud (TSC; Hockney & Eastwood, 1988) interpolation algorithms. Each of the mass-assignment schemes can be summarised by writing their weighting functions W for the mass on the grid defined as $M(\mathbf{x}_j) = m \sum_{i=1}^n W(\mathbf{x}_i, \mathbf{x}_j)$. Consider a one-dimensional box of length L . Defining the separation between particle and grid

positions as $\Delta x = |x_i - x_j|$, the weighting functions are (e.g. Hockney & Eastwood, 1988; Bagla & Padmanabhan, 1997)

$$W_{\text{NGP}}(\Delta x) = \begin{cases} 1 & \text{if } \Delta x \leq \frac{L}{2} , \\ 0 & \text{if } \Delta x > \frac{L}{2} . \end{cases} \quad (3.35)$$

$$W_{\text{CIC}}(\Delta x) = \begin{cases} 1 - \frac{\Delta x}{L} & \text{if } \Delta x \leq L , \\ 0 & \text{if } \Delta x > L . \end{cases} \quad (3.36)$$

$$W_{\text{TSC}}(\Delta x) = \begin{cases} \frac{3}{4} - \frac{\Delta x}{L} & \text{if } \Delta x \leq \frac{L}{2} , \\ \frac{9}{8} \left(1 - \left|\frac{2\Delta x}{3L}\right|\right)^2 & \text{if } \frac{L}{2} < \Delta x \leq \frac{3L}{2} , \\ 0 & \text{if } \frac{3L}{2} < \Delta x . \end{cases} \quad (3.37)$$

The three-dimensional weighting takes a product of these kernels along the three axes. NGP assigns the mass of each particle to a single grid point. This is the fastest of the three schemes, but introduces a large truncation error (Efsthathiou et al., 1985). CIC (TSC) divides the mass of each particle between the nearest 8 (27) particles in the three-dimensional cubic grid. The choice between CIC and TSC is a question of balancing between computational speed and accuracy (with the latter being more accurate but slower).

3.1.5. On finding structure in simulations

Solving the gravitational dynamics of macroparticles is only the first step in closing the gap between simulations and observables. The next step is to consider biased tracers of the dark matter overdensity field – halos. In Chapter 2, we introduced halos as units of clustering of dark matter. Specifically, Sec. 2.3 discussed the theoretical modelling of halos. In this section, we briefly discuss the different ways in which halos and subhalos may be identified in N -body simulations. Compared to the operational costs of N -body simulations, these structure-finding algorithms are inexpensive, but storing the halos across simulation snapshots is memory-exhaustive. To combat this, simulations are designed to carry out the structure-finding process on-the-fly, so that the memory-inexpensive data products can be stored at a finer time-resolution.

After the first galaxy-redshift surveys were carried out (see Chapter 1), a series of numerical algorithms were designed to parse the galaxy catalogues to assign group membership to them. The name *group-finding* algorithms may be ascribed to them, and has passed on to structure-finding algorithms in simulations. These can be broadly classified in two categories – peak-finding, and particle-collecting. Algorithms in the former category identify local density peaks, then group the particles

around the peak into halos based on some criteria. Particle-collecting algorithms, on the other hand, first link nearby particles together, and then determine a group center. Once the algorithm finds a set of particles, the set is pruned to remove particles which are not gravitationally bound to the center by a process called *unbinding*. Within these two broad categories, there are a multitude of more nuanced choices that distinguish different codes. The reason for this diversity is simple – there is no universally agreed upon definition of what constitutes a halo, and particularly a halo boundary. What follows is a survey of halo finders, with special emphasis towards the algorithm we employ in this work, viz. the **Amiga Halo Finder** (AHF¹¹; Knollmann & Knebe, 2009). Before we begin, we point the curious reader to Knebe et al. (2011, 2013), and references therein, for a definitive account of the various halo finders and how their predictions compare.

Halos as density peaks

As we discussed in Sec. 2.3, the simplest way to think about halos is to assume a spherical cluster of objects. Following Press & Schechter (1974), the *spherical overdensity* (S0) criterion defines a halo to be a spherical volume whose internal mass distribution is overdense compared to the mean density of the Universe by a factor Δ . For example, to concur with the prediction of top-hat spherical collapse, $\Delta \approx 178$. An algorithm to find halos in this way can be summarised as follows (Lacey & Cole, 1994)

1. Find the local number density for each particle in the simulation. This is done by drawing a spherical volume of radius r_N around a particle, where r_N is the distance to its N^{th} nearest neighbour.
2. From a sorted list ℓ , pick the particle with the highest local density as a candidate center of the first halo. Draw spherical shells around this center until the mean overdensity within the shells first falls below Δ .
3. Shift the center to the center of mass of the particles in this sphere, and iterate over the procedure in step two. Repeat this process until the shift in the centers is bigger than a user-defined threshold ΔR .
4. Once the iterations stop, all particles that fall in the current spherical shell are labelled as members of a group. This process is repeated for the next particle in ℓ which does not already belong to a group, and is continued until all particles are resolved.

While S0 is conceptually simple, it suffers from an important drawback – it defines halos to be spherical, while real gravitationally bound objects are expected to be shaped by tidal forces into aspheroids. In fact, algorithms which do not insist on

¹¹<http://www.popia.ft.uam.es/AMIGA>

a halo shape (see e.g. **FoF** below) find halos in simulations to be tri-axial (Jing & Suto, 2002; Allgood et al., 2006). Using this as motivation, Despali et al. (2013) extend the **S0** algorithm to allow for triaxiality. The idea hinges on a mass tensor whose eigenvectors are the axes of the ellipsoid. The procedure starts by running the **S0** algorithm on the simulation data to identify spherical halos, then iteratively relaxes the sphericity criterion such that the eigenvalues of the mass tensor converge between iterations to $< 1\%$. Group membership is decided based on the ellipsoidal distance of a particle from the group center.

Let us now turn to the **AHF** algorithm that we use to catalogue the halos in this work. **AHF** is an improvement on the **MLAPM** halo finder (Gill et al., 2004; Knollmann & Knebe, 2009) which used the adaptive grid of the **MLAPM** N -body code to identify local overdensities iteratively at each level of refinement. The unbinding was performed on-the-fly with each iteration. The improvement, relative to **MLAPM**, is mainly computational. **AHF** allows efficient parallel-processing of a large amount of data. The algorithm proceeds as follows

1. Start with a user-specified grid resolution and cover the simulation volume with this grid. Assign mass (density) to particles on the grid using a TSC assignment scheme. Given a user-specified refinement threshold, construct finer grids if the particle density is higher than the threshold. Recompute the particle density at each refinement, in this way obtaining a hierarchical grid that follows the density field until all cells have been refined.
2. Going to the finest level, identify and mark isolated regions (density peaks) as halo candidates. Then proceed along the grid hierarchy to mark halo candidates at coarser refinements, while mapping halo candidates from the finer grid onto their superset in the coarser grid.
3. From this nested tree, generate a list of candidate particles that might belong to the halo. First, all particles that fall in the spherical region encompassing matter with overdensity Δ (i.e., the **S0** criterion) are members of this list.
4. Using a criterion based on the local escape velocity, iteratively remove particles from this initial sample until only the gravitationally bound particles remain.

In this way, **AHF** finds halos within the simulation volume, and generates a list of particles that belong to each halo. Although **AHF** can also make halo catalogues on-the-fly, we use it as a stand-alone halo finder on the simulation snapshots.

Halos as collections of particles

An alternative class of structure-finding algorithms does not assume halo geometry, but links neighbouring particles together to form groups. The variation of algorithms within the class stems from the definition of ‘neighbouring’. Let us sketch one such

algorithm, described e.g. in Huchra & Geller (1982), who were trying to collect galaxies catalogued in a redshift survey into gravitationally bound groups.

1. Start with a random galaxy in the catalogue. Defining a maximum projected separation, identify other galaxies within that separation as companions, or – as Press & Davis (1982) call them – *friends*. This maximum separation is called the linking length (conventionally, b), and expressed in units of the mean-interparticle distance in the simulation. To be precise, for galaxies i and j to be companions,

$$\frac{(\mathbf{x}_i - \mathbf{x}_j)^2}{b^2} < 1 . \quad (3.38)$$

2. All companions are added to a list which is to become the group. Galaxies without companions are identified as *isolated*.
3. For every companion, identify companions which also satisfy the separation criterion and add them to the group.
4. For these *friends-of-friends*, further search for companions that satisfy the criterion, and repeat this until no further members can be found.

These steps are then repeated for each galaxy that has not yet been assigned to a group. An unbinding procedure based on particle potential energies may finally be carried out. In this way, the algorithm organises the galaxies into a collection of groups, and a set of isolated galaxies. Press & Davis (1982) applied this algorithm to an N -body simulation, and it is termed friends-of-friends (FoF). As pointed out in Frenk et al. (1988) and elaborated upon in More et al. (2011), FoF sketches out an isodensity surface in the limit of large particle number. This surface has the overdensity (or rather, $1 + \delta$) value $\approx mb^{-3}$, with $m = \mathcal{O}(1)$. For example, Frenk et al. (1988) found $m = 2$, while Lacey & Cole (1994) quote $m = 3/2\pi \approx 0.48$. The more recent figure is $m \approx 0.65$ from More et al. (2011). With the almost universally used $b = 0.2$, the latter figure implies an overdensity approximately 80 times the mean density of the universe – a few times smaller than the threshold for spherical collapse. FoF has several appealing features. It is, for one, algorithmically simple. Only one free parameter fixes the scheme¹², and a fixed b results in a unique halo catalogue. Further, it does not define a fixed halo shape, thus allowing for triaxial halos, which is consistent with what we expect from real halos. Despite this, FoF suffers from some numerical artefacts. In particular, the masses of halos systematically depend on the mass resolution of the simulation (e.g., More et al., 2011; Leroy et al., 2021), and accounting for this still leaves us with a number of discreteness effects (Warren et al., 2006; Lukić et al., 2009; More et al., 2011). Namely, the performance of the correction itself depends on redshift and cosmology,

¹²In practice, it is important to set a minimum number of particles required for a collection to be labelled as a halo, in order to avoid spurious numerical artefacts.

and there is a small but systematic error dependent on the softening scale (Ludlow et al., 2019; Ondaro-Mallea et al., 2022).

FoF also has a tendency to link two structures together with unphysical particle bridges (see e.g. Lacey & Cole, 1994). One way to fix this issue is to identify those particles which – if removed from the group – will unlink the distinct structures. Another more physically motivated approach is to exploit the full phase-space information available in the simulation. This is complicated by the fact that there is no unique notion of a phase-space distance. The simpler approach is to define a criterion for converting the velocity-space separation to a distance, and then modify Eq. (3.38) to include this criterion. Defining b_v as a velocity-space linking length expressed in units of velocity dispersion, two particles i and j are considered friends if

$$\frac{(\mathbf{x}_i - \mathbf{x}_j)^2}{b^2} + \frac{(\mathbf{v}_i - \mathbf{v}_j)^2}{b_v^2} < 1 . \quad (3.39)$$

This was used in the 6DFoF algorithm (Diemand et al., 2006), and in VELOCIRAPTOR (Elahi et al., 2019). There are three free parameters – the two linking lengths, and the minimum particle number for a collection to be called a halo. Using a fixed linking length is limiting, especially when finding substructure. A better approach is to adapt the linking length depending on the structure in question. This is the idea behind ROCKSTAR¹³ (Behroozi et al., 2012), which also stresses on maintaining a consistent accuracy across multiple snapshots, thus creating better merger trees.

Despite the many conceptual and algorithmic choices made by halo finders, they largely agree on the positions of halos and subhalos, and their statistical distribution (Onions et al., 2012; Knebe et al., 2013). However, investigating major mergers reveals strong biases in predicted children masses among the different halo finders (Behroozi et al., 2015), with phase-space finders performing better than their configuration-space counterparts.

It is often desirable to chart the evolutionary history of specific halos across simulation snapshots. In order to do this, halo finders can keep track of the parents and grandparents of each halo. This creates a hierarchical structure, whose root node is the full volume at the beginning of the simulation, and traversing it downward to $z = 0$ leads to the halos identified by the halo finder. This is called the merger tree, and is useful, e.g. to populate halos with galaxies. The halo finder achieves this by tracking particle IDs across snapshots to identify the evolution of a halo. Often, mergers can result in two parent halos birthing a daughter halo. In this case, the main progenitor is identified by maximising a merit function. Given a halo i and a progenitor j , AHF defines the merit function as

$$M_{ij} = \frac{N_{i \cap j}^2}{N_i N_j} , \quad (3.40)$$

¹³<https://bitbucket.org/gfcstanford/rockstar>

where N is the number of particles in the indexed halo, and $N_{i \cap j}$ is the number of particles shared by halos i and j . The progenitor which maximises the merit function is labelled as the main progenitor. There are many stand-alone merger tree codes with different merit functions and algorithms. For a systematic comparison, the reader may refer to Srisawat et al. (2013); Wang et al. (2016). For our work, we calculate the merger tree using AHF’s native algorithm.

3.1.6. Convergence and other numerical issues

Before we discuss the specifics of the simulations we used in this work, it is important to ask whether numerical simulations can faithfully resolve phase-space dynamics. This is a fundamental problem in numerical simulations, and goes under the name of *convergence*. There are a few ways in which convergence may be tested. One common approach is to vary the mass resolution of a simulation and look at the behaviour of a test statistic to evaluate the numerical stability of the simulation. Another approach is to compare the results among different N -body codes with identical initial conditions and numerical setups. Such systematic tests were carried out in Heitmann et al. (2005, 2008), who found a general agreement in the test statistics (e.g., the matter power spectrum). However, the disagreements in several cases were of the order of 10%, which far exceeds the precision that the near-future observations were projected to achieve. A more recent code comparison was performed in anticipation of the *Euclid* survey (Schneider et al., 2016). This project compared the matter power spectrum obtained using RAMSES, PKDGRAV-3, and GADGET-3, and found a sub-percent agreement up to $k = 1 \, h \text{ Mpc}^{-1}$. This sub-percent agreement bodes well for the convergence of N -body codes, and we can be reasonably confident in their results.

Another concern is the validity of the N -body discretisation itself. Earlier in this chapter, we alluded to the fact that the choice of the gravitational softening scale is subject to debate. If the softening scale is much smaller than the mean interparticle separation, two-body collision can arise as spurious numerical artifacts (Melott et al., 1997). A more recent study suggests that the softening scale should be about twice the interparticle distance (Romeo et al., 2008). As discussed before, simulations typically set the softening scale to smaller than unity. Romeo et al. (2008) also point out that fortunately, the discreteness effects introduced on small scales typically do not propagate to larger scales at later times. A robust, conclusive test of the N -body approximation would be to compare results of N -body codes with direct solution of the Vlasov equation. One such test in one-dimension found the N -body halo density profile to be inconsistent with the corresponding Vlasov solution. To our knowledge, there is only one implementation in three-dimensions, viz. CoLDICE (Sousbie & Colombi, 2016). CoLDICE solves the Vlasov equation by tessellating the six-dimensional phase space, such that the vertices of the tessellation satisfy the equations of motion. Colombi (2021) compares the early phases of halo

formation between `CoLDICE` and a `PM` code in an idealised, symmetric setting. Given the resolution criterion – at least one particle per cell – is satisfied, both simulations agreed on the properties of the proto-halos. Although this avenue is hampered by the cost of running a Vlasov simulation, the early signs are encouraging. In summary, despite some unsettled issues, there are good reasons to rely on N -body codes to simulate the non-linear regime of structure formation.

3.2. Simulations used in this work

3.2.1. Simulations in one-dimension

To approximate the effects of gravity in one-dimension, we can imagine a mass configuration consisting of infinite two-dimensional slabs. The gravitational potential due to a single slab is constant, and depends only on the density of the slab. The net gravitational force at a point due to N slabs can be calculated by simply counting the slabs on either side of this point. If the number of slabs on either side is equal, the force is zero. Otherwise, it is proportional to the number of excess slabs. Thus, the exceptional symmetry allows us to bypass the direct force calculation that makes three-dimensional simulations expensive. For our tests of the `EFTofLSS`, we run one-dimensional simulations in an EdS universe using the `cosmo_sim_1d` code¹⁴ (Rampf et al., 2021).

Let us demonstrate the calculation of gravitational acceleration in this system of N identical slabs. Given a simulation volume bound by the periodic one-dimensional interval $[0, 1)$, consider a slab located at the COM. This slab has an equal density of matter on both sides, and feels no acceleration. Thus, the acceleration on a generic slab is proportional to its position relative to the COM. In `cosmo_sim_1d`, this calculation is simplified by sorting the positions. Once the acceleration is calculated, the code updates the positions and velocities of the slabs. This is accomplished using a second-order, symplectic, DKD integrator (see Sec. 3.1.2).

The publicly available version of `cosmo_sim_1d` only outputs the final state of the system at the end of the simulation run. We modify the code to return snapshots at intervals equally-spaced in $\log a$. Each snapshot contains the updated particle positions and velocities. We interpolate the particles to a uniform grid in order to evaluate the moments and cumulants of the Vlasov distribution. A development version of the code¹⁵ also outputs members of the Vlasov hierarchy, which we use to validate our computation.

¹⁴https://bitbucket.org/ohahn/cosmo_sim_1d

¹⁵https://bitbucket.org/ohahn/cosmo_sim_1d/src/develop/

	box100	box300
Ω_m	0.315	0.315
Ω_Λ	0.685	0.685
Ω_b	0.0493	0.0493
H_0 [km s ⁻¹ Mpc ⁻¹]	67.4	67.4
Box Size [h^{-1} Mpc]	100	300
Number of particles	512 ³	512 ³
Mass resolution [$h^{-1} M_\odot$]	2.068×10^9	5.584×10^{10}

Table 3.1. – Simulation parameters for box100 and box300. The only difference between the two is the box size, which leads to a different mass resolution.

3.2.2. GADGET-3, and simulations in three-dimensions

For our exploration of the assembly of halos, we use three-dimensional cosmological simulations run using **GADGET-3** (Springel et al., 2009). The initial conditions for the simulations were set using second-order LPT via the **MUSIC** code¹⁶ (Hahn & Abel, 2011). We assume that inflation sets the initial conditions of the Universe such that the density fluctuations are fully described by Gaussian statistics. Numerical simulations must begin at a much later redshift. For a given wavenumber, evolution between these two epochs is dictated by when it enters the cosmic horizon. Let k_{eq} denote the wavenumber at matter-radiation equality. Then, perturbations at $k < k_{\text{eq}}$ were outside the horizon in the radiation-dominated epoch and grow differently from perturbations at $k \geq k_{\text{eq}}$. It is conventional to define a propagator that captures the spatial evolution of different scales from some initial time to the current epoch (or to the beginning of a simulation). This propagator is called the *transfer function*. **MUSIC** generates the initial overdensity field by sampling from a Gaussian distribution with zero mean and unit variance, and convolving the sampled field with the appropriate real-space transfer function. In the Lagrangian frame, the time-evolution of perturbations is handled by the displacement and velocity fields. This requires solving the Poisson equation, for which **MUSIC** uses a multigrid method (see Sec. 3.1.3).

The exact form of the initial perturbation is determined by a cosmology-dependent power spectrum. We use two dark-matter only simulations run using **GADGET-3** in this work. Both share identical initial conditions, with the cosmology set by results from *Planck* (Aghanim et al., 2020). The difference between the two simulations is their box sizes. We label them as **box100** and **box300** for box sizes of 100 and 300 h^{-1} Mpc respectively. The parameters of the two simulations are summarised in Table 3.1. Given the initial perturbations at $z = 99$ ($a = 0.01$), **GADGET-3** tessellates the simulation volume using a Peano-Hilbert curve. The large-scale force is calculated using a PM algorithm, and a hierarchical multipole expansion is used on

¹⁶<https://github.com/cosmo-sims/MUSIC2>

small-scales. A symplectic integrator then displaces the particles and updates their velocities. The simulation outputs snapshots which are equally spaced out in $\log a$. Between $a = 0.01$ and $a = 1$, sixty snapshots are generated. For assigning densities to the particles in a given snapshot, we implement a CIC-based interpolation scheme.

We run the AHF algorithm on the snapshots to generate a catalogue of halos for each snapshot. Each halo is characterised by the set of its constituent particles, which may be used to track the halo backward in time. They determine all halo properties such as mass, size, and position of the center. To calculate a halo’s COM, we choose one of its particles as the origin of a new frame of reference, and then shift the position of the other particles accordingly. This is necessary to treat halos which cross over the edge of the box to the other side due to periodic boundary conditions. Once the COM is calculated in the new frame, it is transformed to the original frame of the simulation.

Since our objective is to study the formation of halos, we must calculate their accretion history. We do this using AHF’s merger tree implementation. For each halo in a catalogue C_i , the merger tree algorithm identifies its major progenitor in the previous catalogue C_{i-1} , based on the merit function defined in Eq. (3.40). In this way, the algorithm works its way backward to the first halo catalogue. At the end of this process, each $a = 1$ halo can be traced back in time to $a = 0.01$, and consequently to its proto-halo patch in the initial conditions of the simulation. The accretion history of a halo is then a compilation of the masses of its parent proto-halo(s) at each scale factor.

The other necessary ingredient for our study in Chapter 5 is the linear deformation tensor associated with each proto-halo. To calculate this quantity, we start by smoothing the three-dimensional linear overdensity field (δ_L) with a real-space top-hat filter. The mass scales associated with the filter range from $10^{10} - 10^{16} h^{-1} M_\odot$ (for `box300`, we set the lower limit to $10^{10} h^{-1} M_\odot$). We integrate the Poisson equation for δ_L to obtain the linear deformation tensor field. To associate this with an individual halo, we interpolate this quantity onto the COM of its proto-halo at $a = 0.01$. The eigenvalues of this tensor are then calculated (the sum of the eigenvalues gives back δ_L at the COM). In Chapter 5, we attempt to use a neural network to find a map between these eigenvalues and the accretion history of the halo.

3.3. Artificial neural networks

This section provides a brief overview of artificial neural networks. We introduce the basic idea behind neural networks, and sketch the key algorithm that powers them. For comprehensive reviews of machine learning techniques applied in Physics and Astronomy, see e.g., Carleo et al. (2019); Mehta et al. (2019); Smith & Geach (2023).

An artificial neural network is machine learning model made up of discrete, interconnected nodes called *neurons*. We may think of a neural network (NN) as an arrangement of knobs in an electrical circuit, such that the output is an unknown, non-linear function of the positions of the knobs. The objective is to find the configuration of knob states that results in the least error in the output. This error is defined by a *loss function* appropriate to the task. A neuron is an individual knob, which is turned according to a parameter called its *weight*. An additional *bias* parameter offsets the input of each neuron to provide further flexibility. The neuron inputs are also called *features*. The process of fine-tuning the knob turnings to minimise the error is called *training*. A NN receives data during its training which allows it to iteratively arrive at an optimal configuration of weights and biases. A well-trained network can then respond to new data, and generalise to provide an output appropriate to the learning task.

Mathematically, a NN is a directed graph, i.e. a collection of vertices and directed edges such that

- each vertex i is characterised by a state variable (n_i) and a real-valued bias b_i ,
- an edge connecting vertices i and j carries a real-valued weight w_{ji} ,
- a non-linear function σ determines the future state of each node. Its arguments are the weights of each edge pointing to the node, the bias value of the node, and its current state.

Then, the state of a neuron at the next time-step may be written as

$$n_i(t_1) = \sigma \left[\sum_{j \neq i} w_{ji} n_j(t_0) - b_i \right] \quad (3.41)$$

In their earliest form, neurons in a network could either be ON or OFF, and this binary state variable was determined by comparing the output of the neuron with a fixed threshold value (McCulloch & Pitts, 1943). Given a specific task, the key problem was to identify the set of weights that would carry out the task. Rosenblatt (1958) coined the term *perceptron* for a hypothetical NN in the brain. A perceptron was thought of as a NN with two layers such that information would flow unidirectionally from the input layer to the output layer. Thus, they are *feed-forward* NNs. Rosenblatt and collaborators also devised an early algorithm for calculating the weights of a network (Rosenblatt, 1961; Block, 1962). It was pointed out that two-layer perceptrons could not solve some elementary problems, e.g. representing the exclusive-OR logic gate. It is now known that a three-layer feed-forward NN can solve this problem, although no efficient algorithm for finding the optimal weights existed at that time (Müller et al., 2012). The major break came with the discovery of such an algorithm. It was first published in a PhD thesis under the name *dynamic feedback* (Werbos, 1974), and later rediscovered (Rumelhart et al., 1986; Le Cun,

1986) and given its now-common name – *backpropagation*. We will describe this algorithm in the following section.

3.3.1. Backpropagation

Let us frame this problem in a more concrete way, and explain how the learning process works. Consider a multi-layer NN whose task is to minimise a user-specified loss function L . The output of a given layer of this network can be expressed as a linear combination of its inputs (analogous to Eq. 3.41)

$$\mathbf{y} = \sigma(W\mathbf{x} + \mathbf{b}) , \quad (3.42)$$

where the matrix W contains the layer weights, and the vector $\mathbf{b}$ contains the biases. The numbers in W and $\mathbf{b}$ constitute a high-dimensional hyperparameter space. The task of the network is to find the minimum of the loss function in this space. Ideally, this is the global minimum of the loss, but in practice it is likely to be a local minimum. Like all optimisation problems, we can start at a random point in the hyperparameter space, and iteratively move in the direction of the negative gradient, until we find a minimum. This technique called *gradient descent*. Let the starting set of weights and biases be denoted by $\{W(0), \mathbf{b}(0)\}$. Then, the new weights and biases are given as

$$W(1) = W(0) - \lambda \left. \frac{\partial L}{\partial W} \right|_{W(0)} , \quad (3.43a)$$

$$\mathbf{b}(1) = \mathbf{b}(0) - \lambda \left. \frac{\partial L}{\partial \mathbf{b}} \right|_{\mathbf{b}(0)} , \quad (3.43b)$$

where λ is a step-size, conventionally called the *learning rate*. The challenge for an efficient gradient descent implementation is an inexpensive evaluation of the derivatives in Eq. (3.43). This is exactly what backpropagation accomplishes. Let us consider a general element of the derivative matrix, and use the chain rule of derivatives to expand it (we ignore the activation function σ for the moment)

$$\begin{aligned} \frac{\partial L}{\partial W_{ji}} &= \frac{\partial L}{\partial y_l} \frac{\partial y_l}{\partial W_{ji}} \\ &= \frac{\partial L}{\partial y_j} x_i , \end{aligned} \quad (3.44)$$

where the second equality can be seen to hold recalling $y_l = W_{li}x_i + b_l$, and we renamed the dummy index l back to j . Similarly, for the biases,

$$\frac{\partial L}{\partial b_i} = \frac{\partial L}{\partial y_i} . \quad (3.45)$$

This can be trivially extended to include the activation function. Given the loss function, its derivative with respect to the output variable can often be calculated analytically. At each iteration of updating the weights, the gradients are computed inexpensively using Eqs (3.44) and (3.45). Given the output of the last layer, we find the gradients to update its weights and biases. Then, the output of the next-to-last layer updates *its* weights and biases, and so on, constituting a backward pass through the network. At each step, the calculation is simply a matrix multiplication, lending backpropagation its efficiency. With this relatively simple prescription, NNs with multiple layers can generalise to solve a large variety of problems. In a multi-layer NN, the first and last layers are the input and output layers respectively. The others are so-called *hidden layers*. Allowing an arbitrary number of hidden layers, NNs are *universal functions approximators*. More precisely, given a space of functions, there exists a sequence of NNs which can approximate any function in the space (e.g. Hornik et al., 1989).

3.3.2. Learning tasks

Broadly, learning tasks may be split in two categories – supervised learning, in which the dataset is labelled, and the expected output of the network is known, and unsupervised learning, which attempts to extract latent patterns from unlabelled data. The latter includes algorithms for dimensionality reduction, clustering, and tasks like image and text generation. Our attention in this work is focussed on supervised learning methods. These may further be split into classification and regression tasks. For a classification task, the NN must associate its input with a discrete class label. On the other hand, a NN designed for a regression task identifies the (possibly non-linear) relationship between the dependent and independent variables. As remarked upon in Chapter 1, our motivation is to learn the mapping between the initial overdensity field of a cosmological simulation, and the properties of the collapsed objects which form at late times. More precisely, if I_1 represents the initial proto-halo patch of halo 1, and H_1 represents its accretion history, we wish to find the map f such that $H_1 = f(I_1)$. This map is affected by the non-linear physics that governs halo formation, and is thus analytically intractable. We train a NN to find a numerical approximation for f . Ultimately, this approximation is mimicked by the sets of weights, biases, and activation functions and an expression for it is unattainable¹⁷. However, the utility of a well-trained network is that it is general enough for $H_i = f(I_i)$ to hold for a diverse set of halos. We will show in Chapter 5 that a simple NN with one or two layers is sufficient to find such a mapping.

¹⁷A special machine learning technique called *symbolic regression* may be used to obtain an analytical relationship between dependent and independent variables (e.g. Udrescu & Tegmark, 2020). However, the resulting expressions are often lengthy and unintuitive (see e.g. Delgado et al., 2022; Ondaro-Mallea et al., 2022 for related works).

3.3.3. Activation functions

For a classification task, the output of the neural network has to be bounded, conventionally between zero and one. The output layer has as many neurons as there are class labels, and each output is a probability. The weights and biases of a network do not have these boundedness properties. An activation function enforces such properties for each neuron. The function also serves another crucial purpose. Consider the expression for the output of a neuron (Eq. 3.42) without the activation function – the output is a linear combination of weights and biases. This intrinsic linearity would make the network incapable of modelling any non-linear behaviour. Thus, to model non-linear phenomena, one must use an activation function that makes the output of a neuron non-linear. Below, we highlight three commonly employed activations. In Fig. 3.2, we show their behaviour for uniform input in the range $[-5, 5]$. For recent surveys of activation functions and their comparative behaviour, see Apicella et al. (2021); Dubey et al. (2022).

- The sigmoid or logistic function is defined as

$$\text{sigmoid}(x) = \frac{1}{1 + e^{-x}} . \quad (3.46)$$

Thus, the sigmoid transforms its input to be bounded in the interval $[0, 1]$. This is ideal for the output layer of classification problems. However, for large input values (positive or negative), the gradient of the sigmoid vanishes. This vanishing gradient stalls out the learning process.

- The hyperbolic tangent is defined as

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} . \quad (3.47)$$

The tanh output is bounded in the interval $[-1, 1]$. Since it is centred at 0, it is often a better choice than sigmoid for the hidden layers. However, it also suffers from the vanishing gradient problem.

- Rectified linear unit, or ReLU (Nair & Hinton, 2010) activates a neuron for positive inputs, where its output scales linearly with the input, and is zero otherwise. Thus,

$$\text{ReLU}(x) = \frac{x + |x|}{2} . \quad (3.48)$$

ReLU has two key advantages over other activations. First, its zero for many neurons. This leads to sparse matrices, which reduces the computational overhead. Further, since the derivative of ReLU is either zero or one, the gradients




Figure 3.2. – An illustration of three common activation functions via their action on uniform input in the range $[-5, 5]$.




Figure 3.3. – An illustration of the comparison between training and validation losses. The three panels show three possible scenarios. *Left*: Underfitting – both training and validation losses remain high. *Center*: Good fit – both losses decrease and then plateau, remaining close to each other. *Right*: Overfitting – both losses decrease similarly at first, then the validation loss starts to increase.

do not vanish as easily as in the sigmoid function¹⁸. For these reasons, ReLU is often preferred over other activation functions for modern networks.

3.3.4. Training a neural network

Once the input features are chosen for a specific learning task, we can generate the input and output data for the network. The latter contains the true values used to evaluate the network loss. Then, we can fix a preliminary architecture, which can be fine-tuned based on the results of the trained network. With this in hand, the training process proceeds as follows

¹⁸The problem of vanishing gradients still persists for ReLU if the input is negative. An example of a modified version of ReLU which mitigates this is LeakyReLU (Maas et al., 2013), which keeps the non-zero values, but scales them by a factor (typically 0.01).

1. The dataset is split into training, validation, test sets. The rule of thumb for the fraction of the split is 80/10/10, or 70/20/10. The choice usually depends on the size and quality of the full dataset. It is best practice to normalise the training and validation sets. This reduces the dynamic range of the dataset, and helps with faster convergence of the gradient descent algorithm.
2. The training input is fed to the network in batches, set by a user-defined *batch size*. A small batch size allows the network to learn from individual samples, while a large batch size allows for easier generalisation. Starting with a random instance of weights and biases, the network processes the input in a *forward pass*.
3. The output of the forward pass is compared with the truth, and a *training loss* is computed. Using the input, output, and the loss, the network makes a backward pass using backpropagation to update the weights and biases.
4. When the network has processed all the training samples in this way, an *epoch* is completed. The number of epochs is a hyperparameter that controls the length of the training process. After each epoch, the network processes the validation set, and calculates and stores a *validation loss*.
5. Once this process has repeated for the specified number of epochs, the training process stops. At this point, we can compare the training and validation loss to determine the convergence of the network¹⁹. Examples of possible cases of these plots (called *learning curves* or *loss curves*) are shown in Fig. 3.3. The ideal loss curve is one where the training and validation losses both decrease with epoch, then plateau as the network settles into a minimum.
6. Finally, the trained network is used to make predictions using the test data which the network has never seen. If the performance is unsatisfactory, we may modify the network architecture, or change the hyperparameters until the predictions are up to the desired level of accuracy.

In this section, we have provided some useful background to understanding simple neural networks. We end the section by highlighting the implementation of neural networks that we use in Chapter 5. The primary module we exploit is an open-source `Python` implementation called `Keras` (Chollet et al., 2015). We define a class that inherits from the `Model()` class from `Keras` to customise our neural network. Although a simple sequential network is sufficient for this work, `Keras` also contains native implementation of convolutional and recurrent neural networks. Additionally, we use `scikit-learn` (Pedregosa et al., 2011) for preprocessing our data.

¹⁹Most implementations of neural networks also display the training and validation losses as they are calculated for each epoch.

CHAPTER 4

Testing the assumptions of the EFTofLSS

The following chapter is adapted from Karandikar et al. (2024). The pre-print of the paper is reproduced in full in Appendix D. Although this chapter relies on the background established in Sec. 2.2, it may also be read (almost) independently, using Sec. 2.2 to fill-in occasional gaps.

There are broadly two kinds of theoretical techniques to study structure formation: (i) using perturbative methods to accurately make analytical predictions on large (super-Mpc) scales, and (ii) using numerical simulations to make small-scale predictions. The EFTofLSS (e.g. Baumann et al., 2012; Carrasco et al., 2012, 2014b) is a hybrid approach that attempts to push the validity of perturbative methods to smaller scales by using information from small-scale numerical simulations. In Section 2.2, we discussed the theoretical underpinnings of the EFTofLSS. Specifically, we constructed a bottom-up, perturbative EFT for the evolution of the dark matter field. In this chapter, we use numerical simulations to rigorously test this construction. Using `cosmo_sim_1d`, we run a series of one-dimensional cosmological simulations to compare the bottom-up and top-down predictions of the EFTofLSS. We will start by describing the setup of these simulations in the following section.

4.1. Simulation setup

To explore the relationship between the large and small scales, we consider a universe populated by two fluctuations – a low-amplitude, long-wavelength perturbation, and a high-amplitude, short-wavelength perturbation. This serves as a simple toy model for our real Universe, which has more power on small scales. The resulting linear

Simulation	A_i	$k_i [k_f]$
<code>sim_1</code>	0.05	1
<code>sim_1_7</code>	0.05, 0.5	1, 7
<code>sim_1_11</code>	0.05, 0.5	1, 11
<code>sim_1_15</code>	0.05, 0.5	1, 15
<code>sim_1_6_15</code>	0.05, 0.25, 0.5	1, 6, 15

Table 4.1. – Details of the simulations used to test the EFTofLSS in this chapter. For each simulation, we have eight runs, such that $\varphi_1 = \pi$, and φ_2 is uniformly drawn from $[0, 2\pi)$.

overdensity is a superposition of two cosine waves. Within a periodic simulation box of comoving size L , this is given as

$$\delta_L(x, a = 1) = A_1 \cos(k_1 x - \varphi_1) + A_2 \cos(k_2 x - \varphi_2) , \quad (4.1)$$

where the A_i are positive, real-valued amplitudes, the k_i are integer multiples of the fundamental wavenumber $k_f = 2\pi/L$, and the $\varphi_i \in [0, 2\pi)$ are the phases. We have assumed – without loss of generality – that the planar symmetry leads to all relevant dynamics occurring along the x -direction. The overdensity contains four Fourier modes, viz. $\pm k_1, \pm k_2$, whose amplitudes are paired to keep δ_L real-valued. For our primary simulation, we use the following perturbations – ($A_1 = 0.05, k_1 = k_f$), ($A_2 = 0.5, k_2 = 11 k_f$). The phase $\varphi_1 = \pi$, and we consider eight realisations of the simulation with φ_2 uniformly distributed in $[0, 2\pi)$. This gives us an ensemble of simulations for which the average over the linear, short-wavelength modes is zero. We call this simulation `sim_1_11` (for the wavenumbers in units of k_f). We run additional simulations with different configurations which we list in Table 4.1. One notable entry is `sim_1_6_15`, where we insert an additional perturbation at $6 k_f$ to assess its impact on the effective speed of sound (see Sec. 4.4). We run these simulations from $a = 0.5$ to $a = 6$. Although $a = 0.5$ might appear to be a late starting point, it is perfectly valid in one-dimension. This is because 1LPT is exact in one-dimension before shell crossing, and we may start our simulations at such a late time. In Fig. 4.1, we show the evolution of relevant dynamical fields in `sim_1_11`. From left to right (respectively), we plot the overdensity, mean velocity, and velocity dispersion fields. Each row corresponds to a specific scale factor. At $a = 0.5$, the overdensity peaks show the imprint of the long-wavelength perturbation – the centremost peak is the highest, and peak height reduces on either side. Shell crossing in the central peak occurs around $a = 1.82$, and subsequently in each peak. After this point, velocity dispersion becomes non-zero (right column), and 1LPT (the Zeldovich approximation) predicts that the streams keep expanding. In reality, gravity overcomes the momentum of the streams, and they contract and re-collapse. This leads to the complex phase-space structure as seen in the last two rows of Fig. 4.1.




Figure 4.1. – Evolution of dynamical fields in `sim_1_11`. *Left*: the overdensity field. *Center*: the mean velocity. *Right*: the velocity dispersion. Shell crossing first occurs at $a = 1.82$ (second row), followed by the formation of one-dimensional halos. In the right panel, ν indicates the fraction of particles located in multi-stream regions.

4.2. The non-linear scale, and effective stress

Our objective is to understand the evolution of the long-wavelength mode (k_f) using the EFTofLSS. An important choice we must make at this stage, is the value of the coarse-graining scale Λ . To aid in this choice, we visualise the evolution of the non-linear scale k_{NL} with time for `sim_1_11`. We estimate k_{NL} by the inverse of the mean (median) particle displacement in the simulation (relative to particle positions at $a = 0$). The result is shown in Fig. 4.2. We observe that $k_f \ll k_{NL}$ at $a < 6$. Thus, we can treat the Fourier modes with wavenumber $|k| = k_f$ perturbatively. Then, a robust choice of Λ would be much smaller than the short-wavelength mode $11 k_f$, but larger than k_f . For our primary results, we make the choice $\Lambda = 3 k_f$. The EFTofLSS would then be applicable as long as $\Lambda/k = 3$ and k_{NL}/Λ are close to each other. This means that we can be confident about our results until $k_{NL} = 9 k_f$, or $a = 3$. We caution the reader that this estimate is used merely as a heuristic, and in practice, our results remain consistent at later times. Further, we have verified their robustness to change in Λ up to $\Lambda = 8 k_f$.

Given the choice of Λ , we may filter our fields using a window function (see Eq. 2.45). In the EFTofLSS literature, both a sharp- k filter, and a Gaussian filter are commonly used. Some features of the derivation of the EFT equations use properties of a Gaussian filter (Baumann et al., 2012). However, the choice of filter




Figure 4.2. – Time evolution of k_{NL} for `sim_1_11`. Here, k_{NL} denotes the wavenumber corresponding to the mean (solid blue line) or median (dash-dotted red line) particle displacement. We indicate three important scales with horizontal dashed lines: $k = k_f$ is the scale of interest in our work, $k = 3k_f$ is the coarse-graining scale we mostly use, and $k = 9k_f$ is the smallest k_{NL} which satisfies the ordering of scales $k \ll \Lambda \ll k_{\text{NL}}$. The vertical dashed line highlights the epoch of first shell crossing.

should not affect our results. The curious reader is encouraged to refer to K24 for details on the comparison between the two filter choices. For brevity, we will only consider the sharp- k filter in the rest of this work. As in Sec. 2.2, we indicate the coarse-grained fields by the subscript ‘ ℓ ’. Now, let us construct the effective stress in one-dimension. The planar symmetry of our setup means that the moments of the Vlasov distribution discussed in Sec. 2.2 have a unique scalar component. More precisely, $U^i = U\delta_1^i$, $\tau^{ij} = \tau\delta_1^i\delta_1^j$, and $\Xi_\ell^{ij} = \Xi_\ell\delta_1^i\delta_1^j$. Thus, we can safely refer to each of these by the scalar component, and Eq. (2.50) has the following one-dimensional analogue

$$\tau = \underbrace{\Xi_\ell - \rho_\ell U^2}_{\tau_k} + \underbrace{\frac{1}{8\pi G a^2} \left\{ [(\delta_x \phi)_\ell^2] - (\delta_x \phi_\ell)^2 \right\}}_{\tau_g}, \quad (4.2)$$

where τ_k and τ_g are the kinetic and gravitational components of τ respectively. The kinetic component captures the effect of smoothing on the velocity dispersion of the effective fluid, while the gravitational part quantifies departures from mean-field acceleration. To understand how these components (and indeed the total effective stress) evolve with time, we plot all terms of Eq. (4.2) at three different scale factors in Fig. 4.3. We further split Ξ_ℓ into contributions from particles which are in the single-stream (s) or multi-stream (m) regions. It is instructive to estimate the linear




Figure 4.3. – Decomposition of the effective stress in `sim_1_11` into its constituent pieces. The superscript over Ξ_ℓ indicates whether the contribution is made by particles within single-stream (s) or multi-stream (m) regions.

theory prediction for the evolution of τ . Since the peculiar gravitational potential ϕ is independent of time, τ_g scales as a^{-2} . The peculiar velocity U scales as $a^{1/2}$, and the density as a^{-3} . Thus, the kinetic part also scales as a^{-2} . To extract the non-linear evolution from our measured τ -components, we scale the linear theory contribution out by dividing by a^2 . Before shell crossing (left panel), the multi-stream contribution is zero, and the kinetic part is larger than the gravitational part. Over time, the contributions from the two components are equalised. The total non-linear stress becomes smaller with time, and develops oscillatory behaviour as power leaks into high- k modes.

4.3. EFTofLSS estimators

Our construction of the effective stress in the EFTofLSS has been bottom-up. To test the consistency of the assumptions we have made along the way, we compare this construction with top-down calculations within the EFTofLSS. To make these calculations using the effective stress, the key expression is the one-dimensional analogue of Eq. (2.52), viz.

$$\tau = \langle \tau \rangle + J \approx p_b + \rho_b c_s^2 \delta_\ell + \rho_b c_v^2 \theta_\ell + \dots + J, \quad (4.3)$$

where the Wilson coefficient is a combination of the effective speed of sound c_s^2 and the viscosity coefficient c_v^2 . In Sec. 2.2, we highlighted the practical difficulties in constructing the ensemble-averaged quantity $\langle \tau \rangle$. The gist of the trouble is that the expectation value must be taken over the short-wavelengths at fixed long-wavelength

fields. As the wavelengths are always coupled to each other, it is impossible to build such an ensemble. We suggest two ways to skirt around this difficulty, corresponding to two classes of estimators we define in this section

- One idea is to approximate the ensemble average with a spatial average (Baumann et al., 2012; Carrasco et al., 2012; McQuinn & White, 2016). This allows us to use Eq. (4.3) to extract c_s^2 and c_v^2 (see Sec. 4.3.1).
- Alternatively, it is possible to create an ensemble such that the *linear* long-wavelength modes are fixed. Then, we may average over the short-wavelength modes to obtain an approximation to $\langle \tau \rangle$. In practice, we use the eight realisations of `sim_1_11`, measure τ in each, and average over the eight measurements to get $\langle \tau \rangle$. Then, we use least-squares fitting to estimate c_s^2 and c_v^2 (see Sec. 4.3.2).

4.3.1. Estimates from spatial cross-correlations

For this class of estimators, we replace the ensemble average with a spatial average $\langle \dots \rangle_x$. Assuming that the stochastic term J is uncorrelated with the long-wavelength fields, we may take correlations of τ with δ_ℓ and θ_ℓ to get

$$\langle \delta_\ell^2 \rangle_x c_s^2 + \langle \theta_\ell \delta_\ell \rangle_x c_v^2 = \rho_b^{-1} \langle \tau \delta_\ell \rangle_x , \quad (4.4)$$

$$\langle \theta_\ell \delta_\ell \rangle_x c_s^2 + \langle \theta_\ell^2 \rangle_x c_v^2 = \rho_b^{-1} \langle \tau \theta_\ell \rangle_x . \quad (4.5)$$

An important consideration is that the above equations are degenerate at early times. This is because $\delta_\ell \simeq -\theta_\ell$ in the linear and quasi-linear regimes. Then, there is only one dynamical variable. We have previously defined this as $c_{\text{tot}}^2 = c_s^2 - c_v^2$. Nonetheless, the coefficient matrix of the above equations is invertible at $a \geq 0.5$, and we can individually estimate c_s^2 and c_v^2 . We label this estimate as ‘SC’. To verify that this does not lead to any numerical errors, we also combine Eqs (4.4) and (4.5) to get an estimator for c_{tot}^2 which we call ‘SC δ ’

$$c_{\text{tot}}^2 = \frac{1}{\rho_b} \frac{\langle \tau \delta_\ell \rangle_x}{\langle \delta_\ell^2 \rangle_x} . \quad (4.6)$$

A similar estimator in Fourier-space was presented in McQuinn & White (2016), where

$$c_{\text{tot}}^2 = \frac{1}{\rho_b} \frac{\sum_{k < \Lambda} \tilde{\tau}(k) \tilde{\delta}_{L,\ell}(-k)}{\sum_{k < \Lambda} |\tilde{\delta}_{L,\ell}(k)|^2} , \quad (4.7)$$

where $\delta_{L,\ell}$ is the coarse-grained, linear overdensity. We also estimate c_{tot}^2 using this ‘M&W’ estimator as a comparison.




Figure 4.4. – Fits to the averaged effective stress at three different times. The solid blue line is the stress estimated from the simulations, and the dashed lines represent the fits using the F3P (black) and the F6P (red) models.

4.3.2. Estimates from fitting

The eight runs of `sim_1_11` differ only in the phase of short-wavelength mode, and the linear, long-wavelength mode is fixed. Thus, taking an average of τ over the short-wavelength modes gives us an approximation for $\langle \tau \rangle$. Given this approximation, we may fit Eq. (4.3) for the coefficients. At linear-order in δ_ℓ , we get

$$\langle \tau \rangle = C_0 + C_1 \delta_\ell + C_2 \theta_\ell, \quad (4.8)$$

where $\langle \tau \rangle$, δ_ℓ , and θ_ℓ are all obtained by averaging over the eight simulations. We fit the above equation for the C_i using a least-squares fit, i.e. by minimising the squared difference between the measured $\langle \tau \rangle$ and the fit. We refer to this estimator as ‘F3P’ (short for ‘fit using three parameters’). As a consistency-check, we also perform a fit up to second-order in δ_ℓ , which includes six parameters (F6P). We show the accuracy of these fits in Fig. 4.4. The F3P fit is able to reproduce the simulation measurement reasonably well, except for the last epoch, where the higher-order terms in the expansion become important. At this time, F6P continues to be a good fit. Using the estimated C_i , we can infer the values of the EFT coefficients. As before, C_0 is a background pressure that has no dynamical consequences, and C_1 and C_2 are degenerate at early times. The only dynamically important parameter is the combination $c_{\text{tot}}^2 = \rho_b^{-1}(C_1 - C_2)$. In the following section, we compare the values of c_{tot}^2 obtained from the top-down estimators we have discussed.

4.4. Effective speed of sound

The individual estimates of c_s^2 and c_v^2 are instructive in their own right. The curious reader may find them in Fig. 8 of K24, with a brief discussion about their relation-




Figure 4.5. – Evolution of c_{tot}^2 with time for `sim_1_11` using the top-down estimators. The dash-dotted line corresponds to the scaling $c_{\text{tot}}^2 \propto a$ generally assumed in applications of the EFTofLSS. The dashed vertical lines are instances of shell crossing as detected in the simulation snapshots.

ship. We omit their inclusion here, and instead turn directly to c_{tot}^2 . Fig. 4.5 shows the evolution of c_{tot}^2 with time for `sim_1_11`. The qualitative behaviour of all five estimators is the same. The speed of sound increases linearly with a before shell crossing – as expected from perturbation theory. As the first set of shell crossings occur (indicated by vertical grey lines), the value of c_{tot}^2 decreases from a peak, and continues falling until the end of the simulation. Two distinct groups of estimators emerge – estimates from M&W and $\text{SC}\delta$ are consistently higher than that from F3P, F6P, and SC. This may be explained by noting that the first group uses only δ_ℓ for their predictions, while the second group uses both δ_ℓ and θ_ℓ . The additional information from θ_ℓ translates to a lower value of c_{tot}^2 .

Now, let us try to understand this behaviour of c_{tot}^2 . As a first interpretation, c_{tot}^2 may be sourced by perturbations that are entering the non-linear regime. This has already been suggested, most notably in McQuinn & White (2016), but never verified. The peak around $a \approx 2$ signals that the short-wavelength mode ($11 k_f$) has entered the non-linear regime. After this epoch, c_{tot}^2 continues to drop as there is no new perturbation that enters the non-linear regime by the end of the simulation. To concretely test this hypothesis, let us first see the long-term evolution of the perturbations in `sim_1_11`. As we let the system evolve to later times, we expect long-wavelength modes to eventually become non-linear. Note that although no modes other than k_f and $11 k_f$ are present in the linear overdensity, mode coupling makes the intermediate modes non-zero. We extend the run of `sim_1_11` to $a = 12$




Figure 4.6. – Long-term evolution of c_{tot}^2 for `sim_1_11`. The vertical lines are as defined in Fig. 4.5. A second peak appears around $a = 12$, likely as a consequence of the $k = 3 k_f$ mode entering the non-linear regime.

to see if these modes influence c_{tot}^2 . The resulting¹ c_{tot}^2 is shown in Fig. 4.6. The appearance of a second increase is clearly seen before $a = 12$. The top axis shows the value of Λ/k_{NL} , and we can see that this ratio exceeds unity shortly after $a = 10$. The increase in c_{tot}^2 may therefore be attributed to the $\Lambda = 3 k_f$ mode entering the non-linear regime.

As a final test, we introduce a third perturbation at $k = 6 k_f$ with a linear amplitude half that of the short-wavelength perturbation. To see the effect of this more clearly, we push the latter out to $k = 15 k_f$. The estimated c_{tot}^2 from this simulation (`sim_1_6_15`) is shown in Fig. 4.7. The presence of the intermediate mode introduces a clear peak around $a = 6.75$, shortly after the ratio $6 k_f/k_{\text{NL}}$ exceeds unity. The subsequent increase in c_{tot}^2 as the ratio exceed two (i.e., as the $3 k_f$ mode enters the non-linear regime) is also visible here. Once the modes are in the non-linear regime, they make no further contribution to c_{tot}^2 , as evidenced by the valleys between the three peaks. This confirms our hypothesis that c_{tot}^2 is sourced only by modes which are entering the non-linear regime.

¹Two points may be noted here: (i) the F6P estimator makes unphysical predictions at late times, and we omit it from this and the next figure, and (ii) the M&W estimators fails as it assumes linear theory, which is no longer valid in this regime.




Figure 4.7. – Evolution of c_{tot}^2 with time for `sim_1_6_15` using the top-down estimators. The vertical lines are as defined in Fig. 4.5. Addition of a third, intermediate mode introduces a clear peak when the mode enters the non-linear regime around $a = 6.75$.

4.5. The EFTofLSS power spectrum

In Sec. 2.2.2, we discussed the power spectrum in perturbation theory up to next-to-leading order. Specifically, Eqs (2.58) and (2.59) quantified the correction to the SPT power spectrum induced by the effective stress, captured in the term α_c . Our top-down estimates of c_{tot}^2 may be integrated numerically in accordance with Eq. (2.59) to calculate these corrections. For the integration, we assume that $c_{\text{tot}}^2 \propto a$ until $a = 0.5$ where our simulation starts, and we use the measured c_{tot}^2 thereafter. Additionally, we can equate the EFT and N -body power spectra to get a bottom-up estimate of α_c as follows

$$\hat{\alpha}_c = \frac{P_{N\text{-body}} - P_{\text{tSPT}}}{2k^2 \widetilde{W}_\Lambda^2 P_{\text{lin}}}. \quad (4.9)$$

This estimate is guaranteed to match the simulation results, and is therefore exact. If, instead of measuring c_{tot}^2 from simulations, we use its linear-theory estimate, we can integrate Eq. (2.59) analytically to get

$$\alpha_c = -\frac{c_0^2 a^2}{9H_0^2}, \quad (4.10)$$

where c_0^2 is the estimated value at $a = 1$. We are now in a position to compare these three distinct classes of estimates for α_c . We show this comparative evolution of α_c in Fig. 4.8 (scaled by k^2 to make the quantity dimensionless). The first point of




Figure 4.8. – Comparison of α_c between the bottom-up estimator (solid green line) and the top-down estimators (dashed lines). The dashed vertical lines are as in Figure 4.5.

note is that α_c is consistently negative. This implies that the power spectrum will be dampened relative to the tSPT prediction. At first, the linear-theory estimate (dark grey, dash-dotted line) agrees with both the top-down and bottom-up estimates, which are clustered together – $|\alpha_c|$ grows as a^2 . This persists until well-after shell crossing, unlike the corresponding behaviour in Fig. 4.5. The reason for this is that α_c is the integral of c_{tot}^2 over the past scale factors, and therefore retains the memory of the linear regime. Eventually, this effect is overcome, and the a^2 behaviour no longer holds. The top-down estimates agree with the exact bottom-up prediction up to, and well after $a = 3$ – the epoch at which our naive estimate of k_{NL} indicated that the EFTofLSS may fail. This lends credence to the veracity of the assumptions we made in setting up the theory.

Let us now turn to the next-to-leading power spectrum. First, let us examine the predictions from SPT. The perturbative solution to the perfect-fluid equations is conventionally obtained in Fourier space. In one-dimension, the overdensity takes the following form (McQuinn & White, 2016, K24)

$$\tilde{\delta}(k, a) = a \tilde{\delta}_1(k) + a^2 \tilde{\delta}_2(k) + a^3 \tilde{\delta}_3(k) + \dots, \quad (4.11)$$

with

$$\tilde{\delta}_n(k) = \int F_n(k_1, \dots, k_n) \tilde{\delta}_1(k_1) \dots \tilde{\delta}_1(k_n) \delta_{\text{D}}\left(k - \sum_{i=1}^n k_i\right) dk_1 \dots dk_n, \quad (4.12)$$

4. Testing the assumptions of the EFTofLSS




Figure 4.9. – Evolution of the non-linear power in the $k = k_f$ mode for the system simulated by `sim_1_11`. *Left:* The measurement from `sim_1_11` (solid blue line) compared against the SPT (dotted pink line) and tSPT (dash-dotted brown line) predictions. The vertical lines indicate the time of first shell crossing (dashed teal line), and the time at which the short-wavelength mode becomes non-linear (dotted magenta line). *Right:* The N -body power compared with the five top-down EFT estimates. The bottom panels show the percentage deviation of the theoretical predictions from the measured power. Note the change in the y -axis scale, and that we omit tSPT in the bottom right panel to better visualise the EFT errors. Here, all vertical lines indicate shell crossings.

such that $a\tilde{\delta}_1$ is the solution to the linearised fluid equations. Here, the function F_n quantifies the coupling between Fourier modes. At $k = k_f$, the next-to-leading order power spectrum in SPT is given as

$$P_{\text{SPT}}(k_f) = P_{11}(k_f) + 2P_{13}(k_f), \quad (4.13)$$

where $P_{ij}(k) = a^{i+j}\tilde{\delta}_i(k)\tilde{\delta}_j(-k)$.

In the left panel of Fig. 4.9, we show the evolution of power in the $k = k_f$ mode. To isolate the effect of non-linear evolution, we scale out the linearity by dividing by a^2 . SPT makes accurate predictions even after the short-wavelength mode enters the non-linear regime – the epoch indicated by the dotted magenta line. This implies that the coupling between the long- and short-wavelengths is weak at this point. Around $a = 3$, the SPT power starts to deviate from the N -body measurement, and by $a = 6$, the absolute percentage deviation (bottom panel) is $\approx 2.2\%$. In Eq. (4.12), one conceptual deficiency of SPT is made explicit. The integral is over *all* Fourier modes – even those with $k > k_{\text{NL}}$. These modes are, by definition, non-perturbative, and to make progress we must truncate the integral to $k = \Lambda$. This truncated form

of SPT (tSPT) eliminates the effect of the short-wavelength mode. This is shown by the dash-dotted brown line in Fig. 4.9. While this truncation rids us of the conceptual problem of non-perturbative modes, tSPT consistently overpredicts the power, and is less accurate than SPT until $a \approx 5$. At later times, the tSPT error is lower than SPT, with the percentage deviation at $a = 6$ being $\approx 1.5\%$.

In the right panel of Fig. 4.9, we correct the tSPT power using the top-down EFT estimators via Eq. (2.58). Overall, the EFT power shows excellent agreement with the N -body power. At $a = 6$, even the worst-performing EFT estimators (M&W, SC δ) agree with the N -body power to within 0.1%, a remarkable improvement over the SPT predictions. With these results, we can confidently claim that the EFTofLSS successfully models the coupling of long- and short-wavelengths within our setup. There are quantitative differences among the top-down estimates – as expected from the differences in c_{tot}^2 – but these differences are minor when compared to the SPT predictions. An important subtlety bears mentioning at this point. Although the late-time EFTofLSS predictions are much more accurate than their SPT counterpart, they lack one appealing feature. Namely, unlike SPT in the low- a limit, the EFT prediction is never asymptotically correct in any regime. Furthermore, it is unclear if this can be fixed at higher-orders.

Having established that the EFTofLSS provides an accurate match to the N -body power, it is interesting to examine how this accuracy changes when the simulation setup is altered. We have been examining the coupling of the long- and short-wavelength modes, and specifically how the presence of the former alters the dynamics of the latter. The logical next step in this exercise is to vary the separation between the two modes, and measure the impact of this on the accuracy of the EFT predictions. To that end, we run three separate simulations alongside `sim_1_11` – a reference simulation with only the long-wavelength mode (`sim_1`), and two simulations with the same amplitude ratio as `sim_1_11`, but with smaller and larger mode separations (`sim_1_7` and `sim_1_15` respectively). The resulting power spectra are shown in Fig. 4.10. The left panel shows the N -body and SPT power (solid and dashed lines respectively) for the four simulations. For `sim_1`, the SPT predictions remain accurate throughout the run, as the only linearly non-zero mode has a small amplitude, and thus grows slowly. Adding a short-wavelength mode introduces a feedback effect on $k = k_f$, and the SPT prediction deviates from the measured N -body power. The degree of deviation is more when the two modes are closer to each other, as the influence of the high-amplitude, short-wavelength mode grows with decreasing separation. When the integral of Eq. (4.12) is truncated, all SPT predictions collapse to the one without the short-wavelength mode (right panel). The tSPT error, therefore, is larger for lower separation. The right panel further shows the EFT prediction for each simulation. Specifically, in the bottom panel, we can clearly see the order-of-magnitude improvement that the EFT correction brings. We only show the SC prediction to avoid overcrowding the plot – the other




Figure 4.10. – As in Fig. 4.9, but for simulations with differences in separation between the long- and short-wavelength perturbations. The colours in both panels refer to the different simulations. *Left:* Comparison between the SPT and N -body power. *Right:* Comparison between tSPT, N -body, and EFT power. Note that the tSPT prediction for all simulations is identical.

estimators are in good agreement with the SC estimate. This result demonstrates that the EFTofLSS predictions are robust to change in separation of scales.

In this chapter, we have scrutinised aspects of the EFTofLSS in a simple setting. By comparing the top-down and bottom-up flavours, we find that the EFTofLSS makes reasonable assumptions that hold up to our scrutiny. We summarise our results and place them in the broader landscape of theories of structure formation in Chapter 6.

CHAPTER 5

Predicting halo formation using a neural network

We introduced the theory of excursion sets in Sec. 2.3, along with the physical models of spherical and ellipsoidal collapse. Taken together, these allow us to make analytical predictions about the formation and growth of halos. Specifically, we may predict whether a Lagrangian patch within an N -body simulation will collapse to form a halo, and – if it *does* collapse – predict its accretion history. In this chapter, we exploit the power of neural networks to make the same predictions. For a three-dimensional N -body simulation, we train a neural network to predict the accretion history of halos, given properties of the initial proto-halo patch. Along the way, we aim to answer the following questions

1. What is the optimal method (analytical or numerical) to recover the accretion history of a halo from the properties of the proto-halo?
2. How well do the different candidate methods predict *derived* properties of the halo population (e.g., formation time, mass function)?
3. What property of the proto-halo provides most information about its subsequent evolution?

In Sec. 5.1, we describe the setup of our neural network. The following four sections detail our results, as we compare the predictions made by the analytical model against that of the neural network.

5.1. Network properties

5.1.1. Input features

In Sec. 2.3, we raised the question of determining which quantity best captures the impact of environment on halo formation. In the language of neural networks, this translates to asking which input features lead to the best prediction for the accretion history. In our baseline model, the input features are the linear overdensities δ_L evaluated at fifty smoothing scales. We call this model `delta_box100` (`delta_box300`) for the 100 (300) h^{-1} Mpc simulation. As the next point of comparison, we use all three eigenvalues λ_i as input features, instead of using δ_L . We call this model `evals_box100` (or `evals_box300`). For our primary results, we compare the performance of the two `box100` models (Secs 5.2 and 5.3). We extend our analysis to a higher mass range using the `box300` models (Sec. 5.5). Additionally, we use the tidal anisotropy parameter α defined in Eq. (2.98) as an input feature along with δ_L . This model is labelled as `alpha_box100`. Comparing the three `box100` models allows us to answer the third question raised at the beginning of the chapter, and we discuss this in Sec. 5.4.

5.1.2. Architecture

Our objective is to find the mapping between the input features and the accretion history for each halo. At the beginning of the learning process, the output of the neural network is a non-linear combination of its input features, whose exact form is given by the initial sets of weights, biases, and activations. As the network trains, it attempts to minimise the loss function. For a regression task, the mean squared error between the predicted and true accretion histories is a good choice for the loss function, and it is defined as

$$L = \frac{1}{N} \sum_{i=1}^N \left(H_i^{\text{pred}} - H_i^{\text{true}} \right)^2, \quad (5.1)$$

where H_i^{pred} (H_i^{true}) is the predicted (true) accretion history for halo i , and the sum is over all halos in the training or validation set. The training samples are fed to the network in batches, whose size is a hyperparameter. As mentioned in Sec. 3.3.4, it is useful to normalise the training and validation data to improve the stability of the training process. We standardise the data (i.e. reduce it to zero mean and unit variance) using the preprocessing module from `scikit-learn`. Further, in order to combat any vanishing/exploding gradients, we add a ‘Batch Normalisation’ (Ioffe, 2015) layer following the input layer. Further, we add a dropout layer after the hidden layers. This randomly de-activates a small fraction of neurons in every pass, ensuring that the network predictions do not suffer from outliers. Finally, the output layer returns the predicted accretion history, which is then used to calculate




Figure 5.1. – The architecture of the neural networks for `delta_box100` (top), and `evals_box100` (bottom). The two architectures differ only in the input layer size.

the loss function and initiate backpropagation. We run comparisons among the ReLU, sigmoid and tanh activation functions, and find that this choice changes the shape of the output distribution. However, the impact is insignificant compared to that due to other hyperparameters, and we fix ReLU as our activation function.

For our regression task, the depth of a neural network depends on the complexity of the relationship between the input and output variables. This complexity also governs the size of the network layers. Taken together, layer sizes and depth define the *capacity* of a neural network. Typically, a larger dataset requires a larger capacity for the network to converge. The optimal architecture may be found by varying the layer capacity in combination with other hyperparameters to find a configuration that minimises the prediction error. Since we require a relatively shallow network for our task, it is sufficient to optimise the network by hand. We find that using a single dense (fully-connected) layer is sufficient for the neural network to converge (see Sec. 5.1.3). We keep the two architectures identical, so that the impact of the input features may be directly discussed. Thus, the two networks differ only in the size of the input layer, which is 50 and 150 for `delta_box100` and `evals_box100` respectively. We visualise the full architecture of the two networks in Fig. 5.1.

5.1.3. Convergence

To monitor the convergence of the neural network, we observe how the training loss compares with the validation loss. In Fig. 5.2, we plot the loss curves for `delta_box100` (left panel) and `evals_box100` (right panel), for networks trained to 100 epochs. Note that both losses drop off from their starting values, and then plateau, remaining close to each other consistently. Thus, we may determine that the network has converged. The variance of the training (validation) loss is dictated by the batch size. A smaller batch size would lead to noisier loss curves than we see in Fig. 5.2, where our batch size is set to 128.




Figure 5.2. – Loss curves for models `delta_box100` (left) and `evals_box100` (right). They demonstrate the convergence of the two networks.




Figure 5.3. – The halo mass function for `box100` as predicted by the two primary models (dotted lines), and the true mass function of the test set (solid red line). We also show the EC prediction (dot-dashed line) for an excursion-set theory comparison.

5.2. Mass function and halo masses

We begin our survey of the neural network predictions by looking at the halo masses obtained from the two models. First, let us look at the halo population as a whole. The halo mass function (HMF) quantifies the number density of halos per unit mass, and informs us about the distribution of mass in the cosmological volume. To obtain the mass function, we bin the test-set of halos in the range $10^{12} - 10^{13} h^{-1} M_{\odot}$. The output of the neural network gives us the final mass of the test-set halos, from which we compute the mass function in the defined bins. In Fig. 5.3, we plot the HMF of the test set¹ (solid red line). We compare this against the HMF predicted by both

¹Contrasting with Fig. 2.1, it is important to note that we only consider the test set halos in this section. Since the number of test-set halos is $< 10\%$ of the total halo population in `box100`, the




Figure 5.4. – Halo masses predicted by the models vs the masses measured in the simulation. We bin the halos to get a sense of the underlying correlation. The Pearson correlation coefficients are indicated in the top-left corner of each subplot. *Top*: The neural network predictions. *Bottom*: Excursion-set theory predictions. The former have a smaller scatter, and higher r -values.

`delta_box100` and `evals_box100` (dotted lines in Fig. 5.3). Both models predict the mass function reasonably well. The largest discrepancy is shown by `delta_box100` for the bin centred at $10^{12} h^{-1}M_{\odot}$. As an illustration, we plot the excursion-set theory prediction using ellipsoidal collapse (EC). We calculate this by binning the halo masses predicted by EC in the same mass range ($10^{12} - 10^{13} h^{-1}M_{\odot}$). The EC prediction is similar to the neural network prediction, except for one mass bin where the former significantly overpredicts the mass function.

To investigate the halo masses further, we can compare the predicted and true masses of individual halos. In Fig. 5.4, we plot the predicted masses against the true masses of test set halos. To clarify the underlying correlation, we collect the halos in hexagonal bins, coloured by the number of halos in each bin. The top row

mass function in Fig. 5.3 is an order of magnitude smaller than the corresponding amplitude in Fig. 2.1.

shows results from the two neural networks, and the bottom row shows the SC and EC results. For brevity, we omit the corresponding Bo17 plot, which may be found in App. C. First, let us consider the two neural networks. Both `delta_box100` (top left) and `evals_box100` (top right) make accurate predictions for the halo masses, with a small scatter and a strong positive correlation, indicated by the Pearson correlation coefficient (r). Comparing between the two models, the `delta_box100` masses have a larger scatter. The additional information due to using all three eigenvalues as input (`evals_box100`) improves the correlation significantly. We also note that the hexagons closely follow the 1:1 correlation line (dashed red), especially for `evals_box100`. Next, let us examine the SC (bottom left) and EC (bottom right) predictions. SC consistently overpredicts the halo mass, suggesting that the barrier height of $\delta_c = 1.686$ is too low in practice. This trend was also spotted, e.g., in Sheth et al. (2001). Generalising to ellipsoidal collapse and a non-constant barrier leads to considerable improvement in the predictions. The r -value of the EC predictions is still slightly worse than `delta_box100`, even though the latter only uses δ_L as its input. To make a clear comparison, we keep the scale of the colour bars consistent across the four subplots. The excursion-set predictions are spread out over more hexagonal bins compared to the neural network predictions, as evidenced by the scarcity of bins with more than 30 halos.

5.3. Accretion history and formation time

The target statistic for the trained neural network is the halo accretion history, which allows us to track how the halos – individually, and as a population – gain mass over time. It is instructive to begin by analysing the distribution of halos according to their accretion history. To this end, we use a derived quantity – the 50% formation time (a_{50}). This is defined as the scale factor at which the proto-halo first acquires 50% of its final mass. Then, we can visualise the distribution of halos in this quantity to summarise how the halo population grows with time. In Fig. 5.5, we plot this distribution for halos in the test set (solid red line). Additionally, we show the predictions from the two primary neural network models (dotted lines), viz. `delta_box100` (light blue) and `evals_box100` (black). To make this plot (and subsequent plots of this type), we construct the histogram of halo counts binned in a_{50} , and estimate their density using a Gaussian kernel. The true distribution peaks at $a_{50} \approx 0.42$, and attains its median value at $a_{50} \approx 0.47$. A majority of halos form between $a_{50} = 0.4$ and $a_{50} = 0.7$. We can see a long tail in the high- a_{50} limit, whereas the distribution drops off much more quickly below $a_{50} = 0.4$. The predicted distributions peak near the median of the test set, but they are much more uniform, and do not replicate the high- a_{50} tail, or the low- a_{50} behaviour. This discrepancy is a consequence of two possible effects. Firstly, neural networks generally overpredict




Figure 5.5. – The distribution of test-set halos in formation time. We compare the predictions from the two neural network models (dotted lines) against the true distribution (solid red line). The predictions peak at the median value of the true distribution, but do not capture the underlying patterns either in the low- or high- a_{50} limit.

the majority examples², as this tends to lower the overall loss. The tails of the distribution (particularly at high- a_{50}) are undersampled relative to its bulk, leading to an imbalanced dataset. Secondly, the information that the network receives may not be enough to accurately predict the formation time of halos. This may be a limitation of the neural network, or of the input features.

If we are to draw robust conclusions about the halo formation process using our setup, we must nullify the bias in the neural network predictions. In machine learning, class imbalance has been extensively discussed in reference to classification problems (see e.g. Haixiang et al., 2017 for a review). The issue is less prevalent in regression studies, but has also been discussed, e.g. in Branco et al. (2016, 2017). Broadly, there are two approaches to better model imbalanced data: algorithmic modifications, and data preprocessing. In the former case, the regression algorithm is modified to remove the bias towards the majority samples. This may involve modifying the loss function to penalise certain predictions. The second approach keeps the algorithm unchanged, but modifies the training data to either undersample the over-represented group of examples, or oversample the under-represented examples. Since our dataset is relatively small, it is more practical to oversample the halos in the tails of the distribution in a_{50} . Specifically, we upweight the halos with $a_{50} < 0.4$ and $a_{50} > 0.7$ by a factor of two. This is applied to both the training and the validation set. During the training process, the network now processes more rare examples, and we expect a more balanced distribution when predicting over the test set.

²By this, we refer to the most common examples that lie near the median of the distribution.




Figure 5.6. – Comparison of the predicted and true distributions (solid red line) of test-set halos in formation time (a_{50}). *Left:* The distributions predicted by the two neural network models after oversampling. *Right:* The distributions predicted by excursion-set theory using spherical and ellipsoidal collapse.

We plot the distributions for the two models with oversampling in the left panel of Fig. 5.6. With oversampling, the predicted distributions provide a much better match to the simulation. The bulk of the true distribution is faithfully reconstructed, and the performance towards the tails is greatly improved. However, the high- a_{50} tail is still somewhat underrepresented in the predictions, especially beyond $a_{50} \approx 0.8$, where the true distribution shows a small bump. Similarly, neither model predicts halos with $a_{50} < 0.3$, while a few such halos exist in the simulation. Now, let us contrast these neural networks predictions with those using excursion-set theory, which are shown in the right panel of Fig. 5.6. The SC prediction (dashed teal line) is strongly biased relative to the true distribution, and favours late-forming halos. EC (dot-dashed dark blue line) peaks close to the true median, and predicts the bulk of the distribution reasonably well. However, the underprediction of halos with $a_{50} < 0.5$ is noticeable. Finally, Bo17 (dot-dashed magenta line) is skewed towards low- a_{50} , and the number of halos with $a_{50} > 0.35$ is damped compared to the simulation. Compared to the excursion-set predictions, the neural networks (post oversampling) are able to better reproduce the formation times of the halo population.

Now, we may examine the full accretion history of the halo population. To condense the individual histories, we take their median across the test set. In Fig. 5.7, we compare this to the median of the predicted histories using both the two neural networks (left panel, dotted lines), and the three excursion-set predictions (right panel, dashed and dot-dashed lines). In the simulation (solid red line), the median proto-halo grows quickly at low- a , and has gained half of its final mass around $a \approx 0.47$. After this, the growth slows down, and it smoothly accretes the rest of its mass. Both `delta_box100` (light blue) and `evals_box100` (black) replicate this pattern of growth remarkably well. The excursion-set models, on the other hand,




Figure 5.7. – Comparison of the predicted accretion history vs that obtained from `box100` (solid red line in both panels). We take the median of all halos in the test set to condense the information. *Left*: The two machine learning models (dotted lines). *Right*: The excursion-set predictions from spherical collapse (dashed teal line) and the two ellipsoidal collapse models (dot-dashed lines).

tend to underpredict the growth of proto-halos at very early times, but the subsequent predicted growth is much steeper compared to the simulation. SC (teal), in particular, strongly overpredicts proto-halo growth, and the median predicted halo mass is ≈ 2.5 times its simulation counterpart. This is consistent with our expectation from Fig. 5.4, where SC overpredicts the final mass of halos except in rare examples. In the EC (dark blue) and Bo17 (magenta) models, proto-halo growth – on average – stalls around $a = 0.5$, and their predictions tend to get closer to the simulation with time. Around $a = 1$, the two models are a good match to the simulation, as shown by their r -values for halo mass.

As a quantitative measure of the models, we calculate the coefficient of determination, which enumerates the fraction of variation in a target variable that can be explained by an independent predictor. For our accretion histories, we define two quantities – the sum of squares of residuals (SS_{res}), and the total sum of squares (SS_{tot}) as follows

$$SS_{\text{res}} = \sum_{i=1}^N (H_i^{\text{pred}} - H_i^{\text{true}})^2, \quad SS_{\text{tot}} = \sum_{i=1}^N (H_i^{\text{true}} - \bar{H})^2, \quad (5.2)$$

where $\bar{H}$ is the mean of the true accretion histories, and the sum is over all halos of the test set. Then, the coefficient of determination (R^2) is calculated as

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}. \quad (5.3)$$

When comparing different regression models, a larger value is more desirable, with the upper bound $R^2 = 1$ implying that the predictor explains all variability in the target. We summarise the R^2 values from the five models we have discussed so far

	SC	EC	Bo17	<code>delta_box100</code>	<code>evals_box100</code>
R^2 (MAH)	0.28	0.79	0.53	0.82	0.93

Table 5.1. – Comparison of model performance for `box100` in predicting the accretion histories. We use the coefficient of determination (R^2) as our metric. The highest value is bold for emphasis.

in Table 5.1. Both neural network models outperform the excursion-set predictions, with `evals_box100` achieving the highest R^2 -value. Among the excursion-set models, EC performs the best, and is only slightly worse than `delta_box100`. Both SC and Bo17 are significantly worse than the other three, with SC having the lowest R^2 .

Finally, it is instructive to compare the a_{50} predictions for individual halos. In Fig. 5.8, we plot the predicted and true formation times for all halos in the test set. We repeat the hexagonal binning of Fig. 5.4, and relegate the Bo17 plot to App. C as before. Let us start with the two neural network predictions (top row). The predicted and true a_{50} values show a positive correlation for both models, with r -values of ≈ 0.56 and ≈ 0.59 for `delta_box100` and `evals_box100` respectively. A notable observation is that neither correlation is as strong as in the case of halo mass. There are two possible reasons for this – either, a network trained for the accretion history struggles to predict the derived a_{50} values, or, the input features do not contain enough information to predict a_{50} with great accuracy. In order to test the former, we train another neural network with a_{50} as the target variable, and find no improvement in the correlation. This suggests that the eigenvalues of the tidal tensor are insufficient to accurately predict formation time. This is somewhat backed up by the excursion set predictions (bottom row), which also have a lower correlation compared to their final mass predictions. SC, in particular, has the lowest r -value. As in the case of halo mass, the additional ellipsoidal information in EC improves the correlation significantly.

5.4. Effect of tidal anisotropy

One of our objectives in this study is to identify the tidal property of the proto-halo that has the most predictive power for the accretion history of the halo. We defined these properties in Sec. 2.3.2. Among them, tidal anisotropy (α ; Paranjape et al., 2018) is particularly interesting, as it contains complementary information to δ_L . So far in this chapter, we have used either δ_L , or the three eigenvalues λ_i as input features for our network. In this section, we will see how the predictions from α compare against those discussed earlier. First, we designed a network to accept α as the only input feature. However, this network fails to accurately predict the accretion history. Thus, we use both δ_L and α as input features and assess how this model – named `alpha_box100` – fares in comparison to `delta_box100` and




Figure 5.8. – Predicted vs true formation times (a_{50}), with halos collected in hexagonal bins. We also indicate the corresponding Pearson correlation coefficient (r) in the top-left corner of each subplot. *Top:* The two neural network predictions – `delta_box100` (left) and `evals_box100` (right) – have relatively high r -values. *Bottom:* SC (left) and EC (right) are poorer predictors of a_{50} , as expected from the distributions in Fig. 5.6.

`evals_box100`. With the same architecture as the latter models, the network fails to converge. To get meaningful predictions for `alpha_box100`, we add a second hidden layer with 64 neurons to the original network. The updated architecture is shown in Fig. 5.9. In order to avoid overfitting, we stop the training process after 80 epochs. Let us discuss the predictions made by this trained network, with special emphasis on how they compare to `delta_box100`.

First, let us consider the masses of the halo population. In Fig. 5.10, we plot the properties of the mass distribution. In the left panel, we plot the mass function obtained by binning the predicted mass in the $10^{12} - 10^{13} h^{-1} M_{\odot}$ mass range. We compare the `alpha_box100` prediction (dotted indigo line) to the `delta_box100` prediction (dotted light blue line). For the lowest mass bin, the amplitude of the `alpha_box100` prediction is closer the simulation result. However, for the rest of the bins, `delta_box100` remains closer to the simulation. Despite this,

5. Predicting halo formation using a neural network




Figure 5.9. – Architecture of the neural network for `alpha_box100`. We add a second dense layer with 64 neurons to the existing architecture of Fig. 5.1.




Figure 5.10. – Prediction of halo masses using the `alpha_box100` model. *Left:* The predicted mass function (dotted indigo line) vs that measured from `box100` (solid red line). We also show the `delta_box100` (dotted light blue line) HMF as a comparison. *Right:* Predicted vs true halo masses with hexagonal binning. The Pearson correlation is indicated in the top-left corner.

the `alpha_box100` predictions of individual halo masses have a lower scatter than their `delta_box100` counterpart. This is shown in the right panel of Fig. 5.10, where the indicated correlation ($r \approx 0.86$) is slightly higher than the $r \approx 0.83$ for `delta_box100` (see Fig. 5.4). This apparent discrepancy may be explained by the fact that the lower-mass bins contain a much larger fraction of halos than the higher-mass bins where the `alpha_box100` predictions are most discrepant.

For completeness, let us briefly consider the accretion histories predicted using `alpha_box100`. We calculate the coefficient of determination to be $R^2 \approx 0.78$. Using Table 5.1 as a reference, this is worse than `delta_box100`, and close the corresponding EC value.

As we did in the previous section, let us visualise the distribution of halos in formation time. We oversample the training and validation data in the same way as before, giving higher weight to the tails of the a_{50} distribution. In the left panel of Fig. 5.11, we compare the result (dotted indigo line) to the distribution obtained from the simulation (solid red line), and to `delta_box100` (dotted light blue line). The distribution predicted by `alpha_box100` peaks around the median of the true distribution, but cannot reproduce the high- a_{50} tail – there are too few halos with $a_{50} > 0.6$. This is borne out in the binned plot in the right panel, which shows a




Figure 5.11. – Prediction of accretion history using the `alpha_box100` model. *Left:* The distribution of halos in a_{50} from the predicted and true accretion histories. We also show the `delta_box100` (dotted light blue line) distribution as a comparison. *Right:* Predicted vs true formation times, with hexagonal binning for the halos.

cluster of late-forming halos, whose formation time is underpredicted. The correlation is positive, but it is slightly weaker than that in the `delta_box100` case. From the analyses of Fig. 5.10 and Fig. 5.11, it is clear that the addition of tidal anisotropy as an input feature does not improve the predictions of the neural network. In fact, the network requires additional complexity and makes slightly worse predictions for formation time. This is in contrast to `evals_box100`, where the additional input features clearly enhance the network predictions. Thus, among the three neural network models we train, using the three eigenvalues as an input yields the best performance, followed by the model using δ_L alone.

5.5. Extension to higher masses

In addition to our primary simulation `box100`, we also analyse a secondary run with a box size of $300 h^{-1} \text{Mpc}$. We use this simulation (`box300`) to explore the neural network predictions at higher masses. The summary of this exploration is that the network predictions remain accurate in the higher-mass regime, even though the mass resolution of the simulation is lower. Interestingly, in this regime, the excursion-set predictions are better aligned with the simulation. As before, we consider two neural networks with two sets of input features – `delta_box300` with only δ_L as input, and `evals_box300` with the three eigenvalues as inputs. We oversample the training and validation data for examples with $a_{50} < 0.45$ and $a_{50} > 0.75$. For halos in the test set of the neural network models, we apply excursion-set theory to get predictions for their accretion history. The SC and EC models remain the same as before, while we modify the Bo17 criteria such that proto-halos stop




Figure 5.12. – As in Fig. 5.7, but for models based in the `box300` simulation. *Left:* Predictions from the two neural networks. *Right:* The excursion-set theory predictions.

accreting mass if their ellipticity exceeds 0.3. Let us now examine the predictions made by these models.

First, we consider the accretion history of halos in `box300`, condensed by taking the median over the test set. As shown in Fig. 5.12, the `box300` median halo (solid red line in both panels) is more massive than its `box100` counterpart. As was the case for `box100` (Fig. 5.7), both neural networks predict the accretion history accurately (left panel). The excursion set predictions (right panel) miss the initial stages of mass accretion – similar to their behaviour for the `box100` proto-halos. At later times, SC overpredicts the proto-halo mass, but both EC and Bo17 are closer to the truth than their `box100` analogues. EC in particular, reproduces the accretion history remarkably well for $a > 0.6$.

Next, we compare the predicted and true masses for the test-set halos in Fig. 5.13. As before, we relegate the Bo17 plot to App. C for brevity. Both `delta_box300` (top left) and `evals_box300` (top right) predictions correlate strongly with the true halo masses, and have lower scatter compared to their `box100` analogues. This trend is even stronger for the excursion-set predictions. The SC predictions (bottom left) are markedly improved, with a correlation coefficient of ≈ 0.83 , compared to ≈ 0.67 in the `box100` case. On a similar note, the r -value of the EC predictions is higher than `delta_box300`, despite a cluster of underpredictions at lower masses. Note also the increase in the number of bins with more than 30 halos, indicating that the halos are not as spread out as in the `box100` case.

Finally, let us consider the formation times of the test-set halos. Analogous to Fig. 5.8, we plot the predicted a_{50} vs the truth for individual halos in Fig. 5.14. For the neural networks (top row), the r -values are similar to the `box100` case (slightly smaller for `delta_box300` and larger for `evals_box300`). The excursion-set predictions are greatly improved compared to their `box100` versions, with lower scatter and a stronger correlation.




Figure 5.13. – Predicted vs true halo masses, as in Fig. 5.4 but for models based in `box300`. *Top*: The two neural network predictions – `delta_box300` (left) and `evals_box300` (right). *Bottom*: The excursion-set theory predictions – SC (left) and EC (right).

This pattern of comparatively lower scatter for the excursion-set predictions is curious, and may be due to the range of masses probed by `box300`. To explore this further, let us revisit how the excursion-set predictions are made. A halo is assigned mass whenever an uncorrelated random walk crosses the collapse barrier at a specific mass scale. For higher-mass halos, the corresponding smoothing scales are larger, and the variance of the linear overdensity field is smaller. Thus, the trajectories of δ_L are more consistent and predictable. As a consequence of this lowered stochasticity at higher masses, the corresponding excursion-set predictions have low scatter.

In this chapter, we have tested different analytical models of halo formation. We designed neural networks to find a map between properties of the initial proto-halo and the accretion history of the halo that forms from the subsequent collapse. We compared these models against the traditional excursion-set predictions, and found that the neural networks outperform excursion-set theory in most of our tests. Further, using the three eigenvalues λ_i as inputs for the neural network provides the

5. Predicting halo formation using a neural network




Figure 5.14. – Predicted vs true formation times, as in Fig. 5.8, but for models based in `box300`. *Top*: The two neural network predictions – `delta_box_300` (left) and `delta_box_300` (right). *Bottom*: The excursion-set theory predictions – SC (left) and EC (right).

most information on the halo formation process. In Chapter 6, we summarise the results of this study, and provide a brief outlook into avenues of future research.

CHAPTER 6

Conclusions

Starting from a hot, dense, primordial seed, the Universe has evolved a great diversity and richness of observable structures. The analytical, numerical, and observational tools with which we study our Universe reveal a consistent story of its evolution. However, there remain gaps in our understanding of structure formation, and we do not yet have precise, complete knowledge of the non-linear physics that governs clustering. The overarching goal of cosmology is to bridge this gap between our theoretical models of the universe, and the observations we make.

The growth and clustering of dark matter is driven, primarily, by the gravitational amplification of small fluctuations in the early Universe. Clustering occurs in a top-down manner, with smaller structures forming first. Given our current understanding, we expect dark matter to clump into halos in which observable galaxies form. In the linear regime, where the overdensity field is small, perturbation theory (PT) provides an accurate picture of clustering. As the overdensity within a region grows, non-linear effects start to dominate. In the quasi- and non-linear regimes, the limitations of PT are laid bare. The most pressing concern, is that PT truncates the Vlasov cumulant hierarchy at second-order, while simulations suggest that all higher-order cumulants are generated during non-linear gravitational collapse (e.g. Pueblas & Scoccimarro, 2009).

At the beginning of this work, we outlined four specific questions that enquire about the evolution of structure in the regimes where PT fails. Through the course of this thesis, we have attempted to answer these questions in depth. In this concluding chapter, we summarise the results of our study. For completeness, let us reiterate the key questions here

- How do different analytical theories of structure formation compare with each other in the quasi-linear and non-linear regimes?

- Are the assumptions that these theories rely on justified?
- Gravitational instability leads to the formation of virialised structures called halos. How do their statistical and individual properties depend on the environments in which they are born?
- More generally, what is the best theoretical approach to map the initial dark matter overdensity field to collapsed structures?

A. Perturbative theories of structure formation

In the first half of this work, we investigated an extension to PT – viz. the effective field theory of large-scale structure (EFTofLSS; Baumann et al., 2012; Carrasco et al., 2012, 2014a) – which claims to alleviate the conceptual problems that PT suffers from. Specifically, we designed a one-dimensional toy universe populated by a high-amplitude, short-wavelength perturbation, and a low-amplitude long-wavelength perturbation. By analysing the impact of the former on the gravitational evolution of the latter, we were able to test the key assumptions that the EFTofLSS makes. The EFTofLSS assumes that dark matter is an effective fluid, whose mathematical description is provided by an effective stress tensor (scalar in one-dimension). This stress tensor, when treated with a physical model, contains unknown coefficients which generate corrections to the standard PT result. There are two formulations of the EFTofLSS – a top-down one where the correction to the power spectrum is calculated by fitting the EFT prediction to numerical/observational data, and a bottom-up one where the correction is measured explicitly from simulations. Both formulations share the same equations, and differ only in their methodology. Thus, we expect their predictions to agree with each other. The bottom-up construction makes critical assumptions about the form of the effective stress, and the separation of scales. In Chapter 4, we compared the two predictions as a test of these assumptions. Below, we collect our results in a brief summary

- The bottom-up EFT correction to PT relies on the measurement of the theory’s Wilson coefficients. In one-dimension, there is only one Wilson coefficient, a generalised, effective speed of sound parameter – c_{tot}^2 . We have performed a first measurement of its time-evolution. We found that c_{tot}^2 grows linearly with scale factor a until the first shell crossing, at which point it reaches a local maximum and subsequently drops in value (Fig. 4.5).
- Following leads in the existing literature, we hypothesised that c_{tot}^2 is sourced by perturbations entering the non-linear regime. We ran tests to check the viability of this hypothesis, and found evidence indicating that this is indeed the case (Figs 4.6 and 4.7).
- We calculated the evolution of power in the long-wavelength mode in both PT and the EFTofLSS. Comparing this to the measured power from N -body

simulations shows that the EFTofLSS accurately models the feedback of the short-wavelength mode on the long-wavelength mode. Quantitatively, while the PT power at $a = 6$ is discrepant by $\approx 2.2\%$, the EFT prediction is accurate to within 0.1% (Fig. 4.9). It must be noted, however, that only PT is asymptotically correct in the low- a regime.

- We tested the dependence of the EFT predictions on the change in separation of scales, and found the results to be robust to this modification (Fig. 4.10).

Additionally, there are three takeaways that we did not mention in Chapter 4, but are nonetheless interesting findings of our work. Details of these may be found in Karandikar et al. (2024) (hereafter, K24)

- In this thesis, we have ignored the effects of the stochastic term J that appears in the expansion of the effective stress (e.g., Eq. 4.3). We explicitly measured the effect of J on the power spectrum, and found it to be subdominant (Sec. 5.2 in K24).
- We also tested the idea of renormalisation within our setup, but found it difficult to apply due to the absence of a continuous power spectrum (Sec. 5.4 in K24).
- We have performed all of our analysis using both a sharp- k filter, and a Gaussian filter. We have only tabulated the results of the former in this thesis. The latter also provides qualitatively identical results with some subtleties and minor numerical differences (Appendix A in K24).

A possible extension of our work is to generalise these results in a three-dimensional setting. This would serve as an even stronger verification of the soundness of the underlying assumptions. However, the generalisation would come at a greater computational cost, and an equivalent three-dimensional toy model will be vastly more complicated. A more reasonable middle ground may be to remain in one-dimension, but allow for a continuous initial power spectrum drawn from a Gaussian probability distribution. This would test the EFT assumptions in a more realistic setting, without the added complexity of three-dimensions. An important reason to generalise our setup is to extract a physically meaningful effective speed of sound. As shown in Sec. 4.4, the speed of sound contains information about Fourier modes which are entering the non-linear regime. A promising step in this direction was recently taken by Nascimento et al. (2025), who obtain a semi-analytic estimate the speed of sound, and its cosmology dependence.

The EFTofLSS started as an Eulerian-space extension to PT, designed to predict the dark matter power spectrum. In the intervening decade or so, the theory has been extended to the Lagrangian frame (Porto et al., 2014; Vlah et al., 2015), to include biased tracers (Assassi et al., 2014; Senatore, 2015; Angulo et al., 2015a;

Mirbabayi et al., 2015; Donath & Senatore, 2020), and to treat redshift-space distortions (Lewandowski et al., 2018; D’Amico et al., 2024). We have also seen efforts towards speeding-up the EFT loop calculations (Simonović et al., 2018; Anastasiou et al., 2024), and using machine learning techniques in combination with the EFTofLSS (Donald-McCann et al., 2023; Tucci & Schmidt, 2024). The theory will also be applied to the treasure-trove of data expected from the next generation of cosmological surveys – DESI (Adame et al., 2024) and *Euclid* (Scaramella et al., 2022), among others (see e.g. Maus et al., 2024).

From an analytical standpoint, there may be other viable alternatives to PT. For instance, Kinetic Field Theory (KFT; Bartelmann et al., 2019) models the dark-matter particle ensemble statistically using a generating functional. This object encodes phase-space dynamics, and time-evolution is furnished by taking functional integrals over it. Although KFT has been favourably compared to SPT (Kozlikin et al., 2021), its viability in a conceptual and computational sense remains unclear. More recently, Garny et al. (2023a,b) proposed an extension to PT that self-consistently treats the whole Vlasov cumulant hierarchy. This approach, termed Vlasov Perturbation Theory (VPT) avoids the small-scale issues that standard PT suffers from, by including the higher-order Vlasov cumulants in the perturbative expansion. Although these alternatives are yet to be fully developed, a diversity of analytical techniques is an encouraging sign.

When comparing our theoretical models to observational data, the canonical approach has been to use the N-point correlators of the matter overdensity field, chiefly the power spectrum and bispectrum. Pushing this to higher-order correlators has, thus far, been prohibitively expensive. While work towards faster computation of these statistics goes on (e.g. Philcox & Slepian, 2022), alternative summaries of the non-linear field are also being explored. These include topological structures in the overdensity field (Pranav et al., 2019; Wilding et al., 2021), one-point statistics (Uhlemann et al., 2020), and the wavelet scattering transform (Valogiannis & Dvorkin, 2022). At the same time, field-level inference techniques – exploiting the full information available in the overdensity field – are being used to draw conclusions about cosmology (Hahn et al., 2024; Lemos et al., 2024; Nguyen et al., 2024).

B. Models of clustering of dark matter halos

In the second half of the thesis, we explored the use of neural networks to predict the formation of dark matter halos, and compared these predictions against those from excursion-set theory. Our objective was to train the neural network on properties of the initial proto-halo population, and predict the formation history of halos. Below, we summarise our key findings, which serve as answers to the last two of our key questions

- A shallow neural network with one (or occasionally two) hidden layer(s) is sufficient to accurately predict, on average, the accretion history of halos (Fig. 5.7).

	SC	EC	Bo17	delta_box100	evals_box100	alpha_box100
R^2 (MAH)	0.28	0.79	0.53	0.82	0.93	0.78
$r(M_h)$	0.67	0.81	0.67	0.83	0.94	0.86
$r(a_{50})$	0.26	0.38	0.38	0.56	0.59	0.54

Table 6.1. – Comparison of model performance for **box100**. The rows contain our metrics, and the columns represent the six models. The highlighted values are the highest values in each row.

	SC	EC	Bo17	delta_box300	evals_box300
R^2 (MAH)	0.33	0.84	0.52	0.86	0.92
$r(M_h)$	0.83	0.88	0.84	0.86	0.94
$r(a_{50})$	0.41	0.47	0.44	0.54	0.62

Table 6.2. – As in Table 6.1, but for models based in **box300**. The excursion-set models perform better than their **box100** counterparts.

The prediction for the halo mass is accurate with low scatter, and the mass function of the halo population is also recovered (Figs 5.3 and 5.4).

- To match the distribution of halos in formation time, oversampling the training dataset was necessary. Predictions for the formation time are less accurate than those for the final halo mass. This is consistently true across models and throughout the explored mass range (Figs 5.8 and 5.14). The reason for this is likely that the input features of the network (and the input of the excursion-set models) are insufficient to predict this derived quantity. However, the neural network predictions are still reasonably accurate, with the predicted-to-true correlation hovering between 0.55 – 0.6 for different input features.
- Using only the linear overdensity (δ_L) as an input feature, the network outperforms spherical collapse (which also uses the same information) and ellipsoidal collapse (albeit these predictions are similar at higher masses). The additional information provided by using the eigenvalues of the linear deformation tensor further improves the predictions. However, a combination of δ_L and the tidal anisotropy does not improve the predictions over using δ_L alone (Figs 5.10 and 5.11).
- At higher masses, the excursion-set theory predictions improve in comparison to the lower-mass regime. This is likely due to the trajectories of δ_L being more consistent (Figs 5.12, 5.13, and 5.14).

We tabulate the key numbers for models of **box100** and **box300** in Table 6.1 and Table 6.2 respectively. Our main conclusion is that neural networks outperform excursion set theory when predicting the accretion history of halos, and also recover derived properties of the halo distribution reasonably well. Out of the three input

feature sets, using the eigenvalues of the linear deformation tensor yields the best performance.

There are several ways in which our work may be extended/built upon. From a computational point of view, ours is one of a series of recent works that exploit the power of neural networks to offer insight into the halo formation process. Our work most resembles that of Lucie-Smith et al. (2019), who use a neural network to predict the final mass and concentration of halos. Instead, we have predicted the full accretion history, shedding light on the formation of halos. The obvious question one may ask, is how can we extract some physical understanding from our neural network. One avenue is to try and understand which spatial scales are most important in driving halo formation. This may be achieved by ranking the effect of the input features on the output distribution, thus finding which input feature (i.e. smoothing scale) affects halo formation the most. This is called importance sampling (see Tokdar & Kass, 2010 for a review), and was used in Lucie-Smith et al. (2019) to determine which tidal property – out of linear overdensity, ellipticity, and prolateness – was most informative on the final mass.

Another possibility is to study the role of halo properties such as their density profiles, concentration, and spin. For instance, Lucie-Smith et al. (2022, 2024) has used neural networks to understand the density profile of halos, and probed the role of this assembly bias using a convolutional neural network (Lucie-Smith et al., 2023). The interplay of halo accretion and the properties of the galaxies that reside within them is also active field of research (Contreras et al., 2021; Montero-Dorta et al., 2021; Delgado et al., 2022).

On the analytical side, a novel idea was proposed by Musso & Sheth (2021), who used peaks in the *energy* overdensity field as seeds of halo formation. In this framework, the analogue of the tidal deformation tensor is called the energy shear tensor. Eigenvalues of this tensor may be used to generate analytical predictions for the collapse threshold (Musso & Sheth, 2023). Further, the halo boundary may be found by considering equipotential surfaces around the proto-halo center of mass Musso & Sheth (2024). It would be relatively straightforward to adapt our work to this new framework, and make systematic comparisons with what we have discussed here.

Given the progress of our theoretical models of structure formation, the wealth of numerical techniques at our disposal, and the anticipation of vast swathes of informative data, our field finds itself in a position of extraordinary richness. Concurrently, our cosmological framework faces increased scrutiny in the face of this richness, with inconsistencies between theory and observations that point towards updating or usurping it. In this thesis, we have shed light on a few of these open questions, filling some gaps between the way we model our Universe, and the way we observe it to be.

APPENDIX A

Statistical description of particles

In this appendix, we summarise the key ideas behind the statistical treatment of a system of identical, classical particles that is applied to DM particles as discussed in Chapter 2. Ordinary vectorial mechanics treats particles as individual entities, with forces acting on each particle specified by Newton’s second law of motion. Analytical mechanics, on the other hand, considers the system of particles as a statistical ensemble, with a single scalar function (viz. the Lagrangian, or equivalently the *action*) containing all information about the system. The equations of motion can be reconstructed from the action by means of a minimising principle.

A.1. The Euler-Lagrange equations

Let us start with a generic ensemble of N particles in three-dimensional space, completely specified by the Lagrangian which (generally) depends on the positions, velocities and time – $L(\mathbf{x}, \dot{\mathbf{x}}, t)$. The action of this system is

$$S = \int dt \, L(\mathbf{x}, \dot{\mathbf{x}}, t) . \quad (\text{A.1})$$

The dynamics of this system can be discussed in three different frames

- the three-dimensional configuration space of all possible positions $\mathbf{x}$ (we will loosely include the time-coordinate here),
- the six-dimensional phase-space of all possible positions and momenta $(\mathbf{x}, \mathbf{p})$ where the state of the system is described by N points – one corresponding to each particle, or

- the $6N$ -dimensional Γ -space specified by the possible positions and momenta of all particles taken together. A point in Γ -space represents the state of the system.

Motion in phase- (Γ -) space is specified by trajectories of N (one) points. In configuration space, given two fixed points – $(\mathbf{x}_1, t_1)$ and $(\mathbf{x}_2, t_2)$, the motion of the system leaves the action invariant in accordance with Hamilton’s principle. More precisely, the system follows the trajectory that *minimises* the action integral. Hamilton’s principle is sufficient to generate the dynamical evolution of the system. The corresponding Euler-Lagrange equations (see e.g. Lanczos, 2012 for the derivation)

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{x}}} \right) - \frac{\partial L}{\partial \mathbf{x}} = 0 . \quad (\text{A.2})$$

A.2. Hamilton’s equations

Given the positions and velocities $(\mathbf{x}, \dot{\mathbf{x}})$ of the system, consider a coordinate transformation

$$(\mathbf{x}, \dot{\mathbf{x}}) \rightarrow (\mathbf{x}', \dot{\mathbf{x}}') \quad (\text{A.3})$$

Let’s express Hamilton’s principle for both sets of coordinates as

$$\delta \int_{t_1}^{t_2} dt L(\mathbf{x}, \dot{\mathbf{x}}, t) = 0 , \quad (\text{A.4})$$

$$\delta \int_{t_1}^{t_2} dt L'(\mathbf{x}', \dot{\mathbf{x}}', t) = 0 , \quad (\text{A.5})$$

where δ is a virtual displacement called the *variation*. For a given coordinate transformation, if Eq. (A.4) implies Eq. (A.5), then it is called a *canonical* transformation. Often, instead of dealing with positions and velocities, it is advantageous to work in the phase-space. To make this leap, we perform the Legendre transformation

$$(\mathbf{x}, \dot{\mathbf{x}}) \rightarrow (\mathbf{x}, \mathbf{p}) , \quad \text{where} \quad \mathbf{p} = \frac{\partial L}{\partial \dot{\mathbf{x}}} , \quad (\text{A.6})$$

by defining the Hamiltonian

$$H = \sum \frac{\partial L}{\partial \dot{\mathbf{x}}} \dot{\mathbf{x}} - L , \quad (\text{A.7})$$

where the sum is over the coordinates. The Euler-Lagrange equation can then be recast in this new language, resulting in Hamilton’s equations of motion

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}} , \quad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}} . \quad (\text{A.8})$$

One subtlety we gloss over from here on is that the Lagrangian may be explicitly time-dependent. This is only true for dissipative (i.e., non-conservative) systems and therefore does not apply to our treatment the dark matter ensemble. If we are dealing with a dissipative system, the partial time-derivatives of the Lagrangian and Hamiltonian are equal apart from a negative sign. Two variables linked via an equation of the form of Eq. (A.8) are called canonically conjugate variables.

A.3. Poisson brackets

It is useful at this point to introduce the notation of Poisson brackets. Given two quantities A and B , the Poisson bracket is defined as

$$\{A, B\} = \sum \left(\frac{\partial A}{\partial \mathbf{p}} \cdot \frac{\partial B}{\partial \mathbf{x}} - \frac{\partial A}{\partial \mathbf{x}} \cdot \frac{\partial B}{\partial \mathbf{p}} \right), \quad (\text{A.9})$$

for the conjugate position-momentum pair (q, p) . These conjugate variables follow the following poisson bracket relations

$$\{q_i, q_j\} = 0, \quad \{p_i, p_j\} = 0, \quad \{q_i, p_j\} = \delta_{ij}. \quad (\text{A.10})$$

In this language, Hamilton's equations may be re-expressed as

$$\dot{\mathbf{x}} = \{\mathbf{x}, H\}, \quad \dot{\mathbf{p}} = \{\mathbf{p}, H\}. \quad (\text{A.11})$$

A.4. Liouville's theorem

Liouville's theorem is primarily a mathematical result that has substantial physical consequences. Given the canonical transformation

$$(\mathbf{x}, \mathbf{p}) \rightarrow (\mathbf{x}', \mathbf{p}'), \quad (\text{A.12})$$

it states that the Jacobian of this transformation is unity (e.g. Liboff, 2003), i.e.

$$J \left(\frac{\mathbf{x}', \mathbf{p}'}{\mathbf{x}, \mathbf{p}} \right) = 1. \quad (\text{A.13})$$

At any point in phase-space, the Jacobian quantifies the amount by which the transformation stretches the neighbourhood around the point. Consider a phase-space volume $d\Omega = d\mathbf{p} d\mathbf{x}$, and define the function $f_N(\mathbf{p}, \mathbf{x}, t)$ such that its integral over $d\Omega$ counts the number of particles in that volume at any time t . We may write

$$f_N = \frac{dN}{d\Omega} \quad (\text{A.14})$$

Now let the system evolve from $t \rightarrow t + \Delta t$. Given that the coordinates and momenta uniquely specify the state of the system, system trajectories may never cross each other in phase-space. As a consequence of this, the number of particles within the volume does not change during time-evolution. In fact, this time-evolution is a canonical transformation, and by Liouville's theorem,

$$d\Omega(t) = d\Omega(t + \Delta t) . \quad (\text{A.15})$$

Following Eq. (A.14),

$$\frac{df_N}{dt} = 0 . \quad (\text{A.16})$$

This is called Liouville's equation. Reformulated in this way, Liouville's theorem expresses the invariance of phase-space volume along trajectories which the system follows. We can also write this in terms of the Hamiltonian as follows

$$\frac{\partial f_N}{\partial t} - \mathcal{L}f_N = 0 , \quad (\text{A.17})$$

where $\mathcal{L} = \{\cdot, H\}$ is the Liouville operator, and can be related to the Hamiltonian operator defined in Sec. 3.1.2 as $i\mathcal{L} = \mathcal{H}$.

A.5. BBGKY hierarchy and the Vlasov equation

The function f_N can be interpreted as the joint-probability density of N system points in phase-space. Let us introduce the notation $d\ell = d\mathbf{x}_\ell d\mathbf{p}_\ell$. Thus, $f_N d1 \dots dN$ represents the probability that the ℓ^{th} particle occupies position $\mathbf{x}_\ell$ and has momentum $\mathbf{p}_\ell$. Let us then define the m -particle distribution ($m < N$) by integrating over the remaining $N - m$ particles

$$f_m(1, \dots, m) = \int d(m+1) \dots dN f_N(1, \dots, N) . \quad (\text{A.18})$$

For identical particles, f_m is symmetric in its arguments. The evolution of the m -particle distribution is linked to the $(m+1)$ -particle distribution as (see Liboff 2003 for the derivation)

$$\left(\frac{\partial}{\partial t} - \mathcal{L} \right) f_m + (N - m) \sum_{i=1}^m \frac{\partial}{\partial \mathbf{p}_i} \cdot \int d(m+1) \mathbf{G}_{i,m+1} f_{m+1} , \quad (\text{A.19})$$

where $\mathbf{G}_{i,j} = -\partial_j \phi_{ij}$ is the gravitational acceleration due to the interaction between the i^{th} and j^{th} particle. The set of equations is called the BBGKY hierarchy (after Bogoliubov, Born, Green, Kirkwood, and Yvon) and gives us the full time-evolution of the N -particle system. This is an infinite hierarchy, and the only way to make meaningful progress is to truncate it and obtain a closed set of equations. Consider

the first equation in the hierarchy that relates the one- and two-particle distribution functions

$$\left(\frac{\partial}{\partial t} - \mathcal{L} \right) f_1 + (N-1) \frac{\partial}{\partial \mathbf{p}_1} \cdot \int d2 \mathbf{G}_{12} f_2(1,2) . \quad (\text{A.20})$$

Since dark matter is assumed to be collisionless particles of mass m , we can assume the system to be weakly coupled, i.e. the potential energy to be much smaller than the kinetic energy. This implies that $f_2(1,2) = f_1(1)f_2(2)$. Further, considering only long-range interactions, we may expand Eq. (A.20) as

$$\left(\frac{\partial}{\partial t} + \frac{\partial \mathbf{x}}{\partial t} \cdot \frac{\partial}{\partial \mathbf{x}} + \frac{\mathbf{G}}{m} \cdot \frac{\partial}{\partial \mathbf{p}} \right) f(\mathbf{x}, \mathbf{p}, t) = 0 , \quad (\text{A.21})$$

where we drop the subscript ‘1’ for convenience. This is known as the Vlasov equation, and governs the time-evolution of the one-particle distribution function, physically representing a collisionless system of N particles. The equation is a cornerstone of plasma physics and, more importantly for us, also describes the system of dark matter particles.

APPENDIX B

Predicted accretion histories by population

In this appendix, we attempt to get a more granular understanding of the accretion histories predicted by our neural networks, and using excursion-set theory. We split the halo population into three subsets based on a_{50} , reflecting our undersampling during the training process. The largest of these subsets contains a majority of the population (central panel of Fig. B.1), for whom $0.4 \leq a_{50} \leq 0.7$. The other two subsets represent the minority tails (left and right panels). We indicate the fraction of total test-set-halos in each of the subsets by the number f in Fig. B.1. For halos with $a_{50} \leq 0.4$, all models systematically overestimate the growth of halos at late times. This hints at a physical process that suppresses the growth of these halos, which is missed by only considering the eigenvalues of the tidal tensor. For the majority subset, the predictions are largely as we expect from Fig. 5.7. The only outlier is EC, which is much closer to the simulation history than in Fig. 5.7. This is explained by the failure of EC to predict the accretion of the high- a_{50} halos, as seen in the right panel. For this subset, the excursion-set models predict that the proto-halo grows much later, which suggests that the Lagrangian patch is less overdense (compared to halos in the majority subset). In the simulation, the patch has a relatively ordinary growth until $a = 0.75$, after which it undergoes a period of strong accretion. This is likely due to merging events at late times. We may speculate that this class of proto-halo is surrounded by other overdense patches, making mergers more likely. This information is not captured by the tidal field around the halo. The neural network predictions remain reasonable until the strong growth at $a > 0.75$. By $a = 1$, the `evals_box100` mass is closest to the simulation, underpredicting it by a factor of ≈ 1.4 .




Figure B.1. – The median accretion history predicted by the two machine learning models (dotted lines), compared against the true median accretion history of the test set (solid red line). We also show the excursion-set predictions from both spherical collapse (dashed teal line) and two ellipsoidal collapse models (dot-dashed lines). *Left:* Low- a_{50} tail. *Centre:* Majority subset containing halos with $0.4 \leq a_{50} \leq 0.7$. *Right:* High- a_{50} tail. We also indicate the fraction of halos in each subset (f).

APPENDIX C

Comparison plots from Bo17

In this appendix, we show predictions from the excursion-set Bo17 model that we omitted in the main text. In Fig. C.1, we compare the predicted and true halo masses (left), and formation times (right). Bo17 predicts the halo mass better than SC, although there are several outliers. The mass of many halos in the $5 \times 10^{11} - 5 \times 10^{12} h^{-1} M_{\odot}$ range is overpredicted, and there is also cluster of underpredictions for masses $< 10^{12} h^{-1} M_{\odot}$. In comparison to EC, which also uses ellipsoidal information, Bo17 performs slightly worse. Similar conclusions may be drawn about the formation time predictions, with an r -value of ≈ 0.38 better than SC, but worse than EC.




Figure C.1. – Predicted vs true halo mass (left) and formation time (right) for Bo17 using box100. We group the halos in hexagonal bins to better illustrate the underlying correlation.

When extended to higher masses, Bo17 shows a similar behaviour as its excursion-set theory cousins, with significant improvements relative to the box100 results. As mentioned in the main text, this is likely due to the reduced stochasticity of the barrier crossing problem that leads to the excursion-set theory results. The

correlation between predicted and true halo masses (left panel of Fig. C.2) is close to that achieved by `delta_box100` (see Table 6.2). However, the formation time predictions (right panel) remain worse than either neural network.




Figure C.2. – Same as Fig. C.1, but for models based in `box300`.

APPENDIX D

Testing the assumptions of the EFTofLSS

In this appendix, we reproduce the pre-print¹ of Karandikar et al. (2024). The published version of the paper may be found as JCAP01(2024)051 on the IOP Publishing website².

¹It may be accessed at <https://arxiv.org/abs/2307.11090>

²<https://doi.org/10.1088/1475-7516/2024/01/051>

Testing the assumptions of the Effective Field Theory of Large-Scale Structure

Mandar Karandikar^a Cristiano Porciani^a Oliver Hahn^{b,c}

^aArgelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany

^bDepartment of Astrophysics, University of Vienna, Türkenschanzstraße 17, 1180 Vienna, Austria

^cDepartment of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

E-mail: mandar@astro.uni-bonn.de, porciani@astro.uni-bonn.de,
oliver.hahn@univie.ac.at

Abstract. The Effective Field Theory of Large-Scale Structure (EFTofLSS) attempts to amend some of the shortcomings of the traditional perturbative methods used in cosmology. It models the evolution of long-wavelength perturbations above a cutoff scale without the need for a detailed description of the short-wavelength ones. Short-scale physics is encoded in the coefficients of a series of operators composed of the long-wavelength fields, and ordered in a systematic expansion. As applied in the literature, the EFTofLSS corrects a summary statistic (such as the power spectrum) calculated from standard perturbation theory by matching it to N -body simulations or observations. This ‘bottom-up’ construction is remarkably successful in extending the range of validity of perturbation theory. In this work, we compare this framework to a ‘top-down’ approach, which estimates the EFT coefficients from the stress tensor of an N -body simulation, and propagates the corrections to the summary statistic. We consider simple initial conditions, viz. two sinusoidal, plane-parallel density perturbations with substantially different frequencies and amplitudes. We find that the leading EFT correction to the power spectrum in the top-down model is in excellent agreement with that inferred from the bottom-up approach which, by construction, provides an exact match to the numerical data. This result is robust to changes in the wavelength separation between the two linear perturbations. However, in our setup, the leading EFT coefficient does not always grow linearly with the cosmic expansion factor as assumed in the literature based on perturbative considerations. Instead, it decreases after orbit crossing takes place.

Contents

1	Introduction	1
2	Setup	5
3	N-body simulations	6
4	Top-down EFT coefficients	7
4.1	Transport equations for the effective fluid	8
4.2	Coarse-graining and non-linear scales	11
4.3	Effective stress	11
4.4	Effective speed of sound	12
4.4.1	Estimates from spatial cross-correlations	14
4.4.2	Estimates from least-squares fitting	16
4.4.3	Comparing the different estimates	16
5	Performance of the EFTofLSS	18
5.1	Next-to-leading-order power spectrum in SPT and EFTofLSS	18
5.2	Stochastic term	22
5.3	Dependence on the separation of scales	23
5.4	Considerations about renormalisation	23
5.5	Long-term evolution of the effective speed of sound	25
6	Summary	26
	References	29
A	Results with Gaussian smoothing	32

1 Introduction

On scales smaller than the Hubble radius, dark matter dynamics is governed by the Vlasov-Poisson (VP) system of equations. In order to study the formation and evolution of the large-scale structure of the Universe, two approaches are widely used to solve the VP system: N -body simulations and perturbative techniques. The former approximate the phase-space distribution by a set of macroparticles that are gravitationally evolved [e.g. 1, 2]. In the latter, we take velocity moments of the Vlasov equation, and solve the resulting equations perturbatively in the Eulerian or Lagrangian frame [e.g. 3, 4]. To make this viable, the hierarchy of Vlasov moments is truncated at second order, by assuming the vanishing of the second cumulant – the stress tensor. This truncation is justified in the cold limit before the crossing of distinct streams of matter (orbit crossing). The physical picture is that of a pressureless, perfect fluid.

Perturbation theory has been successful in advancing our understanding of structure formation on large scales [e.g. 5–10]. However, it has two important shortcomings. First, non-linear gravitational clustering breaks the assumption of negligible velocity dispersion on successively larger scales, and second, it treats non-linear scales as perturbative. These are

non-trivial problems, as at the instant of orbit crossing, all higher-order Vlasov cumulants are simultaneously generated, giving rise to a closure problem for the hierarchy. If we are interested in the overdensity and velocity fields, we can close the system of equations by modelling the stress tensor (and its time evolution).

Theoretical advances in recent years have focused on constructing an effective field theory of large-scale structure (EFTofLSS) [e.g. 11–13] that attempts to extend the reach of perturbation theory to larger wavenumbers. An effective theory is an approximate description of a physical system. It assumes the existence of a hierarchy of scales characterising the system it seeks to describe. Given such a hierarchy, the physics on short-wavelength (UV) scales may be described separately from that on long-wavelength (IR) scales. A UV-complete theory captures the physics on all scales, while the EFT describes only the IR scales, making no claims about the UV physics. The coupling between scales is accounted for by using coefficients that are measured from the UV-complete theory, or treated as free parameters. The starting assumption of the EFTofLSS is that there exists a length scale at which perturbation theory breaks down. This ‘non-linear scale’ is usually described in terms of the associated wavenumber k_{NL} and approximately corresponds to the typical distance crossed by dark-matter particles until a given epoch. The goal of the EFTofLSS is to build a perturbative expansion for summary statistics of observable quantities (e.g. their correlation functions or power spectra) organised in powers of k/k_{NL} so that it converges in the perturbative regime $k \ll k_{\text{NL}}$. The IR scales are described by modifying Eulerian¹ standard perturbation theory (SPT, see e.g. [3] for a review). Assuming that the UV and IR scales are well-separated, the EFTofLSS proceeds in the following way:

1. COARSE-GRAINING. The Vlasov equation is smoothed over a length scale much larger than the non-linear scale. Following the convention in the literature, we characterise this smoothing scale in terms of its corresponding wavenumber $\Lambda \ll k_{\text{NL}}$. As routinely done in fluid mechanics, a macroscopic model is obtained by taking the system of moment equations and truncating it at a suitable order. A cornerstone of the EFTofLSS is that the Vlasov hierarchy for the coarse-grained system (which only includes perturbations with wavenumber $k \ll \Lambda \ll k_{\text{NL}}$) can be safely truncated at second order and is thus described by two ‘fluid’ equations. The key idea here is that, due to the finite age of the Universe and their relatively low speeds, dark-matter particles did not have the time to cross long distances. From scaling arguments, it follows that the length scale associated with k_{NL} plays a similar role as the mean free path in a neutral gas dominated by collisions, i.e. higher-order moments of the phase-space density are suppressed by powers of k/k_{NL} [11]. However, contrary to the case of the rarefied gas (where the suppression is due to the frequent molecular collisions), the higher-order moments in the EFTofLSS simply did not have time to develop within the age of the Universe. The time evolution of the coarse-grained system is thus regulated by the continuity equation and the Cauchy momentum equation containing an effective stress tensor τ . A number of arguments indicate that virialised structures decouple completely from the long-wavelength theory and do not alter the expansion history of the background Universe [11]. The largest contribution to the effective stress tensor is thus expected from scales that have recently turned non-linear but did not yet develop into virialised structures (i.e. Fourier modes with $k \simeq k_{\text{NL}}$).

¹An effective theory in Lagrangian space can also be constructed [14, 15].

2. **INTEGRATING OUT SMALL-SCALE FLUCTUATIONS.** In a specific realisation, the effective stress tensor depends on both the long-wavelength perturbations with $k \leq \Lambda$ (which alter the particle trajectories through tidal effects) and the short-wavelength perturbations with $k > \Lambda$. In order to eliminate the latter dependence, the EFTofLSS computes the evolution of the Fourier modes with $k \ll \Lambda$ by replacing $\boldsymbol{\tau}$ with its conditional average $\langle \boldsymbol{\tau} \rangle$ which characterises the ‘typical’ influence of the small-scale modes on a patch of size Λ^{-1} with fixed long-wavelength modes. This way, an effective theory that only includes long-wavelength modes is obtained.
3. **STOCHASTIC TERMS.** In each realisation, the difference $\mathbf{J} = \boldsymbol{\tau} - \langle \boldsymbol{\tau} \rangle$ provides a source of noise in the evolution of the large-scale perturbations with respect to the averaged case. This stochastic contribution is not necessarily small with respect to $\langle \boldsymbol{\tau} \rangle$ but, at leading order, generates corrections that should be uncorrelated with the long-wavelength modes and therefore suppressed in averaged statistics like the power spectrum of density fluctuations for $k \ll \Lambda$. On the other hand, the auto-correlation of the stochastic terms does not vanish and it is often assumed to be Poisson-like on scales comparable with k_{NL} [12].
4. **PERTURBATIVE EFT.** Since the smoothed fluid properties (viz. fluctuations in the matter density and velocity potential) lie, by construction, in the quasi-linear regime, it makes sense to study their evolution perturbatively. In order to solve the non-linear equations regulating the effective fluid, $\langle \boldsymbol{\tau} \rangle$ is expanded in terms of the long-wavelength fields and their spatial derivatives. All terms allowed by the symmetries of the problem (i.e. of the UV-complete theory) are considered. This set of interaction terms is ordered by importance using a power-counting technique which yields a systematic expansion in k/Λ or k/k_{NL} . A finite number of relevant interaction terms appear at a given order. At the lowest one, the coefficients of the expansion can be interpreted as an effective background pressure, a speed of sound, and bulk- and shear-viscosity coefficients, in analogy with the phenomenological stress tensor for an imperfect Newtonian fluid (which leads to the Navier-Stokes equation). These time-dependent but spatially homogeneous coefficients are determined by the short-range physics and the coarse-graining scale Λ . They quantify the UV-IR coupling of the perturbations and are not predictable² within the effective theory. The EFTofLSS differs from SPT because of the presence of these additional terms in the fluid equations.
5. **PERTURBATIVE SOLUTION.** Power counting suggests that the EFT corrections to the SPT solutions only appear at non-linear order. The fluid equations of the perturbative EFT are thus solved iteratively by expanding the fields around the linear solution (as in SPT). Diagrammatic techniques analogous to Feynman diagrams can be employed to evaluate correlation functions of the long-wavelength perturbations. Tree diagrams provide the leading-order solution while higher-order corrections are associated with loop diagrams and require integrations over the wavenumber. The coarse-graining scale Λ acts as a cutoff of the loop integrals (smooth or sharp depending on the filter used to isolate the long-wavelength part of the fields). It turns out that some of the loop corrections in SPT show substantial variations with Λ and are thus said to be UV-sensitive. This strong dependence of the IR theory on the UV degrees of freedom likely reflects

²A strong simplifying assumption which is regularly made in the literature is that their time dependence can be inferred in perturbation theory [e.g. 12, 16].

the failure of the SPT macroscopic equations on small scales (i.e. the loops integrate the SPT solutions for the fields in a regime in which they are not valid). The EFTofLSS addresses this issue by choosing a renormalisation scheme for the IR-UV couplings.

6. RENORMALISATION. The UV sensitivity of the SPT loops can be absorbed into the parameters of the effective theory order by order as follows. The EFT expressions for the correlation functions of the long-wavelength perturbations contain a number of coefficients³ $c(\Lambda, a)$ deriving from the expansion of $\langle \tau \rangle$ in terms of the smoothed fields. These ‘Wilsonian’ coefficients are re-written as $c(\Lambda, a) = c_{\text{ren}}(a) + c_{\text{ctr}}(\Lambda, a)$ where the counterterm parts $c_{\text{ctr}}(\Lambda, a)$ are meant to regularise the behaviour of the loop results in SPT as a function of Λ (i.e. cancel out their UV sensitivity), while the renormalised coefficients $c_{\text{ren}}(a)$ are independent of Λ . In order for the cancellation to work, the loop and counterterm corrections must have the same dependence on k (in the limit $k \rightarrow 0$), Λ and time (but, in practice, the last two equalities are assumed to hold [e.g. 12, 16, 19]). After renormalising the coupling coefficients so that the EFT predictions for $k \ll k_{\text{NL}}$ do not depend on Λ , the limit $\Lambda \rightarrow \infty$ is taken. Although one could invoke that UV physics on virialised scales is likely to be decoupled from the IR scales, the rationale behind this last operation is just that the final expressions become simpler as the contributions from higher-derivative terms get suppressed as k/Λ . This is a mathematical trick; the EFT, in any case, is valid only for $k \ll k_{\text{NL}}$. After renormalising the coupling coefficients and taking the limit $\Lambda \rightarrow \infty$, all physical observables (evaluated at some fixed perturbative order) are unambiguously defined within some finite accuracy (i.e. up to corrections suppressed by some power of k/k_{NL}), once expressed in terms of the $c_{\text{ren}}(a)$.
7. UV MATCHING. The unknown coefficients of the effective theory (at a given perturbative order) can be treated as free parameters that are determined by matching a statistic (e.g. the power spectrum) to measurements from observations or numerical simulations (which provide an approximation to the UV-complete theory) at some scale $k_{\text{match}} \ll k_{\text{NL}}$ (or by fitting the model to the data with $k < k_{\text{match}}$).

Following this procedure, the EFTofLSS allows a perturbative treatment of the Fourier modes with $k \ll k_{\text{NL}}$, while taking into account the effect of higher- k modes. This framework has extended the predictive power of the theory to larger wavenumbers and/or later times (compared to SPT) for the Λ CDM matter power spectrum and bispectrum [e.g. 12, 14, 18, 20–23]. For example, the EFTofLSS extends the k -reach of the models for the matter power spectrum by a factor of 2–3 [24]. The formalism has also incorporated biased tracers [e.g. 25–29] (see also [30–32]).

Broadly, EFTs come in two flavours. The literature on the EFTofLSS follows a ‘bottom-up’ construction consistent with the symmetries of the system without reference to the underlying UV-complete theory. This provides justification for introducing additional free parameters in the theoretical model for summary statistics of the large-scale structure. In contrast, a ‘top-down’ construction starts from a UV-complete theory and integrates out degrees of freedom above the cutoff scale Λ . The EFT coefficients can be calculated directly from the

³Further complexity is introduced by the fact that while UV and IR scales are well separated in space, the corresponding perturbations evolve on similarly long timescales, comparable with the Hubble time (before virialisation). This implies that the effective theory is non-local in time [e.g. 12, 17, 18]. As a consequence, the coefficients $c(\Lambda, a)$ that are derived from the same effective-fluid parameter (e.g. the effective speed of sound) do not coincide at different perturbative orders as they correspond to time integrals containing different kernels.

UV-complete theory without free parameters. The two constructions result in the same dynamical equations, and differ only in the method of estimation of the coefficients. Thus, in a consistent theory, their predictions should agree with each other. For the EFTofLSS, N -body simulations provide an approximate UV-complete theory. For example, reference [12] compares the bottom-up and top-down approaches in a three-dimensional universe with a Λ CDM cosmology, and finds them to be in agreement within large error bars. However, this was shown for a fine-tuned range of separations and only at a single fixed time. Reference [33] specialises to the one-dimensional case of interacting sheets with CDM-like initial power spectra, and find similar agreement at a fixed time, while assuming a time-dependence for the coefficients.

Our main motivation in this work is to design the simplest universe in which the bottom-up and top-down constructions can be compared, and run controlled tests to assess the agreement between the two. Similar to [33], we study the non-linear growth of density perturbations with planar symmetry. However, we specialise to an Einstein-de Sitter (EdS) universe and consider much simpler, deterministic initial conditions, viz., two superimposed cosinusoidal perturbations with substantially different wavelengths. Our aim is to understand the impact of the short-wavelength perturbation on the long-wavelength one. The simplicity of the setting allows us to measure the time evolution of the EFT coefficients, in contrast to previous studies which perform their calculations at a fixed epoch.

The paper is structured as follows. In Sections 2 and 3, we make our setup precise and introduce our N -body simulations, respectively. In Section 4, we measure the IR-UV interaction terms from the simulations. These coefficients are then used in Section 5 to compare the next-to-leading order perturbative expressions for the matter power spectrum in SPT and EFTofLSS to the simulations. Finally, in Section 6, we summarise our results.

2 Setup

We build a toy model in which two linear cosinusoidal density perturbations with planar symmetry grow in a background EdS universe. In such a system, the overdensity develops singularities at shell crossing, while the gravitational acceleration remains bounded at all times [34–36]. Our goal is to study how the presence of an initially high-amplitude, short-wavelength fluctuation influences the evolution of a low-amplitude, long-wavelength one. We consider a box of comoving side L with periodic boundary conditions. Thus, the linear overdensity is of the following form

$$\delta_{\text{lin}}(x, a = 1) = A_1 \cos(k_1 x - \varphi_1) + A_2 \cos(k_2 x - \varphi_2) , \quad (2.1)$$

where our time variable is the cosmic expansion factor a of the background, $A_i \in \mathbb{R}^+$, $k_i = n_i k_f$ with $n_i \in \mathbb{N}^+$ and $k_f = 2\pi/L$, and $\varphi_i \in [0, 2\pi)$. Due to the planar symmetry of the perturbations, δ_{lin} only depends on the Cartesian spatial coordinate x . The overdensity in Eq. (2.1) contains four Fourier modes,

$$\delta_{\text{lin}}(x, a = 1) = \frac{1}{2} \left[A_1 e^{-i\varphi_1} e^{ik_1 x} + A_1 e^{i\varphi_1} e^{-ik_1 x} + A_2 e^{-i\varphi_2} e^{ik_2 x} + A_2 e^{i\varphi_2} e^{-ik_2 x} \right] , \quad (2.2)$$

with their amplitudes paired so that δ_{lin} is real valued.

In order for δ_{lin} to represent a realisation of a Gaussian random field, the phases φ_i must be independently drawn from a uniform probability distribution within $[0, 2\pi)$ and the coefficients $A_i/2$ must be independently drawn from a Rayleigh distribution with scale

parameter (i.e. its mode) equal to $\sqrt{P_{\text{lin}}(k_i, a=1)}$ with P_{lin} the power spectrum of the random field. However, we do not consider a Gaussian random field in this work. Rather, we set $A_1 = 0.05$ for the long-wavelength perturbation with $n_1 = 1$, and $A_2 = 0.5$ for the short-wavelength one with $n_2 = 11$. We further set $\varphi_1 = \pi$, and consider eight different values for φ_2 equispaced between 0 and 2π . The purpose of these different initial conditions is to create an ensemble for which the linear short-wavelength modes average to zero.

3 *N*-body simulations

We follow the non-linear evolution of the perturbations in Eq. (2.1) with the publicly available⁴ *N*-body code `cosmo_sim_1d` which is based on an exact force calculation and uses a drift-kick-drift symplectic integrator [37]. We consider $N_p = 125,000$ equal-mass particles (representing two-dimensional slabs) and use the Zel’dovich solution [38] to generate initial conditions⁵ at $a = 0.5$ by displacing the particles from a uniform grid. At this time, the amplitudes of the linear overdensity in the short- and long-wavelength perturbations are (respectively) 0.25 and 0.025.

To calculate the spatial dependence of the moments and cumulants from the *N*-body particle data, we interpolate the particles onto a regular grid with 10,000 grid points. Since all properties of our system scale with the box size L , we express all length scales in units of L , wavevectors in units of k_f , and velocities in units of $H_0 L$ (where H_0 is the Hubble constant of the background). We label the mass of the *N*-body particles as $M_p \propto N_p^{-1}$.

We name our primary simulation (with $\varphi_2 = \pi$) `sim_1_11`, referring to the non-zero modes in the linear overdensity. The left panel of Figure 1 shows the evolution of the overdensity field for this simulation. At $a = 0.5$, we can see the 11 peaks that reflect the short-wavelength mode ($n_2 = 11$) in the overdensity field. The central peak has the largest amplitude, while peaks on either side are progressively weaker. This is due to the long-wavelength mode. At later times, these initial peaks are the sites of formation of 1D halos (3D pancakes). The first shell crossing occurs at $a = 1.82$, leading to a multi-stream structure (the phase-space distribution of the innermost halo is shown in Figure 2). The presence of this region signals the breakdown of the assumption of zero velocity dispersion which characterises SPT. The multi-stream structure first expands due to the momentum of the streams, but gravity slows this expansion, and eventually leads to re-collapse with a second set of shell crossings starting at $a = 4.35$. These are visible as spikes within the 1D halos in the bottom row of Figure 1. We also note that the halos cluster towards the central regions of the box, because of the long-wavelength perturbation.

The same evolution can be traced in the middle panel of Figure 1, where we plot the mean peculiar velocity at fixed Eulerian position. Before shell crossing occurs, the velocity field is single-valued (top row). At later times, we can infer the presence of multi-streaming regions from discontinuities in the mean velocity around the overdensity peaks (bottom two rows). In the right panel, we show the velocity dispersion for our system. Before shell crossing, this quantity is zero, but grows after shell crossing within the 1D halos. We also indicate the fraction of particles in multi-stream regions – ν ; by $a = 6$ (bottom row), $\nu \approx 0.5$.

In Figure 3, we visualise the power spectrum of the overdensity field in `sim_1_11` at two epochs: $a = 1.71$ (the snapshot before the first shell crossing), and $a = 6$ (the final snapshot

⁴https://bitbucket.org/ohahn/cosmo_sim_1d/

⁵Note that the Zel’dovich solution is exact in 1D before shell crossing, removing the need to start the simulations at an earlier epoch.




Figure 1: Time evolution of the overdensity, mean velocity, and velocity dispersion for **sim_1_11**. Shell crossing first occurs at $a = 1.82$, followed by the formation of 1D halos. In the right panel, we indicate the fraction of particles located in multi-stream regions, ν .

in our run). To directly quantify the impact of the high-amplitude, short-wavelength modes, we also run a simulation (hereafter **sim_1**) with $A_2 = 0$ i.e., where only the long-wavelength perturbation with $n_1 = 1$ and $A_1 = 0.05$ is present in the initial conditions. For $k < 6k_f$, the spectra extracted from **sim_1** (red open circles) and **sim_1_11** (blue disks) both decrease nearly exponentially with k as power cascades down to shorter scales from the $k = k_f$ mode at both times. For $k \geq 6k_f$, the **sim_1_11** spectra has excess power generated from the linear modes with $k = 11k_f$, while the **sim_1** spectrum continues to decay exponentially. At $a = 1.71$, the **sim_1_11** power peaks at $k = 11k_f$, while mode coupling suppresses the power at this wavenumber at $a = 6$. In order to better visualise the small deviations at large scales, we plot the fractional difference between the spectra for $k \leq 3k_f$ in the inset. This is the ratio between the power measured in **sim_1** and in **sim_1_11**, minus one. Note that the presence of the high-frequency fluctuation in the initial conditions suppresses the power of the long-wavelength modes. At $a = 1.71$, the offset between the two simulations is between 0.3 and 3%, and, by $a = 6$, the feedback from large- k modes amplifies this to between 1 and 8%. In the remainder of the paper, we will investigate whether perturbation theory and, in particular, the EFTofLSS can accurately and self-consistently model this feedback.

4 Top-down EFT coefficients

In this section, we treat the N -body simulations described above as the UV-complete theory from which we derive the effective dynamics for the coarse-grained momentum moments of the dark matter particles. We start by reviewing the governing equations of the EFTofLSS.




Figure 2: The particle distribution in phase-space for `sim_1_11`. We zoom-in on the innermost structure to show its interior more clearly. The different panels refer to (in order of increasing a) the beginning of the simulation, the first shell crossing, the second shell crossing, and the last snapshot of the simulation.

4.1 Transport equations for the effective fluid

The Vlasov-Poisson system of equations provides a mean-field mathematical model for a self-gravitating collection of a large number of particles. It describes the evolution of the mean⁶ one-particle phase-space density, $f(\mathbf{x}, \mathbf{p}, t)$, and neglects the local fluctuations due to the discreteness of the particles.

In order to study the growth of cosmological perturbations, we take the momentum moments/cumulants of the Vlasov equation to obtain transport equations for a set of ‘fluid-like’ variables. The (physical) mass density is obtained by integrating f over $\mathbf{p}$ and rescaling by the dilution factor a^3 ,

$$\rho(\mathbf{x}, t) = \frac{m}{a^3(t)} \int f(\mathbf{x}, \mathbf{p}, t) d^3p = \rho_b(t) [1 + \delta(\mathbf{x}, t)] , \quad (4.1)$$

where m denotes the particle mass, ρ_b is the background density and δ is the density contrast. In an expanding universe, the canonical momentum conjugate to the comoving coordinate $\mathbf{x}$

⁶The average can be interpreted either over an ensemble of realisations corresponding to the same macroscopic state or as a coarse-graining over (mesoscopic) phase-space patches containing a large number of particles.




Figure 3: The N -body power spectra for `sim_1` (red open circles) and `sim_1_11` (blue disks) at $a = 1.71$ (left panel) and $a = 6$ (right panel). We indicate the linear power spectra evaluated at the same expansion factors with dashed vertical lines. *Inset:* the fractional difference between the spectra for $k \leq 3k_f$, i.e. the ratio of the power measured in `sim_1` to that measured in `sim_1_11` minus one.

is $\mathbf{p} = m a^2 \dot{\mathbf{x}}$, where the dot indicates differentiation with respect to cosmic time t and the product $\mathbf{v} = a \dot{\mathbf{x}}$ is commonly referred to as the peculiar velocity of the particle. Therefore, the average particle velocity is

$$\mathbf{V}(\mathbf{x}, t) = \frac{\int \frac{\mathbf{p}}{m a(t)} f(\mathbf{x}, \mathbf{p}, t) d^3 p}{\int f(\mathbf{x}, \mathbf{p}, t) d^3 p} \quad (4.2)$$

which can be expressed as

$$\mathbf{V}(\mathbf{x}, t) = \frac{\mathbf{P}(\mathbf{x}, t)}{\rho(\mathbf{x}, t)} \quad (4.3)$$

with

$$\mathbf{P}(\mathbf{x}, t) = \frac{1}{a^4(t)} \int \mathbf{p} f(\mathbf{x}, \mathbf{p}, t) d^3 p. \quad (4.4)$$

Similarly, the second velocity moment is given by the dyadic tensor

$$\Sigma(\mathbf{x}, t) = \frac{\int \frac{\mathbf{p}}{m a(t)} \frac{\mathbf{p}}{m a(t)} f(\mathbf{x}, \mathbf{p}, t) d^3 p}{\int f(\mathbf{x}, \mathbf{p}, t) d^3 p} \quad (4.5)$$

which can be written as

$$\Sigma(\mathbf{x}, t) = \frac{\Xi(\mathbf{x}, t)}{\rho(\mathbf{x}, t)}, \quad (4.6)$$

with

$$\Xi(\mathbf{x}, t) = \frac{1}{ma^5(t)} \int \mathbf{p} \mathbf{p} f(\mathbf{x}, \mathbf{p}, t) d^3p . \quad (4.7)$$

This tensor is further decomposed into an unconnected streaming part and a connected dispersion part:

$$\Xi(\mathbf{x}, t) = \rho(\mathbf{x}, t) \mathbf{V}(\mathbf{x}, t) \mathbf{V}(\mathbf{x}, t) + \mathbf{\Gamma}(\mathbf{x}, t) . \quad (4.8)$$

The caveat in taking successive moments in this way is that the equation for the n^{th} moment depends on the $(n+1)^{\text{th}}$ moment. Thus, a closure condition is needed to truncate the Vlasov hierarchy. SPT assumes that structure formation is driven by matter with negligible velocity dispersion (cold dark matter in the single-stream regime) and thus sets the velocity-dispersion tensor $\mathbf{\Gamma}$ to zero (this approximation is expected to break down with time as more and more collapsed structures form [39, 40]).

The governing equations are thus the mass and momentum-conservation equations coupled to the Poisson equation for the peculiar gravitational potential ϕ :

$$\dot{\delta} + \frac{1}{a} \nabla \cdot [(1 + \delta) \mathbf{V}] = 0 , \quad (4.9)$$

$$\dot{\mathbf{V}} + \frac{\dot{a}}{a} \mathbf{V} + \frac{1}{a} (\mathbf{V} \cdot \nabla) \mathbf{V} = -\frac{1}{a} \nabla \phi , \quad (4.10)$$

$$\nabla^2 \phi = 4\pi G a^2 \rho_b \delta . \quad (4.11)$$

In SPT, these equations are solved perturbatively by expanding δ and $\nabla \cdot \mathbf{V}$ (as no vorticity can be generated if $\mathbf{\Gamma} = 0$) about the linear solutions [3].

In the EFTofLSS framework, the Vlasov-Poisson system is first smoothed with a low-pass spatial filter which removes small-scale (but still macroscopic) fluctuations. As a second step, momentum moments of the (coarse-grained) Vlasov equation are taken. As conventional in the literature, we use the subscript ℓ – short for ‘long-wavelength part’ – to distinguish the fluid-like variables that have been smoothed. The governing equations for the first two moments in this case are:

$$\dot{\delta}_\ell + \frac{1}{a} \nabla \cdot [(1 + \delta_\ell) \mathbf{U}] = 0 , \quad (4.12)$$

$$\dot{\mathbf{U}} + \frac{\dot{a}}{a} \mathbf{U} + \frac{1}{a} (\mathbf{U} \cdot \nabla) \mathbf{U} = -\frac{1}{a} \nabla \phi_\ell - \frac{1}{a \rho_\ell} \nabla \cdot \boldsymbol{\tau} , \quad (4.13)$$

$$\nabla^2 \phi_\ell = 4\pi G a^2 \rho_b \delta_\ell , \quad (4.14)$$

where $\mathbf{U} = \mathbf{P}_\ell / \rho_\ell$ and $\boldsymbol{\tau}$ denotes the effective stress tensor which quantifies the dynamical coupling to the degrees of freedom suppressed by the smoothing procedure. The EFTofLSS closes the system of equations by modelling $\boldsymbol{\tau}$.

Considering perturbations with planar symmetry greatly simplifies the mathematical formulation as we can pick a reference system in which all fluid functions depend only on one Cartesian coordinate, say x^1 (simply x hereafter). Moreover, $U^i = U \delta_1^i$, $\tau^{ij} = \tau \delta_1^i \delta_1^j$ and $\Xi_\ell^{ij} = \Xi_\ell \delta_1^i \delta_1^j$ where δ_j^i denotes the Kronecker symbol. From now on, we refer to τ as the *effective stress*.

Reference [41] (see also e.g. [12]) has shown that the effective stress naturally decomposes into two pieces

$$\tau = \tau_k + \tau_g , \quad (4.15)$$

where the kinetic part τ_k accounts for the velocity dispersion within a smoothing patch

$$\tau_k = \Xi_\ell - \rho_\ell U^2, \quad (4.16)$$

and the gravitational part τ_g quantifies the departures from the mean-field peculiar acceleration $-\partial_x \phi_l / a$

$$\tau_g = \frac{1}{8\pi G a^2} \left\{ [(\partial_x \phi)^2]_\ell - (\partial_x \phi_\ell)^2 \right\}. \quad (4.17)$$

4.2 Coarse-graining and non-linear scales

We now return to our N -body simulations and measure the effective stress in their snapshots at various output times using Eqs. (4.15–4.17). First, we smooth all the relevant quantities with a low-pass filter defined by a kernel $W_\Lambda(x)$. Specifically, for a generic field g , we compute the convolution

$$g_\ell(x) = \int W_\Lambda(x - x') g(x') dx'. \quad (4.18)$$

We use the FFT algorithm to perform convolutions in Fourier space, $\tilde{g}_l(k) = \tilde{W}_\Lambda(k) \tilde{g}(k)$. To check whether our results depend on the choice of the kernel, we use both a sharp k -space cutoff, $\tilde{W}_\Lambda(k) = \Theta(\Lambda - k)$ where Θ is the Heaviside step function, and a smooth Gaussian cutoff, $\tilde{W}_\Lambda(k) = \exp[-k^2/(2\Lambda^2)]$. The key difference between these two choices is that the Gaussian filter receives contributions from modes with $k > \Lambda$.

As previously mentioned, the EFTofLSS should be applicable only when the hierarchy of scales $k \ll \Lambda \ll k_{\text{NL}}$ holds true. It is thus necessary to first get an estimate of k_{NL} in order to pick an appropriate value for Λ . We do this by taking the inverse of the mean (or median) displacement of the simulation particles (relative to their initial position at $a = 0$) at each epoch. As an example, we plot the evolution of k_{NL} with a for `sim_1_11` in Figure 4. The non-linear wavenumber rapidly decreases with time and, at the epoch of first shell crossing (highlighted with a vertical dashed line), $k_{\text{NL}} \simeq 15 k_f$. This is a reasonable estimate given that the shell crossing is generated by the collapse of the linear perturbation with $k = 11 k_f$.

In what follows, we will be mainly concerned with the non-linear evolution of the longest perturbation in the simulation box, i.e. $k = k_f$. Since $k_f \ll k_{\text{NL}}$ for $a < 6$, it should be possible to treat the Fourier modes of δ with wavenumber $|k| = k_f$ perturbatively. With the above discussion in mind, we need to pick a value of Λ that is simultaneously much smaller than the wavenumber of the large- k mode ($11 k_f$), and larger than the scale of interest. The latter is the largest scale in our system ($k = k_f$). Thus, $\Lambda = 3 k_f$ or $4 k_f$ are natural choices. For most of our analysis, we will make the choice $\Lambda = 3 k_f$ (i.e. the sharp- k -space filter retains the modes with $|k| = k_f, 2 k_f$, and $3 k_f$) which is three times larger than the scale of interest. Thus, $\Lambda/k = 3$, and k_{NL}/Λ should be similar to ensure that the hierarchy $k \ll \Lambda \ll k_{\text{NL}}$ holds. The EFTofLSS should then be applicable until $k_{\text{NL}} = 9 k_f$ which corresponds to $a \approx 3$ (see Figure 4). In principle, a larger value of Λ may also be chosen, as long as it remains much smaller than $11 k_f$. We have checked that our results do not depend on the smoothing scale for $\Lambda \leq 8 k_f$.

4.3 Effective stress

The effective stress measured in `sim_1_11` is shown in Figure 5 (solid blue line) together with the different contributions defined in Eqs. (4.15–4.17). The coarse-grained second moment Ξ_ℓ is split into contributions from matter in the single-stream regime (Ξ_ℓ^s) and in the multi-streaming regime (Ξ_ℓ^m). At early times, when the flow is in the single-stream regime (i.e.




Figure 4: Time evolution of k_{NL} for `sim_1_11`. Here, k_{NL} denotes the wavenumber corresponding to the mean (solid blue line) or median (dash-dotted red line) particle displacement. We indicate three important scales with horizontal dashed lines: $k = k_f$ is the scale of interest in our work, $k = 3k_f$ is the coarse-graining scale we mostly use, and $k = 9k_f$ is the smallest k_{NL} which satisfies the ordering of scales $k \ll \Lambda \ll k_{\text{NL}}$. The vertical dashed line highlights the epoch of first shell crossing.

$\Gamma = 0$), we can use linear perturbation theory to predict the time dependence of the different contributions to τ . Since the peculiar gravitational potential ϕ does not evolve with time and the peculiar velocity of the fluid scales as $a^{1/2}$, it follows that both τ_k and τ_g should decrease proportionally to a^{-2} . In Figure 5, we therefore multiply all the components by a^2 to cancel out this linear evolution.

The first row displays the sharp-cutoff case. At $a = 1.71$ (left panel), the flow is still in the single-stream regime, and the effective stress receives a slightly larger contribution from the kinetic part than the gravitational part. At later times, we expect the scalings discussed above to not hold, as the system becomes more and more non-linear. After shell crossing, $a^2\tau$ decays with time. In the middle panel ($a = 3.03$), we note a non-zero contribution from the multi-stream regime via Ξ_ℓ^m . In the right panel ($a = 6$), we observe oscillations in the effective stress, with four prominent peaks and valleys.

In the second row, we plot the different contributions when using a Gaussian smoothing. At all times, the total stress τ is larger in comparison to the sharp case. This excess comes from modes with $k > \Lambda$ which provide damped but non-zero contributions even at early times. Despite noticeable differences between the effective stress measured from the two smoothing schemes, the results are in rough qualitative agreement.

4.4 Effective speed of sound

In order to close the system of equations (4.12–4.14), it is necessary to model the effective stress in terms of the fluid variables δ_ℓ and U . This seems to be an impossible task as τ contains a dependence on short-wavelength modes. To address this issue, the EFTofLSS assumes that




Figure 5: The effective stress in `sim_1_11` decomposed into its constituent pieces as described in Eqs. (4.15–4.17) for sharp cutoff (top row) and Gaussian smoothing (bottom row). The superscript over Ξ_ℓ indicates whether the contribution is made by particles within single-stream (s) or multi-stream (m) regions.

the effective stress can be decomposed into a deterministic part (which depends on δ_ℓ and U) and a stochastic part (uncorrelated with the long-wavelength fields). The former quantifies the tidal effects of the long modes onto the particle trajectories and can be thought of as the conditional expectation value of τ over the short-wavelength perturbations at fixed long-wavelength background [11, 12, 17], $\langle \tau | \text{fixed background} \rangle$ (in the remainder of this paper, we denote this quantity with the symbol $\langle \tau \rangle$). However, it is not clear if and how the ensemble can be constructed in practice (e.g. using numerical simulations). In fact, this would require being able to build identical replicas of the non-linearly evolved long modes with different sets of short modes which is a formidable problem as the short modes influence the evolution of the long ones. We discuss two possible ways of addressing this issue in Sections 4.4.1 and 4.4.2.

For the moment, we assume that $\langle \tau \rangle$ exists. The next step is to expand $\langle \tau \rangle$ in terms of the long-wavelength fields. Since the characteristic time scales with which long- and short-wavelength perturbations evolve are similar, $\langle \tau \rangle$ at a given time should depend on the entire past history of, say, δ_ℓ averaged over a kernel with a width of the order of the Hubble time. This concept is expressed concisely by saying that the effective theory is ‘non-local in time’ [e.g. 12, 17, 18]. It turns out, however, that perturbative solutions can be obtained by expanding $\langle \tau \rangle$ as a sum of operators all evaluated at the same time. The non-locality of the theory will be reflected in the fact that the effective coefficients of a given operator will not coincide at different perturbative orders (see also footnote 3). This is not a concern for this paper which only considers the leading-order corrections to the SPT power spectrum. Then, we expand $\langle \tau \rangle$ in derivatives and powers of the long-wavelength perturbations δ_ℓ and θ_ℓ (with $\theta_\ell = \partial_x U/aH$) including all terms allowed by symmetries. Writing explicitly only the terms to linear order in δ_ℓ and θ_ℓ , we obtain

$$\tau = \langle \tau \rangle + J \approx p_b + \rho_b c_s^2 \delta_\ell + \rho_b c_v^2 \theta_\ell + \dots + J, \quad (4.19)$$

where p_b is the effective background pressure⁷ (which, being constant in space, has no dynamical effects), c_s^2 is the speed of sound in the effective fluid, c_v^2 is a (bulk-)viscosity coefficient expressed in units of speed (in doing so, we divide the viscosity by ρ_b), the ellipses represent higher-order terms, and J is the stochastic term. Note that c_s^2 and c_v^2 are assumed to not vary with position but are time dependent. The stochastic term is a stationary random field which varies in space and time.

In the following subsections, we discuss different methods to estimate c_s^2 and c_v^2 from the simulations.

4.4.1 Estimates from spatial cross-correlations

The simplest method to estimate the EFT coefficients is to assume that the leading-order expansion in Eq. (4.19) is exact and measure the spatial cross-correlation between τ and either δ_ℓ or θ_ℓ in a single simulation. This approach circumvents the problem of defining an ensemble in order to measure $\langle \tau \rangle$. Since the stochastic term should be uncorrelated with the long-wavelength fields, we obtain

$$\langle \delta_\ell^2 \rangle_x c_s^2 + \langle \theta_\ell \delta_\ell \rangle_x c_v^2 = \langle \tau \delta_\ell \rangle_x / \rho_b \quad (4.20)$$

$$\langle \theta_\ell \delta_\ell \rangle_x c_s^2 + \langle \theta_\ell^2 \rangle_x c_v^2 = \langle \tau \theta_\ell \rangle_x / \rho_b \quad (4.21)$$

where the symbol $\langle \dots \rangle_x$ denotes averaging over different locations in space (related approaches have been used in the spatial domain by [12], and in the wavenumber domain by [11, 33]).

At early times, when the long-wavelength perturbations are quasi-linear, $\theta_\ell \simeq -\delta_\ell$ and the system given in Eqs. (4.20) and (4.21) is nearly degenerate. In this regime, the expansion of the effective stress reduces to

$$\tau \approx p_b + \rho_b (c_s^2 - c_v^2) \delta_\ell + \dots + J, \quad (4.22)$$

as, in practice, there is only one dynamical variable. Although δ_ℓ and $-\theta_\ell$ differ more and more as the long-wavelength perturbations evolve into the non-linear regime, it turns out that

⁷The literature on the EFTofLSS defines τ with the opposite sign with respect to the standard convention used in fluid dynamics. Therefore, τ coincides with the pressure tensor defined in the study of non-equilibrium continua.




Figure 6: Time evolution of the fields $(\delta_\ell + \theta_\ell)/a$ (solid black) and $(\delta_\ell - \theta_\ell)/a$ (dashed blue) in `sim_1_11` for sharp- k (top row) and Gaussian filtering (bottom row) together with $1 + r$, where r is the linear cross-correlation coefficient between δ_ℓ and θ_ℓ .

these fields nearly coincide at all times probed by our simulations. This is demonstrated in Figure 6 where we show that $(\delta_\ell - \theta_\ell)/a$ is always substantially larger than $(\delta_\ell + \theta_\ell)/a$ when $\Lambda = 3 k_f$. The linear correlation coefficient r between δ_ℓ and θ_ℓ (also reported in the figure) is always very close to -1 (i.e., $1 + r$ is close to zero). Fortunately, the coefficient matrix of the system (4.20) and (4.21) is numerically invertible for $a \geq 0.5$, and we can still obtain an estimate for c_s^2 and c_v^2 . We label this estimator SC (short for spatial correlations). To cross-check that our results are not influenced by the numerical matrix inversion, we also use the following estimator for the combination $c_{\text{tot}}^2 = c_s^2 - c_v^2$,

$$c_{\text{tot}}^2 = \frac{1}{\rho_b} \frac{\langle \tau \delta_\ell \rangle_x}{\langle \delta_\ell^2 \rangle_x}, \quad (4.23)$$

obtained from Eq. (4.22). We dub this estimator $\text{SC}\delta$.

Reference [33] used a similar estimator in the wavenumber domain and further assumed that the linear density perturbation δ_{lin} provides an accurate model on large scales,

$$c_{\text{tot}}^2 = \frac{1}{\rho_{\text{b}}} \frac{\sum_{k < \Lambda} \tilde{\tau}(k) \tilde{\delta}_{\text{lin},l}(-k)}{\sum_{k < \Lambda} \left| \tilde{\delta}_{\text{lin},l}(k) \right|^2}, \quad (4.24)$$

where the tilde denotes a Fourier transform. To connect with the previous literature, we also include this estimator (which we dub the M&W estimator) in our analysis.

4.4.2 Estimates from least-squares fitting

We can also estimate c_{s}^2 and c_{v}^2 by fitting the expansion in Eq. (4.19) to the simulation data. However, if τ , δ_{ℓ} and θ_{ℓ} are measured from a single simulation, the least-squares method gives back Eqs. (4.20) and (4.21) for the best-fitting values.

Alternatively, we can first compute $\langle \tau \rangle$ and then perform the fit. As discussed above, constructing the average $\langle \tau \rangle$ is a considerable challenge as it requires conditioning on non-linearly evolved fields. It is easier, instead, to build an ensemble of realisations in which the *linear* long modes are identical and the short modes vary. This is a well-defined ensemble which can be expressed in terms of a probability distribution for the linear modes. Therefore, we compute $\langle \tau \rangle$ by averaging at fixed x over the eight simulations with different values of the angle φ_2 (i.e. different phases between the $k = k_{\text{f}}$ and $11 k_{\text{f}}$ modes) described in Section 2. We note that this ensemble average is nearly identical to the τ measured from `sim_1_11`, especially in the sharp-cutoff case.

Starting from the δ_{ℓ} , θ_{ℓ} , and $\langle \tau \rangle$ estimated by averaging at fixed x over the eight realisations, we fit the coefficients of the expression

$$\langle \tau \rangle = C_0 + C_1 \delta_{\ell} + C_2 \theta_{\ell} \quad (4.25)$$

using the least-squares method. The fit parameters C_i are related to the expansion in Eq. (4.19). We refer to the resulting values for c_{s}^2 and c_{v}^2 as the F3P estimates (short for ‘fit using three parameters’). To check the robustness of the results, we also fit $\langle \tau \rangle$ using all possible terms up to second-order in δ_{ℓ} and θ_{ℓ} , i.e. by adding $C_3 \delta_{\ell}^2 + C_4 \theta_{\ell}^2 + C_5 \delta_{\ell} \theta_{\ell}$ to the right-hand side of Eq. (4.25). We refer to the estimates for c_{s}^2 and c_{v}^2 obtained this way using the label F6P (short for ‘fit using six parameters’). Figure 7 shows the accuracy of the fits (which are indicated by τ_{fit}) using $\Lambda = 3 k_{\text{f}}$. For the sharp filter (top row), both the three- and six-parameter fits match the estimated $\langle \tau \rangle$ well. At $a = 6$, however, F6P better reproduces the peaks and valleys of the high-frequency oscillations seen in $\langle \tau \rangle$. When using Gaussian smoothing (bottom row), both F3P and F6P accurately reproduce $\langle \tau \rangle$ around the centre of the simulation volume, while F6P provides a better match to the measured stress near the edges of the box at all times.

4.4.3 Comparing the different estimates

At this point, it is instructive to compare the c_{s}^2 and c_{v}^2 values obtained from the two classes of estimators we have discussed. In Figure 8, we show the results obtained applying the F3P (solid black line), F6P (dash-dotted green line) and SC (dashed magenta line) methods to `sim_1_11` for the sharp-cutoff case. The three estimates are in excellent agreement with each




Figure 7: Fits to the averaged effective stress (using Eq. 4.25, F3P, or its extension to quadratic order, F6P), with a sharp cutoff at $\Lambda = 3 k_f$. The solid blue line is the stress estimated from the simulations, and the dashed lines represent the fits using the F3P (black) and the F6P (red) models.

other, thus suggesting that we can robustly measure the IR-UV coupling coefficients of the effective theory.

The best-fit values for c_s^2 and c_v^2 are similar to each other. Thus, their difference c_{tot}^2 (the only combination entering the corrections to the matter power spectrum, see Section 5.1) is much smaller than the individual values. In the inset of Figure 8, we show the error ellipse for c_s^2 and c_v^2 obtained with the F3P method at $a = 3.03$. The statistical uncertainties on the individual fit parameters are approximately 10% but their difference is much more tightly constrained since the coefficients are highly correlated, as expected from the fact that $\delta_\ell \approx -\theta_\ell$.

In Figure 9, we show c_{tot}^2 for **sim_1_11** as a function of time for the top-down estimators in case of a sharp cutoff with $\Lambda = 3 k_f$ (for most of our following results, we will relegate the plots for Gaussian smoothing to Appendix A). The dash-dotted green line represents the scaling $c_{\text{tot}}^2 \propto a$, which is the time-dependence assumed in the EFTofLSS literature based on perturbative considerations [e.g. 12, 13, 16, 33]. All estimates satisfy this relation when the

$$\Lambda = 3 k_f \text{ (sharp cutoff)}$$




Figure 8: The c_s^2 and c_v^2 values obtained with the F3P, F6P, and SC estimators using a sharp filter with $\Lambda = 3 k_f$. The dashed vertical lines indicate the times when a shell crossing is detected in `sim_1_11`. *Inset:* The 68% confidence ellipse for c_s^2 and c_v^2 at $a = 3.03$ for the F3P estimator.

flow is in the single-stream regime. As the first set of shell crossings occur, however, the linear growth stops and c_{tot}^2 steadily decreases with time until $a = 6$. Because of this, c_{tot}^2 assumes its peak value at $a = 2.04$. It has been argued that virialised structures should decouple from the long-wavelength dynamics and stop contributing to c_{tot}^2 [11]. The decrease we see in c_{tot}^2 might then reflect the fact that the 1D halos in `sim_1_11` are approaching or have reached virialisation. However, c_{tot}^2 does not vanish because only half of the particles lie in multi-stream regions by $a = 6$. These conjectures suggest that c_{tot}^2 is sourced by perturbations that are entering the non-linear regime. In a sense, our toy model allows us to identify the response of the effective coupling constant to a single scale becoming non-linear. It is thus reasonable to expect that, in the presence of a continuous spectrum of perturbations, the time evolution of c_{tot}^2 mirrors the flow of power into the non-linear regime. Lastly, we note that the M&W and SC δ estimators give slightly larger values for c_{tot}^2 . This overestimation likely comes from assuming that $\delta_\ell = -\theta_\ell$.

5 Performance of the EFTofLSS

5.1 Next-to-leading-order power spectrum in SPT and EFTofLSS

SPT solves the system of equations (4.9–4.11) perturbatively. It is convenient to work in Fourier space, where spatial derivatives transform to products by the wavenumber. In the 1D




Figure 10: The solid blue line appearing in both panels shows the time evolution of the matter power spectrum in `sim_1_11` evaluated at $k = k_f$. The matter density has been smoothed using a sharp filter with $\Lambda = 3 k_f$. In the left panel, the N -body result is compared to the NTLO predictions from SPT (dotted pink line) and tSPT (dash-dotted brown line). Two important timescales are indicated with vertical lines: the time of the first shell crossing (dashed teal line) and the time at which $k_{\text{NL}}/k_f = 11$ (dotted magenta line). In the right panel, the N -body result is contrasted with the NTLO predictions from the EFTofLSS obtained using five different top-down estimators for c_{tot}^2 . The dashed vertical lines are as in Figure 8. The bottom panels show the percentage deviation of the models from the simulation. Note that the vertical scale differs in the left and right panels. In the bottom right panel, we omit tSPT to better visualise the EFT errors.

Next, we investigate the effect of introducing a cutoff in the SPT loop integrals. The dash-dotted brown line in Figure 10 shows the NTLO SPT spectrum obtained using a sharp cutoff with $\Lambda = 3 k_f$. We refer to this model as truncated SPT (tSPT, in short). With our initial conditions, this is equivalent to only accounting for the mode coupling between the two linear modes with $k = \pm k_f$. Therefore, the difference between the SPT and tSPT models shows the contribution of the linear modes with $k = \pm 11 k_f$ to $P_{\text{SPT}}(k_f)$. Their net effect is to slow down the growth of the long-wavelength perturbation. The tSPT model overpredicts $P(k_f)$ at all times. Early on, when $k_{\text{NL}} < 11 k_f$, and the linear short-wavelength modes in our setup are perturbative, SPT is much more accurate than tSPT. The error in SPT remains smaller than in tSPT until $a \simeq 5$, at which time they cross over. The tSPT error at $a = 6$ is $\approx 1.5\%$.

In the EFTofLSS, the expansion of Eq. (4.19) is substituted into the 1D equivalents of Eqs. (4.12) and (4.13) to obtain (along with the Poisson equation) the dynamical equations

for the effective fluid

$$\dot{\delta}_\ell + \frac{1}{a} \partial_x [(1 + \delta_\ell) U] = 0 , \quad (5.3)$$

$$\dot{U} + \frac{\dot{a}}{a} U + \frac{1}{a} U \partial_x U = -\frac{1}{a} \partial_x \phi_\ell - \frac{c_s^2}{a} \partial_x \delta_\ell - \frac{c_v^2}{a} \partial_x \theta_\ell - \frac{1}{a \rho_\ell} \partial_x J . \quad (5.4)$$

At linear order in the density perturbations,

$$\left(\partial_t^2 + 2 \frac{\dot{a}}{a} \partial_t - 4\pi G \rho_b \right) \delta_{1\ell} = \frac{1}{a^2} c_{\text{tot}}^2 \partial_x^2 \delta_{1\ell} + \frac{1}{a^2 \rho_\ell} \partial_x^2 J , \quad (5.5)$$

which differs from the corresponding tSPT equation because the terms on the right side do not vanish⁸. Neglecting for a moment the contribution from the stochastic term which will be discussed separately in Section 5.2, the leading-order correction to the tSPT solution is

$$\widetilde{\delta}_\ell^c(k, a) = \alpha_c a k^2 \widetilde{\delta}_{1\ell}(k) , \quad (5.6)$$

where

$$\alpha_c = \frac{1}{a} \int_0^\infty G(a, a') c_{\text{tot}}^2(a') a' da' , \quad (5.7)$$

and, in the EdS universe, the retarded Green's function assumes the form [e.g. 16, 20]

$$G(a, a') = \Theta(a - a') \frac{2}{5H_0^2} \left[\left(\frac{a'}{a} \right)^{3/2} - \frac{a}{a'} \right] . \quad (5.8)$$

It follows that the EFTofLSS power spectrum to NTLO is

$$P_{\text{EFT}} = P_{\text{tSPT}} + 2 \alpha_c k^2 \widetilde{W}_\Lambda^2 P_{11} . \quad (5.9)$$

We obtain top-down estimates for α_c by numerically integrating Eq. (5.7) and using the functions $c_{\text{tot}}^2(a)$ measured from the simulations and presented in Figure 9. Since our simulations start at $a = 0.5$, we assume $c_{\text{tot}}^2 \propto a$ for $a < 0.5$. As shown in the right panel of Figure 10, the resulting power spectra provide an excellent match to the N -body results for $k = k_f$. Note, however, that the spectra are never asymptotically correct in any regime, and it is not clear how this could be rectified at higher order. Nonetheless, deviations are always smaller than 0.06% even at $a = 6$ when $\Lambda = 3 k_f$ is only 1.5 times larger than k_{NL} . This result provides evidence that the effective dynamics accurately models the feedback of the short modes onto the growth rate of the long ones. All the top-down estimators for c_{tot}^2 we used essentially give the same result with the sharp filter while some dissimilarities are noticeable for Gaussian smoothing (see Appendix A).

Assuming – as done in the literature [e.g. 12, 13, 16, 33] – that $c_{\text{tot}}^2 \propto a$ for all a , we can write $c_{\text{tot}}^2 = c_0^2 a$ (where c_0^2 is the estimated value at $a = 1$). Using this expression, we can analytically integrate Eq. (5.7), which gives the relation

$$\alpha_c = -\frac{c_0^2 a^2}{9H_0^2} . \quad (5.10)$$

⁸References [42, 43] have recently introduced a formalism which also includes higher-order cumulants in the perturbative expansion.




Figure 11: Comparison of α_c between the bottom-up estimator (solid green line) and the top-down estimators (dashed lines). The dashed vertical lines are as in Figure 8.

In our setup, using this scaling gives inaccurate results for $a > 3$. This is shown in Figure 11 where we plot the time evolution of the EFT correction to the NTLO power spectrum for $k = k_f$ and $\Lambda = 3 k_f$. The solid line represents the value of α_c that perfectly reproduces the N -body power spectra, i.e.

$$\hat{\alpha}_c = \frac{P_{N\text{-body}} - P_{\text{tSPT}}}{2 k^2 \widetilde{W}_\Lambda^2 P_{11}}. \quad (5.11)$$

This is the ‘bottom-up’ approach predominantly used in the EFTofLSS literature. As expected from the right panel in Figure 10, the top-down estimates for α_c cluster around the bottom-up value. However, the approximation in Eq. (5.10) (dark grey dash-dotted line) departs from $\hat{\alpha}_c$ at late times. This happens because the actual value of c_{tot}^2 scales proportionally to a only before shell crossing takes place and decreases afterwards. As α_c evaluated at a given a retains information from all $a' < a$ as we integrate over the past, the scaling $\alpha_c \propto a^2$ continues to hold until $a \approx 3$. However, at later times, α_c grows slower than the a^2 scaling, consistent with a decreasing value of c_{tot}^2 .

5.2 Stochastic term

In our discussion so far, we have ignored the stochastic part of the effective stress as its impact on the matter power spectrum is expected to be suppressed at $k < \Lambda$ with respect to the deterministic NTLO corrections [11]. The reason is twofold. First, the variance of the stochastic fluctuations in a patch of size Λ is determined by the properties of the short-wavelength fluctuations with $k > \Lambda$ and should be dominated by the perturbations that are turning non-linear with $k \simeq k_{\text{NL}}$. Second, a region of size k^{-1} contains many sub-domains of size Λ^{-1} and thus the amplitude of statistical fluctuations is suppressed by averaging over the large number of sampled sub-regions. In this section, we investigate the properties of the stochastic contribution in our simulations. We compute J as the difference between the effective stress measured from `sim_1_11` and that reconstructed from the c_s^2 and c_v^2 using the SC and F3P estimators. For $a < 3$ we find that J only contributes to τ at the sub-percent

level, while the stochastic fraction reaches 40% at $a = 6$ with the sharp filter. However, most of this contribution comes from higher- k modes, and does not enter the $k = k_f$ calculations. Also, note that our measurement of J is degenerate to higher-order terms in the expansion of $\langle \tau \rangle$. Including the stochastic term in Eq. (5.4), gives a correction to the tSPT overdensity [e.g. 20]

$$\tilde{\delta}_J(k, a) = -\frac{k^2}{\rho_b} \int_0^\infty G(a, a') J(k, a') da' , \quad (5.12)$$

where $G(a, a')$ is the Green's function defined in Eq. (5.8). We find that this term makes negligible contributions to the matter power spectrum at $k = k_f$. The auto-spectrum of δ_J (hereafter, P_{JJ}) is always of the order of a part per million, independently of the estimator used to determine c_{tot}^2 .

5.3 Dependence on the separation of scales

An effective theory assumes that the scales it describes must be well-separated from the UV scales. In our simple setup, we can change the separation between wavelengths and see how the different flavours of perturbation theory compare with the numerical solution as a consequence. To this end, we run two further simulations with the same amplitude ratio as `sim_1_11`, but with $n_1 = 1, n_2 = 7$ (hereafter `sim_1_7`), and $n_1 = 1, n_2 = 15$ (hereafter `sim_1_15`). We show the effect of this change in separation in the left panel of Figure 12, by plotting the non-linear evolution of the power in $k = k_f$ obtained from the simulations (solid lines). We also show the SPT power spectra in each case (dashed lines). For `sim_1` (the simulation run with only the long-wavelength mode), SPT matches the N -body spectrum well at all times, as there is no linear high-amplitude mode. Comparing `sim_1` and the other simulations, we see that presence of the linear high-amplitude, short-wavelength mode dampens the N -body spectrum at $k = k_f$. The strength of this dampening increases with decreasing separation, as expected. For all simulations, the SPT and N -body spectra agree well until $a \approx 3$. At later times, however, SPT fails to correctly capture the non-linear evolution and underestimates the power by a few per cent. The absolute value of the error (shown in the bottom panel) decreases with increasing the separation between the linear modes. In the right panel, we plot the tSPT spectra (dashed lines) in addition to the N -body results, when using a sharp cutoff with $\Lambda = 3 k_f$. All the tSPT spectra are identical to each other as truncation removes the high-amplitude mode. It follows that tSPT overestimates the power by several percent and the error with respect to the simulations increases with decreasing separation between the linear modes. Finally, we plot the EFT spectra (dotted lines) obtained with the SC estimator for c_{tot}^2 (all the estimators are in good agreement for all simulations). These predictions remain accurate across simulations to $\approx 0.1\%$ at all times. As expected, the accuracy is higher for the initial conditions with the largest separation between the modes. These results show that the EFT corrections are robust to change in separation between the linear modes.

5.4 Considerations about renormalisation

In general, the tSPT loop integrals P_{13} and P_{22} depend on Λ due to truncation of the linear power spectrum which delimits the range of integration. Similarly, the EFT correction in Eq. (5.9) depends on Λ through α_c (and thus c_{tot}^2) which is determined from the smoothed perturbations. Since P_{13} and the EFT correction have the same k dependence as $k \rightarrow 0$, references [12, 13, 16] proposed to renormalise the coupling coefficient between long- and short-wavelength perturbations to get Λ -independent predictions. This mirrors the renormalisation




Figure 12: As in Figure 10, but for a suite of simulations with different separation between the wavelengths of the non-vanishing linear perturbations. Note that: (i) all the tSPT lines coincide in the top-right panel; (ii) the colors of the lines have different meanings from Figure 10; (iii) contrary to Figure 10, we plot the absolute value of the percentage deviation in the bottom panels.

procedure used in field theory to control UV divergences. For the moment, we only consider the renormalisation of the P_{13} contribution. We later remark on a similar scheme for P_{22} and its applicability in our setup. We begin by choosing a ‘renormalisation scale’ $\mu < \Lambda$ below which one can trust the validity of SPT. The ‘bare’ effective parameter α_c is then decomposed into a (Λ -independent) renormalised part and a (Λ -dependent) counterterm

$$\alpha_c(\Lambda, \mu, a) = \alpha_{\text{ren}}(\mu, a) + \alpha_{\text{ctr}}(\Lambda, \mu, a) , \quad (5.13)$$

so that $\alpha_{\text{ren}}(\mu, a)$ accounts for the actual correction to the dynamics generated by the non-perturbative scales while $\alpha_{\text{ctr}}(\Lambda, \mu, a)$ removes the spurious contributions introduced by using SPT beyond its range of validity within the loop integrals. The renormalised parameter α_{ren} is defined by modifying the bare parameter from the cutoff scale Λ to the renormalisation scale μ , in such a way that the physical observable (in our case the EFT power spectrum) is unaltered⁹. In practice, this requires that the counterterm exactly cancels the difference between the full Λ -dependent P_{13} contribution and the renormalised contribution which depends on μ .

$$2 \alpha_{\text{ctr}}(\Lambda, \mu, a) k^2 \widetilde{W}_\Lambda^2(k) P_{11}(k, a) = -2 [P_{13}^{\text{tSPT}}(k, \Lambda, a) - P_{13}^{\text{tSPT}}(k, \mu, a)] . \quad (5.14)$$

⁹As in field theory, this is governed by the renormalisation group equations. For a review of these techniques, see e.g. [44]. For an example of its implementation for the EFTofLSS, see [17].

Since, in the limit¹⁰ $k \rightarrow 0$,

$$P_{13}^{\text{tSPT}}(k, \Lambda, a) \rightarrow -a^2 f(\Lambda) k^2 P_{11}(k, a), \quad (5.15)$$

where

$$f(\Lambda) = \int \frac{P_{11}(k, a=1)}{k^2} \widetilde{W}_\Lambda^2(k) dk \quad (5.16)$$

is the variance of the matter displacement field¹¹ at $a = 1$, one, finally, gets

$$\alpha_{\text{ctr}}(\Lambda, \mu, a) = -a^2 [f(\Lambda) - f(\mu)]. \quad (5.17)$$

Therefore, the full non-linear power spectrum for $k \rightarrow 0$ (ignoring the P_{22} contribution) is

$$P_{\text{EFT}}(k, a) = [1 + 2 \alpha_{\text{ren}}(\mu, a) k^2] P_{11}(k, a) + 2 P_{13}^{\text{tSPT}}(k, \mu, a). \quad (5.18)$$

The final expression for the power spectrum is thus Λ -independent.

In order to determine the terms appearing in Eq. (5.13), we proceed as follows. First, we use Eq. (5.18) with $k = k_f$ to measure α_{ren} by matching the power spectrum of `sim_1_11`. Subsequently, we determine the counterterm by taking the difference between the bare α_c (measured from the simulations with our top-down approach) and its renormalised counterpart. Figure 13 compares the time evolution of α_c (solid), α_{ctr} (dashed), and α_{ren} (dash-dotted) obtained using $\mu = 2 k_f$. Three things are worth noticing. First, the counterterm is always positive while the other two quantities are negative. Second, α_{ctr} and α_{ren} do not share the same time evolution (contrary to the guess sometimes made in the literature [e.g. 12]). In fact, the measured α_{ctr} deviates from the expected a^2 scaling, especially at late times when the EFT prediction becomes less accurate. Third, the feedback due to the UV modes on the IR modes is much smaller than the non-linear interaction between the modes with $|k|/k_f = 1$ (as evidenced by the relatively small value of α_{ctr}).

A similar renormalisation scheme has been proposed to cancel the Λ -dependence of P_{22}^{tSPT} by using the spectrum generated by the stochastic fluctuations [12, 16] under the assumption that $P_{JJ} \propto k^4$ when $k \rightarrow 0$ in analogy with Peebles' argument for uncorrelated small-scale displacements that conserve momentum [45]. However, this idea is not applicable to our configuration as P_{22}^{tSPT} is non-zero only for a discrete set of wavenumbers while P_{JJ} gets contributions at all wavenumbers. For instance, $P_{22}^{\text{tSPT}}(k_f, \Lambda, a) = 0$, while $P_{JJ}(k_f, a)$ is small but non zero. Therefore, no renormalisation is possible.

5.5 Long-term evolution of the effective speed of sound

Although our main simulations stop at $a = 6$, we run a test simulation with the same setup as `sim_1_11` which extends to $a \approx 12$. The objective of this run is to understand how the EFT coefficients behave at very late times. In Figure 14 we show the evolution of c_{tot}^2 with a . On the x -axis (top) we also show the ratio Λ/k_{NL} . For $\Lambda = 3 k_f$, as this ratio approaches unity, the $3 k_f$ mode enters the non-linear regime. We note that after dropping to a small value after shell crossing c_{tot}^2 starts increasing after $a = 6$. At $a = 6$, $\Lambda/k_{\text{NL}} \approx 0.6$, indicating that Λ and k_{NL} are not well-separated. At later times, the increase of c_{tot}^2 may therefore be attributed to the $3 k_f$ mode entering the non-linear regime. By $a = 12$, $\Lambda > k_{\text{NL}}$, meaning

¹⁰Note that $\widetilde{W}_\Lambda(k) \rightarrow 1$ in the limit $k \rightarrow 0$ for any low-pass filter.

¹¹For `sim_1_11`, $f(\Lambda) = A_1^2 \widetilde{W}_\Lambda(k_f)/2 k_f^2 + A_2^2 \widetilde{W}_\Lambda(11 k_f)/2 (11 k_f)^2$ where the A_i are the amplitudes of the linear overdensity defined in Eq. (2.1), and any Λ -dependence is contained in the smoothing kernel $\widetilde{W}_\Lambda$.




Figure 13: Time evolution of the bare coefficient α_c (solid blue line), the counterterm α_{ctr} (dashed black line), and the renormalised coefficient α_{ren} (dash-dotted red line).

that this epoch is beyond the regime of applicability of the EFTofLSS. We also note that the M&W estimator deviates from the others after $a = 6$. This is because the estimator assumes linear theory which is no longer valid in this regime.

We make one further test by adding an intermediate non-zero perturbation with wavenumber $6k_f$ and amplitude 0.25 to `sim_1_15`. We use this simulation instead of `sim_1_11` to increase the wavenumber separation between the non-zero perturbations so that another peak in c_{tot}^2 corresponding to $6k_f$ is clearly visible. The c_{tot}^2 measured from this run (which we christen `sim_1_6_15`) is shown in Fig. 15. As before, $c_{\text{tot}}^2 \propto a$ holds until the first shell crossing, then c_{tot}^2 decreases slightly. However, the $6k_f$ mode enters the non-linear regime shortly after, and we see a second peak around $a \approx 6.75$. The x -axis (top) shows that $k_{\text{NL}} < 6$ at this epoch. After this, c_{tot}^2 decreases again, until $3k_f$ enters the non-linear regime (when $6k_f/k_{\text{NL}}$ exceeds two) as in `sim_1_11` leading to another increase beyond $a \approx 10$. These results seem to confirm our earlier claim that c_{tot}^2 is sourced by modes entering the non-linear regime.

6 Summary

The EFTofLSS has been successful in extending the range of accurate perturbative predictions of summary statistics of the dark-matter density field to larger wavenumbers [e.g. 11–13, 18, 24]. With extensions to include biased tracers [e.g. 25–29] and redshift-space distortions [e.g. 46, 47], the EFTofLSS is a valuable tool in extracting cosmological parameters from data obtained from current and next-generation surveys [e.g. 48–52]. In this study, our motivation has been to strip away the complexity of the systems that the approach is usually applied to, and make controlled tests in the simplest possible setting.

The EFTofLSS makes three important assumptions: one, that the physical scales are organised in a hierarchy ordered as $k \ll \Lambda \ll k_{\text{NL}}$ (for cutoff scale Λ); two, that the effective stress tensor can be expanded in terms of long-wavelength fields in an expansion of the form given in Eq. (4.19); and three, that the EFT correction to a summary statistic (e.g., the power spectrum) at a given loop order regulates the UV sensitivity of the corresponding SPT term.

We specialise to an EdS universe with plane-parallel cosinusoidal initial conditions. Using five different top-down estimators, all assuming the expansion of the effective stress, we




Figure 14: The long-term behaviour of c_{tot}^2 for `sim_1_11`.




Figure 15: The long-term behaviour of c_{tot}^2 for `sim_1_6_15`.

measure the full time-evolution of the EFT coefficients. For all these estimators, c_{tot}^2 increases linearly with the expansion factor of the background universe until the first shell crossing, which is the scaling expected in the literature. However, after this time, we notice a clear

departure from this behaviour, as c_{tot}^2 decreases after $a \approx 2$ (Figure 9).

We compute the next-to-leading-order power spectrum in SPT and in the EFTofLSS, and compare this to the power spectrum measured in our simulations (Figure 10). We find that for the largest scale in our model universe, the SPT power matches the N -body power accurately until $a \approx 3$ for our primary simulation `sim_1_11`. At later times, we note deviations from the N -body spectrum which rises to $\approx 2.2\%$ at $a = 6$. Introducing a cutoff in the SPT loop integrals leads to an overprediction of the power at all times. At $a = 6$, this tSPT model differs from the N -body spectrum by $\approx 1.5\%$. On the other hand, the EFTofLSS power spectrum, calculated via the top-down coefficients is accurate to within $\approx 0.06\%$ at all times, indicating that the EFTofLSS accurately models the feedback from the short-wavelength modes onto the large scales. Further, the top-down and bottom-up approaches are in good agreement with each other for the bare coefficient α_c (Figure 11).

We also investigate the effect of separation (in wavenumber) between the initially non-zero perturbations on the EFTofLSS predictions (Figure 12). As expected, the tSPT power becomes less accurate (compared to N -body) with smaller separation. The EFT power shows a similar trend, with the highest accuracy at largest separation. However, for all simulations, the EFT prediction is accurate to within $\approx 0.1\%$ of that of the N -body, showing that the EFT predictions are robust to change in separation.

We explicitly measure the stochastic fluctuations of the effective stress within our setup and find that their contribution to the matter power spectrum is much smaller than the leading-order EFT correction. Further, the power spectrum of these fluctuations cannot renormalise the P_{22} term in our setup. This may be due to our initial conditions having too few degrees of freedom.

Using the measured power spectrum of `sim_1_11`, we calculate the renormalised leading EFT correction to the power spectrum (α_{ren}) and the associated counterterm (α_{ctr} ; which cancels the UV sensitivity of P_{13}^{tSPT}), and find that the two terms – unsurprisingly – do not share the same time evolution.

All of our analysis has been performed for both sharp cutoff and for Gaussian smoothings. We find differences between the estimated $\langle \tau \rangle$ from the two kernels (Figure 7), which can be understood via leakage from the non-zero high- k modes that is only a feature of Gaussian smoothing. However, the two cutoffs produce qualitatively similar results for the power spectrum of the overdensity field, with small differences for $a > 3$.

Finally, we assess the long-term behaviour of c_{tot}^2 using a test run with the same setup as `sim_1_11` but extended to $a \approx 12$. We find that c_{tot}^2 increases after $a = 6$, likely due to the $3k_f$ mode entering the non-linear regime (this is our cutoff scale, so modes bigger than $3k_f$ are excluded in the sharp cutoff case). To confirm our suspicion that c_{tot}^2 is sourced by modes entering the non-linear regime, we introduce an intermediate non-zero perturbation (with wavenumber $6k_f$) to `sim_1_15`. This introduces an additional peak in the c_{tot}^2 vs a plot when this mode becomes non-linear, further solidifying our argument.

In summary, the EFTofLSS accurately models the feedback of the small scales onto the large-scale mode in our physical system. However, the procedure of renormalising the stochastic term does not apply in our setup, perhaps due to the simplicity of our initial conditions. Tests with a more general initial setup may provide insight into these aspects of the theory.

Acknowledgments

The authors declare no conflicts of interest or external support in the preparation of this manuscript. MK is part of the International Max Planck Research School in Astronomy and Astrophysics.

References

- [1] M. Davis, G. Efstathiou, C.S. Frenk and S.D. White, *The evolution of large-scale structure in a universe dominated by cold dark matter*, *The Astrophysical Journal* **292** (1985) 371.
- [2] E. Bertschinger, *Simulations of structure formation in the universe*, *Annual Review of Astronomy and Astrophysics* **36** (1998) 599.
- [3] F. Bernardeau, S. Colombi, E. Gaztanaga and R. Scoccimarro, *Large-scale structure of the universe and cosmological perturbation theory*, *Physics reports* **367** (2002) 1 [astro-ph/0112551].
- [4] F. Bouchet, S. Colombi, E. Hivon and R. Juszkiewicz, *Perturbative lagrangian approach to gravitational instability*, *arXiv preprint* (1994) [astro-ph/9406013].
- [5] B. Jain and E. Bertschinger, *Second-order power spectrum and nonlinear evolution at high redshift*, *The Astrophysical Journal* **431** (1994) 495 [astro-ph/9311070].
- [6] R. Scoccimarro, *Cosmological perturbations: Entering the nonlinear regime*, *The Astrophysical Journal* **487** (1997) 1 [astro-ph/9612207].
- [7] M. Crocce and R. Scoccimarro, *Renormalized cosmological perturbation theory*, *Physical Review D* **73** (2006) 063519 [astro-ph/0509418].
- [8] A. Taruya, T. Nishimichi, S. Saito and T. Hiramatsu, *Nonlinear evolution of baryon acoustic oscillations from improved perturbation theory in real and redshift spaces*, *Physical Review D* **80** (2009) 123503 [arXiv:0906.0507].
- [9] J. Carlson, M. White and N. Padmanabhan, *Critical look at cosmological perturbation theory techniques*, *Physical Review D* **80** (2009) 043531 [0905.0479].
- [10] T. Matsubara, *Resumming cosmological perturbations via the lagrangian picture: One-loop results in real space and in redshift space*, *Physical Review D* **77** (2008) 063530 [0711.2521].
- [11] D. Baumann, A. Nicolis, L. Senatore and M. Zaldarriaga, *Cosmological non-linearities as an effective fluid*, *Journal of Cosmology and Astroparticle Physics* **2012** (2012) 051 [1004.2488].
- [12] J.J.M. Carrasco, M.P. Hertzberg and L. Senatore, *The effective field theory of cosmological large scale structures*, *Journal of High Energy Physics* **2012** (2012) 1 [1206.2926].
- [13] J.J.M. Carrasco, S. Foreman, D. Green and L. Senatore, *The effective field theory of large scale structures at two loops*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 057 [1310.0464].
- [14] R.A. Porto, L. Senatore and M. Zaldarriaga, *The lagrangian-space effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 022 [1311.2168].
- [15] Z. Vlah, M. White and A. Aviles, *A lagrangian effective field theory*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 014 [1506.05264].
- [16] M.P. Hertzberg, *Effective field theory of dark matter and structure formation: Semianalytical results*, *Physical Review D* **89** (2014) 043521 [1208.0839].
- [17] S.M. Carroll, S. Leichenauer and J. Pollack, *Consistent effective theory of long-wavelength cosmological perturbations*, *Physical Review D* **90** (2014) 023518 [1310.2920].

- [18] T. Baldauf, L. Mercolli, M. Mirbabayi and E. Pajer, *The bispectrum in the effective field theory of large scale structure*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 007 [1406.4135].
- [19] M.M. Ivanov, *Effective field theory for large scale structure*, *arXiv preprint* (2022) [2212.08488].
- [20] E. Pajer and M. Zaldarriaga, *On the renormalization of the effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2013** (2013) 037 [1301.7182].
- [21] J.J.M. Carrasco, S. Foreman, D. Green and L. Senatore, *The 2-loop matter power spectrum and the ir-safe integrand*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 056 [1304.4946].
- [22] L. Senatore and M. Zaldarriaga, *The ir-resummed effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 013 [1404.5954].
- [23] R.E. Angulo, S. Foreman, M. Schmittfull and L. Senatore, *The one-loop matter bispectrum in the effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 039 [1406.4143].
- [24] D. Alkhanishvili, C. Porciani, E. Sefusatti, M. Biagetti, A. Lazanu, A. Oddo et al., *The reach of next-to-leading-order perturbation theory for the matter bispectrum*, *Monthly Notices of the Royal Astronomical Society* **512** (2022) 4961 [2107.08054].
- [25] V. Assassi, D. Baumann, D. Green and M. Zaldarriaga, *Renormalized halo bias*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 056 [1402.5916].
- [26] L. Senatore, *Bias in the effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 007 [1406.7843].
- [27] R. Angulo, M. Fasiello, L. Senatore and Z. Vlah, *On the statistics of biased tracers in the effective field theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 029 [1503.08826].
- [28] M. Mirbabayi, F. Schmidt and M. Zaldarriaga, *Biased tracers and time evolution*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 030 [1412.5169].
- [29] Y. Donath and L. Senatore, *Biased tracers in redshift space in the eftoflss with exact time dependence*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 039 [2005.04805].
- [30] P. McDonald, *Clustering of dark matter tracers: Renormalizing the bias parameters*, *Physical Review D* **74** (2006) 103512 [astro-ph/0609413].
- [31] P. McDonald and A. Roy, *Clustering of dark matter tracers: generalizing bias for the coming era of precision lss*, *Journal of Cosmology and Astroparticle Physics* **2009** (2009) 020 [0902.0991].
- [32] K.F. Werner and C. Porciani, *Renormalization of linear halo bias in n-body simulations*, *Monthly Notices of the Royal Astronomical Society* **492** (2020) 1614 [1907.03774].
- [33] M. McQuinn and M. White, *Cosmological perturbation theory in 1+ 1 dimensions*, *Journal of Cosmology and Astroparticle Physics* **2016** (2016) 043 [1502.07389].
- [34] S. Colombi, *Vlasov–poisson in 1d for initially cold systems: post-collapse lagrangian perturbation theory*, *Monthly Notices of the Royal Astronomical Society* **446** (2015) 2902 [1411.4165].
- [35] A. Taruya and S. Colombi, *Post-collapse perturbation theory in 1d cosmology–beyond shell-crossing*, *Monthly Notices of the Royal Astronomical Society* **470** (2017) 4858 [1701.09088].

- [36] C. Rampf and O. Hahn, *Renormalization group and uv completion of cosmological perturbations: Gravitational collapse as a critical phenomenon*, *Physical Review D* **107** (2023) 023515 [2211.02053].
- [37] C. Rampf, U. Frisch and O. Hahn, *Unveiling the singular dynamics in the cosmic large-scale structure*, *Monthly Notices of the Royal Astronomical Society: Letters* **505** (2021) L90 [1912.00868].
- [38] Y.B. Zel'Dovich, *Gravitational instability: An approximate theory for large density perturbations.*, *Astronomy & Astrophysics* **5** (1970) 84.
- [39] S. Pueblas and R. Scoccimarro, *Generation of vorticity and velocity dispersion by orbit crossing*, *Physical Review D* **80** (2009) 043504 [0809.4606].
- [40] M. Buehlmann and O. Hahn, *Large-scale velocity dispersion and the cosmic web*, *Monthly Notices of the Royal Astronomical Society* **487** (2019) 228 [1812.07489].
- [41] T. Buchert and A. Dominguez, *Adhesive gravitational clustering*, *Astronomy & Astrophysics* **438** (2005) 443 [astro-ph/0502318].
- [42] M. Garny, D. Laxhuber and R. Scoccimarro, *Perturbation theory with dispersion and higher cumulants: Framework and linear theory*, *Physical Review D* **107** (2023) 063539 [2210.08088].
- [43] M. Garny, D. Laxhuber and R. Scoccimarro, *Perturbation theory with dispersion and higher cumulants: Nonlinear regime*, *Physical Review D* **107** (2023) 063540 [2210.08089].
- [44] T.R. Morris, *The exact renormalization group and approximate solutions*, *International Journal of Modern Physics A* **9** (1994) 2411 [hep-ph/9308265].
- [45] P.J. Peebles, *The gravitational-instability picture and the nature of the distribution of galaxies*, *Astrophysical Journal*, vol. 189, p. L51 **189** (1974) L51.
- [46] L. Senatore and M. Zaldarriaga, *Redshift space distortions in the effective field theory of large scale structures*, *arXiv preprint* (2014) [1409.1225].
- [47] M. Lewandowski, L. Senatore, F. Prada, C. Zhao and C.-H. Chuang, *Eft of large scale structures in redshift space*, *Physical Review D* **97** (2018) 063526 [1512.06831].
- [48] M.M. Ivanov, M. Simonović and M. Zaldarriaga, *Cosmological parameters from the boss galaxy power spectrum*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 042 [1909.05277].
- [49] M.M. Ivanov, M. Simonović and M. Zaldarriaga, *Cosmological parameters and neutrino masses from the final planck and full-shape boss data*, *Physical Review D* **101** (2020) 083504 [1912.08208].
- [50] T. Colas, G. D'amico, L. Senatore, P. Zhang and F. Beutler, *Efficient cosmological analysis of the sdss/boss data from the effective field theory of large-scale structure*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 001 [1909.07951].
- [51] T. Nishimichi, G. D'Amico, M.M. Ivanov, L. Senatore, M. Simonović, M. Takada et al., *Blinded challenge for precision cosmology with large-scale structure: results from effective field theory for the redshift-space galaxy power spectrum*, *Physical Review D* **102** (2020) 123541 [2003.08277].
- [52] S.-F. Chen, Z. Vlah and M. White, *A new analysis of galaxy 2-point functions in the boss survey, including full-shape information and post-reconstruction bao*, *Journal of Cosmology and Astroparticle Physics* **2022** (2022) 008 [2110.05530].




Figure 16: As in Figure 7, for $a = 4.35$ with Gaussian smoothing. Note the oscillations at the outer edges of the box for the F6P fit.

A Results with Gaussian smoothing

In this Appendix, we show the measured EFT coefficients and the power spectrum results for Gaussian smoothing. As pointed out in Section 4.4.2, the primary difference between the two smoothing kernels is that the Gaussian kernel allows modes with $k \gtrsim \Lambda$ to be non-zero. Thus, the smoothed δ and θ include contributions from these modes which introduce oscillatory features. These oscillations appear in the fits to τ , and affect the estimation of the EFT coefficients. The effect is enhanced at second order in the expansion of τ , as terms of order δ^2 and θ^2 are included. We show this for $a = 4.35$ in Figure 16. Note that the estimator fits the measured stress well in the central regions, but the effect of the oscillations is visible on the outskirts of the box. Due to this mis-estimation we omit the six-parameter fit (F6P) for the Gaussian smoothing results that follow, even though it is a good fit to τ for many snapshots of the simulation.

In Figure 17, we show the time evolution of the estimated c_{tot}^2 for four top-down estimators, analogous to Figure 9. We notice the same general behaviour as in the sharp case, viz. that c_{tot}^2 increases linearly until the first shell crossing, reaches a maximum around $a \approx 2$, then decreases with increasing a . However, there are two main differences from the sharp case: one, the estimators have a larger variation among them, and two, the value of c_{tot}^2 plateaus after an initial decrease post the first shell crossing, then shows a slight increase at very late times.

The time evolution of α_c is shown in Figure 18 for the bottom-up (dashed lines) and top-down (solid line) estimators. As in the corresponding sharp-cutoff plot (Figure 11), the top-down estimators are in rough agreement with each other and with the bottom-up estimator. However, there is larger variation among them compared to the sharp case, consistent with the observed behaviour of c_{tot}^2 .

Finally, in Figure 19, we plot the non-linear power in the $k = k_f$ mode in the different theories for the Gaussian case, in analogy with the right panel of Figure 10. The agreement between the top-down and bottom-up estimators is good until $a \approx 3$, after which most




Figure 17: As in Figure 9 but for a Gaussian filter.




Figure 18: As in Figure 11 but for a Gaussian filter.

of our estimators over-correct the tSPT-4 power. The SC estimator remains accurate to within $\approx 0.1\%$ throughout the simulation run. Overall, the results from Gaussian smoothing complement the results from sharp case in the main study. Although the match between the top-down and the bottom-up estimates is slightly worse than in the sharp case (especially at late times), the choice of the smoothing kernel does not influence the results strongly.




Figure 19: As in Figure 10 but for a Gaussian filter.

Bibliography

- Aarseth S. J., Hoyle F., 1963, *Monthly Notices of the Royal Astronomical Society*, 126, 223
- Aarseth S. J., Gott III J. R., Turner E. L., 1979, *Astrophysical Journal*, Part 1, vol. 228, Mar. 15, 1979, p. 664-683. Research supported by the Alfred P. Sloan Foundation, 228, 664
- Abdalla E., et al., 2022, *Journal of High Energy Astrophysics*, 34, 49
- Achilles R., Bonfiglioli A., 2012, *Archive for history of exact sciences*, 66, 295
- Adam R., et al., 2016, *Astronomy & Astrophysics*, 594, A10
- Adame A., et al., 2024, *The Astronomical Journal*, 167, 62
- Adamek J., Daverio D., Durrer R., Kunz M., 2013, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 88, 103527
- Adamek J., Daverio D., Durrer R., Kunz M., 2016, *Journal of Cosmology and Astroparticle Physics*, 2016, 053
- Ade P. A., et al., 2014, *Astronomy & Astrophysics*, 571, A1
- Aghanim N., et al., 2020, *Astronomy & Astrophysics*, 641, A1
- Akrami Y., et al., 2020, *Astronomy & Astrophysics*, 641, A7
- Albrecht A., Steinhardt P. J., 1982, *Physical Review Letters*, 48, 1220
- Alkhanishvili D., Porciani C., Sefusatti E., Biagetti M., Lazanu A., Oddo A., Yankelevich V., 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 4961

- Allgood B., Flores R. A., Primack J. R., Kravtsov A. V., Wechsler R. H., Faltenbacher A., Bullock J. S., 2006, *Monthly Notices of the Royal Astronomical Society*, 367, 1781
- Aluri P. K., et al., 2023, *Classical and Quantum Gravity*, 40, 094001
- Anastasiou C., Bragança D. P., Senatore L., Zheng H., 2024, *Journal of High Energy Physics*, 2024, 1
- Angulo R. E., Hahn O., 2022, *Living Reviews in Computational Astrophysics*, 8, 1
- Angulo R., Fasiello M., Senatore L., Vlah Z., 2015a, *Journal of Cosmology and Astroparticle Physics*, 2015, 029
- Angulo R. E., Foreman S., Schmittfull M., Senatore L., 2015b, *Journal of Cosmology and Astroparticle Physics*, 2015, 039
- Apicella A., Donnarumma F., Isgrò F., Prevete R., 2021, *Neural Networks*, 138, 14
- Appel A. W., 1985, *SIAM Journal on Scientific and Statistical Computing*, 6, 85
- Assassi V., Baumann D., Green D., Zaldarriaga M., 2014, *Journal of Cosmology and Astroparticle Physics*, 2014, 056
- Athanassoula E., Fady E., Lambert J., Bosma A., 2000, *Monthly Notices of the Royal Astronomical Society*, 314, 475
- Bagla J. S., 2002, *Journal of Astrophysics and Astronomy*, 23, 185
- Bagla J. S., Padmanabhan T., 1997, *Pramana*, 49, 161
- Baldauf T., Mercolli L., Mirbabayi M., Pajer E., 2015, *Journal of Cosmology and Astroparticle Physics*, 2015, 007
- Bardeen J. M., Bond J., Kaiser N., Szalay A., 1986, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 304, May 1, 1986, p. 15-61. SERC-supported research., 304, 15
- Barnes J. E., 2012, *Monthly Notices of the Royal Astronomical Society*, 425, 1104
- Barnes J., Hut P., 1986, *nature*, 324, 446
- Bartelmann M., et al., 2019, *Annalen der Physik*, 531, 1800446
- Baumann D., Nicolis A., Senatore L., Zaldarriaga M., 2012, *Journal of Cosmology and Astroparticle Physics*, 2012, 051

- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2012, *The Astrophysical Journal*, 762, 109
- Behroozi P., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 454, 3020
- Bernardeau F., Colombi S., Gaztanaga E., Scoccimarro R., 2002, *Physics reports*, 367, 1
- Bertone G., Hooper D., 2018, *Reviews of Modern Physics*, 90, 045002
- Bertone G., Hooper D., Silk J., 2005, *Physics reports*, 405, 279
- Binney J., Tremaine S., 1987, *Galactic dynamics*. Princeton university press
- Birdsall C. K., Fuss D., 1969, *Journal of Computational Physics*, 3, 494
- Birkhoff G. D., Langer R. E., 1923, *Relativity and Modern Physics*. Harvard University Press
- Block H.-D., 1962, *Reviews of Modern Physics*, 34, 123
- Blumenthal G. R., Faber S., Primack J. R., Rees M. J., 1984, *Nature*, 311, 517
- Bode P., Ostriker J. P., Xu G., 2000, *The Astrophysical Journal Supplement Series*, 128, 561
- Boggess N. W., et al., 1992, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 397, no. 2, p. 420-429., 397, 420
- Bond J., Cole S., Efstathiou G., Kaiser N., 1991, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 379, Oct. 1, 1991, p. 440-460. Research supported by NSERC, NASA, and University of California., 379, 440
- Bond J. R., Kofman L., Pogosyan D., 1996, *Nature*, 380, 603
- Borzyszkowski M., Porciani C., Romano-Díaz E., Garaldi E., 2017a, *Monthly Notices of the Royal Astronomical Society*, 469, 594
- Borzyszkowski M., Bertacca D., Porciani C., 2017b, *Monthly Notices of the Royal Astronomical Society*, 471, 3899
- Bouchet F. R., Hernquist L., 1988, *Astrophysical Journal Supplement Series* (ISSN 0067-0049), vol. 68, Dec. 1988, p. 521-538. Research supported by San Diego Supercomputer Center and University of California., 68, 521
- Bouchet F., Colombi S., Hivon E., Juszkievicz R., 1994, *arXiv preprint*

- Boveia A., Doglioni C., 2018, *Annual Review of Nuclear and Particle Science*, 68, 429
- Branco P., Torgo L., Ribeiro R. P., 2016, *ACM computing surveys (CSUR)*, 49, 1
- Branco P., Torgo L., Ribeiro R. P., 2017, in *First international workshop on learning with imbalanced domains: Theory and applications*. pp 36–50
- Bruni M., Thomas D. B., Wands D., 2014, *Physical Review D*, 89, 044010
- Buchert T., 1989, *Astronomy and Astrophysics* (ISSN 0004-6361), vol. 223, no. 1-2, Oct. 1989, p. 9-24., 223, 9
- Buchert T., Dominguez A., 2005, *Astronomy & Astrophysics*, 438, 443
- Carleo G., Cirac I., Cranmer K., Daudet L., Schuld M., Tishby N., Vogt-Maranto L., Zdeborová L., 2019, *Reviews of Modern Physics*, 91, 045002
- Carlson J., White M., Padmanabhan N., 2009, *Physical Review D*, 80, 043531
- Carrasco J. J. M., Hertzberg M. P., Senatore L., 2012, *Journal of High Energy Physics*, 2012, 1
- Carrasco J. J. M., Foreman S., Green D., Senatore L., 2014a, *Journal of Cosmology and Astroparticle Physics*, 2014, 056
- Carrasco J. J. M., Foreman S., Green D., Senatore L., 2014b, *Journal of Cosmology and Astroparticle Physics*, 2014, 057
- Catelan P., Theuns T., 1996, *Monthly Notices of the Royal Astronomical Society*, 282, 436
- Cheung C., Fitzpatrick A. L., Kaplan J., Senatore L., Creminelli P., 2008, *Journal of High Energy Physics*, 2008, 014
- Chin S., 1978, *Physics Letters B*, 78, 552
- Chisari N. E., Zaldarriaga M., 2011, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 83, 123505
- Chollet F., et al., 2015, *Keras*, <https://keras.io>
- Colombi S., 2021, *Astronomy & Astrophysics*, 647, A66
- Contreras S., Chaves-Montero J., Zennaro M., Angulo R. E., 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 3412
- Cooley J. W., Tukey J. W., 1965, *Mathematics of computation*, 19, 297

- Cooray A., Sheth R., 2002, *Physics reports*, 372, 1
- Copeland E. J., Sami M., Tsujikawa S., 2006, *International Journal of Modern Physics D*, 15, 1753
- Corasaniti P.-S., Achitouv I., 2011a, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 84, 023009
- Corasaniti P.-S., Achitouv I., 2011b, *Physical Review Letters*, 106, 241302
- Couchman H., 1991, *Astrophysical Journal, Part 2-Letters (ISSN 0004-637X)*, vol. 368, Feb. 20, 1991, p. L23-L26. Research supported by NSERC., 368, L23
- Creminelli P., Luty M. A., Nicolis A., Senatore L., 2006, *Journal of High Energy Physics*, 2006, 080
- Crocce M., Scoccimarro R., 2006a, *Physical Review D*, 73, 063519
- Crocce M., Scoccimarro R., 2006b, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 73, 063520
- Croton D. J., Gao L., White S. D., 2007, *Monthly Notices of the Royal Astronomical Society*, 374, 1303
- Cyburt R. H., Fields B. D., Olive K. A., Yeh T.-H., 2016, *Reviews of Modern Physics*, 88, 015004
- D’Amico G., Senatore L., Zhang P., Nishimichi T., 2024, *Journal of Cosmology and Astroparticle Physics*, 2024, 037
- De Lapparent V., Geller M. J., Huchra J. P., 1986, *Astrophysical Journal, Part 2-Letters to the Editor (ISSN 0004-637X)*, vol. 302, March 1, 1986, p. L1-L5. Research supported by the Smithsonian Institution., 302, L1
- Dehnen W., 2000, *The Astrophysical Journal*, 536, L39
- Dehnen W., 2001, *Monthly Notices of the Royal Astronomical Society*, 324, 273
- Dehnen W., 2002, *Journal of Computational Physics*, 179, 27
- Dehnen W., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 1226
- Dehnen W., Read J. I., 2011, *The European Physical Journal Plus*, 126, 1
- Delgado A. M., Wadekar D., Hadzhiyska B., Bose S., Hernquist L., Ho S., 2022, *Monthly Notices of the Royal Astronomical Society*, 515, 2733

- Despali G., Tormen G., Sheth R. K., 2013, *Monthly Notices of the Royal Astronomical Society*, 431, 1143
- Di Bari P., 2018, *Cosmology and the early universe*. CRC Press
- Di Valentino E., et al., 2021, *Classical and Quantum Gravity*, 38, 153001
- Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, *Astrophysical Journal*, vol. 142, p. 414-419, 142, 414
- Diemand J., Moore B., Stadel J., 2005, *Nature*, 433, 389
- Diemand J., Kuhlen M., Madau P., 2006, *The Astrophysical Journal*, 649, 1
- Donald-McCann J., Koyama K., Beutler F., 2023, *Monthly Notices of the Royal Astronomical Society*, 518, 3106
- Donath Y., Senatore L., 2020, *Journal of Cosmology and Astroparticle Physics*, 2020, 039
- Doroshkevich A., 1970, *Astrophysics*, 6, 320
- Doroshkevich A., Kotok E., Novikov I., Polyudov A., Shandarin S., Sigov Y. S., 1980, *Monthly Notices of the Royal Astronomical Society*, 192, 321
- Droste J., 1917, *Ned. Acad. Wet., SA*, 19, 197
- Dubey S. R., Singh S. K., Chaudhuri B. B., 2022, *Neurocomputing*, 503, 92
- Dyson F. J., 1949, *Physical Review*, 75, 486
- Efstathiou G., Eastwood J., 1981, *Monthly Notices of the Royal Astronomical Society*, 194, 503
- Efstathiou G., Davis M., White S., Frenk C., 1985, *Astrophysical Journal Supplement Series* (ISSN 0067-0049), vol. 57, Feb. 1985, p. 241-260. Research supported by Cambridge University, Science and Engineering Research Council of England, and NASA., 57, 241
- Ehlers J., Buchert T., 1997, *General Relativity and Gravitation*, 29, 733
- Einasto J., Kaasik A., Saar E., 1974, *Nature*, 250, 309
- Einstein A., 1915, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, pp 844-847
- Einstein A., Grossmann M., 1913, *Zeitschrift für Mathematik und Physik*, 62, 225

- Elahi P. J., Cañas R., Poulton R. J., Tobar R. J., Willis J. S., Lagos C. d. P., Power C., Robotham A. S., 2019, *Publications of the Astronomical Society of Australia*, 36, e021
- Ewald P. P., 1921, *Annalen der Physik*, 369, 253
- Fall S. M., 1978, *Monthly Notices of the Royal Astronomical Society*, 185, 165
- Fidler C., Rampf C., Tram T., Crittenden R., Koyama K., Wands D., 2015, *Physical Review D*, 92, 123517
- Fields B. D., Olive K. A., 2006, *Nuclear Physics A*, 777, 208
- Finzi A., Pirani F., 1963, *Monthly Notices of the Royal Astronomical Society*, 127, 21
- Fixsen D., 2009, *The Astrophysical Journal*, 707, 916
- Flender S. F., Schwarz D. J., 2012, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 86, 063527
- Fosalba P., Gaztanaga E., 1998, *Monthly Notices of the Royal Astronomical Society*, 301, 503
- Freedman W. L., 2021, *The Astrophysical Journal*, 919, 16
- Freedman B. A., McLerran L. D., 1977, *Physical Review D*, 16, 1169
- Freeman K. C., 1970, *Astrophysical Journal*, vol. 160, p. 811, 160, 811
- Frenk C. S., White S. D., Davis M., Efstathiou G., 1988, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 327, April 15, 1988, p. 507-525. Research supported by the Nuffield Foundation and SERC., 327, 507
- Friedmann A., 1924, *Zeitschrift für Physik*, 21, 326
- Frusciante N., Perenon L., 2020, *Physics Reports*, 857, 1
- Fry J. N., 1984, *The Astrophysical Journal*, 279, 499
- Gao L., Springel V., White S. D., 2005, *Monthly Notices of the Royal Astronomical Society: Letters*, 363, L66
- Garny M., Laxhuber D., Scoccimarro R., 2023a, *Physical Review D*, 107, 063539
- Garny M., Laxhuber D., Scoccimarro R., 2023b, *Physical Review D*, 107, 063540
- Giblin Jr J. T., Mertens J. B., Starkman G. D., 2016, *Physical Review Letters*, 116, 251301

- Gill S. P., Knebe A., Gibson B. K., 2004, *Monthly Notices of the Royal Astronomical Society*, 351, 399
- Gnedin N. Y., 2019, *The Astrophysical Journal Supplement Series*, 243, 19
- Goroff M., Grinstein B., Rey S.-J., Wise M. B., 1986, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 311, Dec. 1, 1986, p. 6-14., 311, 6
- Green S. R., Wald R. M., 2012, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 85, 063512
- Greengard L., Rokhlin V., 1987, *Journal of computational physics*, 73, 325
- Gubitosi G., Piazza F., Vernizzi F., 2013, *Journal of Cosmology and Astroparticle Physics*, 2013, 032
- Gunn J. E., Gott III J. R., 1972, *Astrophysical Journal*, vol. 176, p. 1, 176, 1
- Gunn J. E., Peterson B. A., 1965, *Astrophysical Journal*, vol. 142, p. 1633-1636, 142, 1633
- Guth A. H., 1981, *Physical Review D*, 23, 347
- Hahn O., Abel T., 2011, *Monthly Notices of the Royal Astronomical Society*, 415, 2101
- Hahn O., Paranjape A., 2016, *Physical Review D*, 94, 083511
- Hahn C., et al., 2024, *Nature Astronomy*, pp 1–11
- Hairer E., Hochbruck M., Iserles A., Lubich C., 2006, *Oberwolfach Reports*, 3, 805
- Haixiang G., Yijing L., Shang J., Mingyun G., Yuanyue H., Bing G., 2017, *Expert systems with applications*, 73, 220
- Hayli A., 1967, in *Les Nouvelles Méthodes de la Dynamique Stellaire*. p. 67
- Hayli A., 1974, in *Proceedings of the Conference on the Numerical Solution of Ordinary Differential Equations: 19, 20 October 1972, The University of Texas at Austin*. pp 304–312
- Heinz U., Jacob M., 2000, arXiv preprint nucl-th/0002042
- Heitmann K., Ricker P. M., Warren M. S., Habib S., 2005, *The Astrophysical Journal Supplement Series*, 160, 28
- Heitmann K., et al., 2008, *Computational Science & Discovery*, 1, 015003

- Hernquist L., Barnes J. E., 1990, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 349, Feb. 1, 1990, p. 562-569. Research supported by the Pittsburgh Supercomputing Center and New Jersey High Technology Grant., 349, 562
- Hertzberg M. P., 2014, *Physical Review D*, 89, 043521
- Hockney R., 1966, *The Physics of Fluids*, 9, 1826
- Hockney R. W., Eastwood J. W., 1988, *Computer simulation using particles*. CRC Press
- Hockney R. W., Goel S., Eastwood J., 1974, *Journal of Computational Physics*, 14, 148
- Holmberg E., 1941, *Astrophysical Journal*, vol. 94, p. 385, 94, 385
- Hornik K., Stinchcombe M., White H., 1989, *Neural networks*, 2, 359
- Hoyle F., 1948, *Monthly Notices of the Royal Astronomical Society*, Vol. 108, p. 372, 108, 372
- Hoyle F., Burbidge G., Narlikar J. V., 1995, *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 448, 191
- Hubble E. P., 1925, *Pop. Astr.*; Vol. 33; Page 252-255, 33
- Hubble E., 1929, *Proceedings of the national academy of sciences*, 15, 168
- Huchra J., Geller M., 1982, *Astrophysical Journal*, Part 1, vol. 257, June 15, 1982, p. 423-437. Research supported by Cambridge University., 257, 423
- Huchra J. P., et al., 2012, *The Astrophysical Journal Supplement Series*, 199, 26
- Iannuzzi F., Dolag K., 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 2846
- Ioffe S., 2015, arXiv preprint arXiv:1502.03167
- Iserles A., 2009, *A first course in the numerical analysis of differential equations*. Cambridge university press
- Isserlis L., 1918, *Biometrika*, 12, 134
- Jain B., Bertschinger E., 1994, *The Astrophysical Journal*, 431, 495
- Jain B., Bertschinger E., 1996, *Astrophysical Journal* v. 456, p. 43, 456, 43
- Jebsen J. T., 1921, *Über die allgemeinen kugelsymmetrischen Lösungen der Einstein'schen Gravitationsgleichungen im Vakuum*. Almqvist & Wiksell

- Jenkins A., Frenk C., White S. D., Colberg J. M., Cole S., Evrard A. E., Couchman H., Yoshida N., 2001, *Monthly Notices of the Royal Astronomical Society*, 321, 372
- Jernigan J., 1985, in *Symposium-International astronomical union*. pp 275–284
- Jing Y., Suto Y., 2002, *The Astrophysical Journal*, 574, 538
- Kamionkowski M., Riess A. G., 2023, *Annual Review of Nuclear and Particle Science*, 73, 153
- Kapteyn J. C., 1922, *Astrophysical Journal*, 55, 302
- Kapusta J. I., 1979, *Nuclear Physics B*, 148, 461
- Karachentsev I., 1966, *Astrophysics*, 2, 39
- Karandikar M., Porciani C., Hahn O., 2024, *Journal of Cosmology and Astroparticle Physics*, 2024, 051
- Kimontovich Y., 1967, *The Statistical Theory of Non-Equilibrium Processes in a Plasma*. Pergamon
- Klessen R. S., Glover S. C., 2023, *Annual Review of Astronomy and Astrophysics*, 61, 65
- Klypin A., Shandarin S., 1983, *Monthly Notices of the Royal Astronomical Society*, 204, 891
- Knebe A., et al., 2011, *Monthly Notices of the Royal Astronomical Society*, 415, 2293
- Knebe A., et al., 2013, *Monthly Notices of the Royal Astronomical Society*, 435, 1618
- Knollmann S. R., Knebe A., 2009, *The Astrophysical Journal Supplement Series*, 182, 608
- Kozlikin E., Lilow R., Fabis F., Bartelmann M., 2021, *Journal of Cosmology and Astroparticle Physics*, 2021, 035
- Kragh H., 2013, *Astronomy & Geophysics*, 54, 2
- Kravtsov A. V., Klypin A. A., Khokhlov A. M., 1997, *The Astrophysical Journal Supplement Series*, 111, 73
- Kuhlen M., Vogelsberger M., Angulo R., 2012, *Physics of the Dark Universe*, 1, 50

- Kurzak J., Pettitt B. M., 2006, *Molecular simulation*, 32, 775
- Lacey C., Cole S., 1994, *Monthly Notices of the Royal Astronomical Society*, 271, 676
- Lanczos C., 2012, *The variational principles of mechanics*. Courier Corporation
- Le Cun Y., 1986, in , *Disordered systems and biological organization*. Springer, pp 233–240
- Lemaître G., 1927, *Annales de la Société Scientifique de Bruxelles*, A47, p. 49-59, 47, 49
- Lemaître A. G., 1931a, *Monthly Notices of the Royal Astronomical Society*, 91, 483
- Lemaître G., 1931b, *Nature*, 127, 706
- Lemaître G., 1958, in *Ricerche Astronomiche*, Vol. 5, *Specola Vaticana*, Proceedings of a Conference at Vatican Observatory, Castel Gandolfo, May 20-28, 1957, Amsterdam: North-Holland, and New York: Interscience, 1958, edited by DJK O’Connell., p. 475. p. 475
- Lemos P., et al., 2024, *Physical Review D*, 109, 083536
- Leroy M., Garrison L., Eisenstein D., Joyce M., Maleubre S., 2021, *Monthly Notices of the Royal Astronomical Society*, 501, 5064
- Lesgourgues J., Pastor S., 2006, *Physics Reports*, 429, 307
- Lewandowski M., Senatore L., Prada F., Zhao C., Chuang C.-H., 2018, *Physical Review D*, 97, 063526
- Liboff R. L., 2003, *Kinetic theory: classical, quantum, and relativistic descriptions*. Springer Science & Business Media
- Lifshitz E. M., Khalatnikov I. M., 1963, *Advances in Physics*, 12, 185
- Linde A. D., 1982, *Physics Letters B*, 108, 389
- Livio M., 2011, *Nature*, 479, 171
- Lucie-Smith L., Peiris H. V., Pontzen A., 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 331
- Lucie-Smith L., Adhikari S., Wechsler R. H., 2022, *Monthly Notices of the Royal Astronomical Society*, 515, 2164

- Lucie-Smith L., Barreira A., Schmidt F., 2023, *Monthly Notices of the Royal Astronomical Society*, 524, 1746
- Lucie-Smith L., Peiris H. V., Pontzen A., 2024, *Physical Review Letters*, 132, 031001
- Ludlow A. D., Porciani C., 2011, *Monthly Notices of the Royal Astronomical Society*, 413, 1961
- Ludlow A. D., Schaye J., Bower R., 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 3663
- Lukić Z., Reed D., Habib S., Heitmann K., 2009, *The Astrophysical Journal*, 692, 217
- Maas A. L., Hannun A. Y., Ng A. Y., et al., 2013, in *Proc. icml*. p. 3
- Madau P., Dickinson M., 2014, *Annual Review of Astronomy and Astrophysics*, 52, 415
- Maggiore M., Riotto A., 2010a, *The Astrophysical Journal*, 711, 907
- Maggiore M., Riotto A., 2010b, *The Astrophysical Journal*, 717, 515
- Makino N., Sasaki M., Suto Y., 1992, *Physical Review D*, 46, 585
- Matarrese S., Pietroni M., 2007, *Journal of Cosmology and Astroparticle Physics*, 2007, 026
- Mather J. C., et al., 1990, *Astrophysical Journal*, Part 2-Letters (ISSN 0004-637X), vol. 354, May 10, 1990, p. L37-L40., 354, L37
- Maus M., et al., 2024, arXiv preprint arXiv:2404.07272
- McCulloch W. S., Pitts W., 1943, *The bulletin of mathematical biophysics*, 5, 115
- McDonald P., 2006, *Physical Review D*, 74, 103512
- McQuinn M., White M., 2016, *Journal of Cosmology and Astroparticle Physics*, 2016, 043
- Mehta P., Bukov M., Wang C.-H., Day A. G., Richardson C., Fisher C. K., Schwab D. J., 2019, *Physics reports*, 810, 1
- Melott A. L., Shandarin S. F., Splinter R. J., Suto Y., 1997, *The Astrophysical Journal*, 479, L79
- Michaux M., Hahn O., Rampf C., Angulo R. E., 2021, *Monthly Notices of the Royal Astronomical Society*, 500, 663

- Milillo I., Bertacca D., Bruni M., Maselli A., 2015, *Physical Review D*, 92, 023519
- Mirbabayi M., Schmidt F., Zaldarriaga M., 2015, *Journal of Cosmology and Astroparticle Physics*, 2015, 030
- Monaghan J. J., Lattanzio J. C., 1985, *Astronomy and Astrophysics* (ISSN 0004-6361), vol. 149, no. 1, Aug. 1985, p. 135-143., 149, 135
- Montero-Dorta A. D., Chaves-Montero J., Artale M. C., Favole G., 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 940
- More S., Kravtsov A. V., Dalal N., Gottlöber S., 2011, *The Astrophysical Journal Supplement Series*, 195, 4
- Müller B., Reinhardt J., Strickland M. T., 2012, *Neural networks: an introduction*. Springer Science & Business Media
- Musso M., Sheth R. K., 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 3634
- Musso M., Sheth R. K., 2023, *Monthly Notices of the Royal Astronomical Society: Letters*, 523, L4
- Musso M., Sheth R. K., 2024, arXiv preprint arXiv:2410.06289
- Nair V., Hinton G. E., 2010, in *Proceedings of the 27th international conference on machine learning (ICML-10)*. pp 807–814
- Nascimento C., Jamieson D., McQuinn M., Loverde M., 2025, *Journal of Cosmology and Astroparticle Physics*, 2025, 023
- Neyman J., Scott E., 1952, *Astrophysical Journal*, 116, 144
- Nguyen N.-M., Schmidt F., Tucci B., Reinecke M., Kostić A., 2024, *Physical Review Letters*, 133, 221006
- Nussbaumer H., Bieri L., 2009, *Discovering the expanding universe*. Cambridge University Press
- Ondaro-Mallea L., Angulo R. E., Zennaro M., Contreras S., Aricò G., 2022, *Monthly Notices of the Royal Astronomical Society*, 509, 6077
- Onions J., et al., 2012, *Monthly Notices of the Royal Astronomical Society*, 423, 1200
- Oort J. H., 1932, *Bulletin of the Astronomical Institutes of the Netherlands*, Vol. 6, p. 249, 6, 249

- Öpik E., 1922, *Astrophysical Journal*, 55, 406
- Ostriker J. P., Peebles P. J., Yahil A., 1974, *Astrophysical Journal*, vol. 193, Oct. 1, 1974, pt. 2, p. L1-L4., 193, L1
- Pajer E., Zaldarriaga M., 2013, *Journal of Cosmology and Astroparticle Physics*, 2013, 037
- Paranjape A., Hahn O., Sheth R. K., 2018, *Monthly Notices of the Royal Astronomical Society*, 476, 3631
- Peacock J., Heavens A. F., 1985, *Monthly Notices of the Royal Astronomical Society*, 217, 805
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, 12, 2825
- Peebles P., 1970, *Astronomical Journal*, Vol. 75, p. 13 (1970), 75, 13
- Peebles P. J. E., 1980, *The large-scale structure of the universe*. Princeton University Press
- Penzias A. A., Wilson R. W., 1965, *Astrophysical Journal*, 142, 419
- Perlmutter S., et al., 1999, *The Astrophysical Journal*, 517, 565
- Philcox O. H., Slepian Z., 2022, *Proceedings of the National Academy of Sciences*, 119, e2111366119
- Pietroni M., Mangano G., Saviano N., Viel M., 2012, *Journal of Cosmology and Astroparticle Physics*, 2012, 019
- Porto R. A., Senatore L., Zaldarriaga M., 2014, *Journal of Cosmology and Astroparticle Physics*, 2014, 022
- Potter D., Stadel J., Teyssier R., 2017, *Computational Astrophysics and Cosmology*, 4, 2
- Poulin V., Smith T. L., Karwal T., Kamionkowski M., 2019, *Physical review letters*, 122, 221301
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D., Springel V., Stadel J., Quinn T., 2003, *Monthly Notices of the Royal Astronomical Society*, 338, 14
- Pranav P., et al., 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 4167
- Press W. H., Davis M., 1982, *Astrophysical Journal*, Part 1, vol. 259, Aug. 15, 1982, p. 449-473., 259, 449

- Press W. H., Schechter P., 1974, *Astrophysical Journal*, Vol. 187, pp. 425-438 (1974), 187, 425
- Pritchard J. R., Loeb A., 2008, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 78, 103511
- Pueblas S., Scoccimarro R., 2009, *Physical Review D*, 80, 043504
- Quintana-Miranda E., Monaco P., Tornatore L., 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 5238
- Rampf C., Frisch U., Hahn O., 2021, *Monthly Notices of the Royal Astronomical Society: Letters*, 505, L90
- Riess A. G., et al., 1998, *The astronomical journal*, 116, 1009
- Riess A. G., Casertano S., Yuan W., Bowers J. B., Macri L., Zinn J. C., Scolnic D., 2021, *The Astrophysical Journal Letters*, 908, L6
- Riotto A., Trodden M., 1999, *Annual Review of Nuclear and Particle Science*, 49, 35
- Robertson H. P., 1935, *Astrophysical Journal*, vol. 82, p. 284, 82, 284
- Robertson B. E., 2022, *Annual Review of Astronomy and Astrophysics*, 60, 121
- Robertson B. E., Kravtsov A. V., Tinker J., Zentner A. R., 2009, *The Astrophysical Journal*, 696, 636
- Romeo A. B., Agertz O., Moore B., Stadel J., 2008, *The Astrophysical Journal*, 686, 1
- Rosenblatt F., 1958, *Psychological review*, 65, 386
- Rosenblatt F., 1961, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books
- Rubin V. C., Ford Jr W. K., 1970, *Astrophysical Journal*, vol. 159, p. 379, 159, 379
- Rumelhart D. E., Hinton G. E., Williams R. J., 1986, *nature*, 323, 533
- Saha P., Tremaine S., 1992, *Astronomical Journal* (ISSN 0004-6256), vol. 104, no. 4, p. 1633-1640., 104, 1633
- Sahni V., Coles P., 1995, *Physics Reports*, 262, 1
- Salmon J. K., 1991, *Parallel hierarchical N-body methods*. California Institute of Technology

- Samushia L., Slepian Z., Villaescusa-Navarro F., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 628
- Scaramella R., et al., 2022, *Astronomy & Astrophysics*, 662, A112
- Schaller M., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 2378
- Scherrer R. J., Bertschinger E., 1991, *Astrophysical Journal*, 381, 349
- Schneider A., et al., 2016, *Journal of Cosmology and Astroparticle Physics*, 2016, 047
- Schwarzschild K., 1916, *Sitzungsberichte der königlich preussischen Akademie der Wissenschaften*, pp 189–196
- Scoccimarro R., 1997, *The Astrophysical Journal*, 487, 1
- Sellwood J., 1987, IN: *Annual review of astronomy and astrophysics. Volume 25* (A88-13240 03-90). Palo Alto, CA, Annual Reviews, Inc., 1987, p. 151-186., 25, 151
- Senatore L., 2015, *Journal of Cosmology and Astroparticle Physics*, 2015, 007
- Shandarin S. F., Zeldovich Y. B., 1989, *Reviews of Modern Physics*, 61, 185
- Sheth R. K., Tormen G., 1999, *Monthly Notices of the Royal Astronomical Society*, 308, 119
- Sheth R. K., Tormen G., 2004, *Monthly Notices of the Royal Astronomical Society*, 350, 1385
- Sheth R. K., Mo H., Tormen G., 2001, *Monthly Notices of the Royal Astronomical Society*, 323, 1
- Simonović M., Baldauf T., Zaldarriaga M., Carrasco J. J., Kollmeier J. A., 2018, *Journal of Cosmology and Astroparticle Physics*, 2018, 030
- Slipher V. M., 1913, *Lowell Observatory Bulletin*, vol. 2, no. 8, pp. 56-57, 2, 56
- Smith S., 1936, *Astrophysical Journal*, vol. 83, p. 23, 83, 23
- Smith M. J., Geach J. E., 2023, *Royal Society Open Science*, 10, 221454
- Sousbie T., Colombi S., 2016, *Journal of Computational Physics*, 321, 644
- Spergel D. N., et al., 2003, *The Astrophysical Journal Supplement Series*, 148, 175

- Spergel D. N., et al., 2007, The astrophysical journal supplement series, 170, 377
- Springel V., 2005, Monthly notices of the royal astronomical society, 364, 1105
- Springel V., Yoshida N., White S. D., 2001, New Astronomy, 6, 79
- Springel V., White S. D., Navarro J., Jenkins A., Frenk C. S., Helmi A., Gao L., 2009, in High Performance Computing in Science and Engineering, Garching/Munich 2007: Transactions of the Third Joint HLRB and KONWIHR Status and Result Workshop, Dec. 3–4, 2007, Leibniz Supercomputing Centre, Garching/Munich, Germany. pp 93–108
- Springel V., Pakmor R., Zier O., Reinecke M., 2021, Monthly Notices of the Royal Astronomical Society, 506, 2871
- Srisawat C., et al., 2013, Monthly Notices of the Royal Astronomical Society, 436, 150
- Stark D. P., 2016, Annual Review of Astronomy and Astrophysics, 54, 761
- Starobinsky A. A., 1982, Physics Letters B, 117, 175
- Tao T., 2011, An introduction to measure theory. American Mathematical Soc.
- Taruya A., Nishimichi T., Saito S., Hiramatsu T., 2009, Physical Review D, 80, 123503
- Teyssier R., 2002, Astronomy & Astrophysics, 385, 337
- Tokdar S. T., Kass R. E., 2010, Wiley Interdisciplinary Reviews: Computational Statistics, 2, 54
- Trimble V., 1996, Publications of the Astronomical Society of the Pacific, 108, 1073
- Tucci B., Schmidt F., 2024, Journal of Cosmology and Astroparticle Physics, 2024, 063
- Udrescu S.-M., Tegmark M., 2020, Science Advances, 6, eaay2631
- Uhlemann C., Friedrich O., Villaescusa-Navarro F., Banerjee A., Codis S., 2020, Monthly Notices of the Royal Astronomical Society, 495, 4006
- Undagoitia T. M., Rauch L., 2015, Journal of Physics G: Nuclear and Particle Physics, 43, 013001
- Valogiannis G., Dvorkin C., 2022, Physical Review D, 105, 103534

- Van Albada G., 1960, Bulletin of the Astronomical Institutes of the Netherlands, Vol. 15, p. 165, 15, 165
- Verde L., Treu T., Riess A. G., 2019, Nature Astronomy, 3, 891
- Vishniac E. T., 1983, Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 274, Nov. 1, 1983, p. 152-167., 274, 152
- Vlah Z., White M., Aviles A., 2015, Journal of Cosmology and Astroparticle Physics, 2015, 014
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, Nature Reviews Physics, 2, 42
- Von Hoerner S., 1960, Z. Astrophys. 50, 184-214, 50
- Von Hoerner S., 1963, Z. Astrophys. 57, 47-82, 57
- Walker A. G., 1937, Proceedings of the London Mathematical Society, 2, 90
- Walker T. P., Steigman G., Schramm D. N., Olive K. A., Kang H.-S., 1991, Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 376, July 20, 1991, p. 51-69., 376, 51
- Wang Y., et al., 2016, Monthly Notices of the Royal Astronomical Society, 459, 1554
- Warren M. S., Abazajian K., Holz D. E., Teodoro L., 2006, The Astrophysical Journal, 646, 881
- Wechsler R. H., Zentner A. R., Bullock J. S., Kravtsov A. V., Allgood B., 2006, The Astrophysical Journal, 652, 71
- Werbos P., 1974, PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA
- White S. D., Rees M. J., 1978, Monthly Notices of the Royal Astronomical Society, 183, 341
- Wick G.-C., 1950, Physical review, 80, 268
- Wilding G., Nevenzeel K., van de Weygaert R., Vegter G., Pranav P., Jones B. J., Efstathiou K., Feldbrugge J., 2021, Monthly Notices of the Royal Astronomical Society, 507, 2968
- Wilson K. G., 1983, Reviews of Modern Physics, 55, 583
- Xu G., 1995, Astrophysical Journal Supplement v. 98, p. 355, 98, 355

- Yanagisawa C., 2014, *Frontiers in Physics*, 2, 30
- Zeldovich Y. B., 1970, *Astronomy & astrophysics*, 5, 84
- Zentner A. R., 2007, *International Journal of Modern Physics D*, 16, 763
- Zhang J., Hui L., 2006, *The Astrophysical Journal*, 641, 641
- Zhang T., Liao S., Li M., Gao L., 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 1227
- Zhao G.-B., et al., 2017, *Nature Astronomy*, 1, 627
- Zwicky F., 1933, *Helvetica Physica Acta*, Vol. 6, p. 110-127, 6, 110
- Zwicky F., 1937, *Astrophysical Journal*, 86, 217
- van den Bosch F. C., Ogiya G., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 4066

List of Figures

1.1.	The distance-velocity plot from Hubble (1929). The filled and empty circles represent – respectively – measurements from individual ‘nebulae’ (galaxies), and from combining them into groups. The solid and dashed lines are linear fits to the two sets of data, and the cross is the estimated mean of 22 nebulae for which individual distances were not found. The slope of the best-fit line is the estimated value of H_0 . <i>Credit:</i> PNAS; Hubble (1929).	7
1.2.	The temperature map of the CMB obtained from <i>Planck</i> ’s observations using the SMICA algorithm (Adam et al., 2016). Note the remarkable isotropy – the temperature deviation is of the order 10^{-4} K, compared to the CMB temperature of 2.7 K. <i>Credit:</i> ESA and the <i>Planck</i> collaboration.	9
2.1.	Comparison of excursion-set predictions of the HMF with that obtained from N -body simulations. The spherical collapse prediction (SC; dashed teal line) differs from the simulation HMF through most of the mass range, while the ellipsoidal collapse prediction (EC; dot-dashed dark blue line) is a much better match throughout, except for the highest masses.	39
2.2.	The analytical prediction for the accretion history of an example halo. <i>Left:</i> δ_L vs M trajectories at three different times. We show the SC and EC barriers. The SC barrier is constant (dashed grey line), while its EC counterpart changes with time (dot-dashed lines; colours correspond to different scale factors). The SC and EC up-crossings are marked by discs and stars respectively. <i>Right:</i> The accretion histories predicted by the three models, and the true history obtained from the simulation (box100).	40

3.1.	<i>Left</i> : Barnes-Hut quadtree for ten random points. <i>Right</i> : The corresponding tree graph.	51
3.2.	An illustration of three common activation functions via their action on uniform input in the range $[-5, 5]$	67
3.3.	An illustration of the comparison between training and validation losses. The three panels show three possible scenarios. <i>Left</i> : Underfitting – both training and validation losses remain high. <i>Center</i> : Good fit – both losses decrease and then plateau, remaining close to each other. <i>Right</i> : Overfitting – both losses decrease similarly at first, then the validation loss starts to increase.	67
4.1.	Evolution of dynamical fields in <code>sim_1_11</code> . <i>Left</i> : the overdensity field. <i>Center</i> : the mean velocity. <i>Right</i> : the velocity dispersion. Shell crossing first occurs at $a = 1.82$ (second row), followed by the formation of one-dimensional halos. In the right panel, ν indicates the fraction of particles located in multi-stream regions.	71
4.2.	Time evolution of k_{NL} for <code>sim_1_11</code> . Here, k_{NL} denotes the wavenumber corresponding to the mean (solid blue line) or median (dash-dotted red line) particle displacement. We indicate three important scales with horizontal dashed lines: $k = k_f$ is the scale of interest in our work, $k = 3 k_f$ is the coarse-graining scale we mostly use, and $k = 9 k_f$ is the smallest k_{NL} which satisfies the ordering of scales $k \ll \Lambda \ll k_{\text{NL}}$. The vertical dashed line highlights the epoch of first shell crossing.	72
4.3.	Decomposition of the effective stress in <code>sim_1_11</code> into its constituent pieces. The superscript over Ξ_ℓ indicates whether the contribution is made by particles within single-stream (s) or multi-stream (m) regions.	73
4.4.	Fits to the averaged effective stress at three different times. The solid blue line is the stress estimated from the simulations, and the dashed lines represent the fits using the F3P (black) and the F6P (red) models.	75
4.5.	Evolution of c_{tot}^2 with time for <code>sim_1_11</code> using the top-down estimators. The dash-dotted line corresponds to the scaling $c_{\text{tot}}^2 \propto a$ generally assumed in applications of the EFTofLSS. The dashed vertical lines are instances of shell crossing as detected in the simulation snapshots.	76
4.6.	Long-term evolution of c_{tot}^2 for <code>sim_1_11</code> . The vertical lines are as defined in Fig. 4.5. A second peak appears around $a = 12$, likely as a consequence of the $k = 3 k_f$ mode entering the non-linear regime.	77
4.7.	Evolution of c_{tot}^2 with time for <code>sim_1_6_15</code> using the top-down estimators. The vertical lines are as defined in Fig. 4.5. Addition of a third, intermediate mode introduces a clear peak when the mode enters the non-linear regime around $a = 6.75$	78

4.8.	Comparison of α_c between the bottom-up estimator (solid green line) and the top-down estimators (dashed lines). The dashed vertical lines are as in Figure 4.5.	79
4.9.	Evolution of the non-linear power in the $k = k_f$ mode for the system simulated by <code>sim_1_11</code> . <i>Left</i> : The measurement from <code>sim_1_11</code> (solid blue line) compared against the SPT (dotted pink line) and tSPT (dash-dotted brown line) predictions. The vertical lines indicate the time of first shell crossing (dashed teal line), and the time at which the short-wavelength mode becomes non-linear (dotted magenta line). <i>Right</i> : The N -body power compared with the five top-down EFT estimates. The bottom panels show the percentage deviation of the theoretical predictions from the measured power. Note the change in the y -axis scale, and that we omit tSPT in the bottom right panel to better visualise the EFT errors. Here, all vertical lines indicate shell crossings.	80
4.10.	As in Fig. 4.9, but for simulations with differences in separation between the long- and short-wavelength perturbations. The colours in both panels refer to the different simulations. <i>Left</i> : Comparison between the SPT and N -body power. <i>Right</i> : Comparison between tSPT, N -body, and EFT power. Note that the tSPT prediction for all simulations is identical.	82
5.1.	The architecture of the neural networks for <code>delta_box100</code> (top), and <code>evals_box100</code> (bottom). The two architectures differ only in the input layer size.	85
5.2.	Loss curves for models <code>delta_box100</code> (left) and <code>evals_box100</code> (right). They demonstrate the convergence of the two networks. . . .	86
5.3.	The halo mass function for <code>box100</code> as predicted by the two primary models (dotted lines), and the true mass function of the test set (solid red line). We also show the EC prediction (dot-dashed line) for an excursion-set theory comparison.	86
5.4.	Halo masses predicted by the models vs the masses measured in the simulation. We bin the halos to get a sense of the underlying correlation. The Pearson correlation coefficients are indicated in the top-left corner of each subplot. <i>Top</i> : The neural network predictions. <i>Bottom</i> : Excursion-set theory predictions. The former have a smaller scatter, and higher r -values.	87
5.5.	The distribution of test-set halos in formation time. We compare the predictions from the two neural network models (dotted lines) against the true distribution (solid red line). The predictions peak at the median value of the true distribution, but do not capture the underlying patterns either in the low- or high- a_{50} limit.	89

- 5.6. Comparison of the predicted and true distributions (solid red line) of test-set halos in formation time (a_{50}). *Left*: The distributions predicted by the two neural network models after oversampling. *Right*: The distributions predicted by excursion-set theory using spherical and ellipsoidal collapse. 90
- 5.7. Comparison of the predicted accretion history vs that obtained from **box100** (solid red line in both panels). We take the median of all halos in the test set to condense the information. *Left*: The two machine learning models (dotted lines). *Right*: The excursion-set predictions from spherical collapse (dashed teal line) and the two ellipsoidal collapse models (dot-dashed lines). 91
- 5.8. Predicted vs true formation times (a_{50}), with halos collected in hexagonal bins. We also indicate the corresponding Pearson correlation coefficient (r) in the top-left corner of each subplot. *Top*: The two neural network predictions – **delta_box_100** (left) and **delta_box_100** (right) – have relatively high r -values. *Bottom*: SC (left) and EC (right) are poorer predictors of a_{50} , as expected from the distributions in Fig. 5.6. 93
- 5.9. Architecture of the neural network for **alpha_box100**. We add a second dense layer with 64 neurons to the existing architecture of Fig. 5.1. 94
- 5.10. Prediction of halo masses using the **alpha_box100** model. *Left*: The predicted mass function (dotted indigo line) vs that measured from **box100** (solid red line). We also show the **delta_box100** (dotted light blue line) HMF as a comparison. *Right*: Predicted vs true halo masses with hexagonal binning. The Pearson correlation is indicated in the top-left corner. 94
- 5.11. Prediction of accretion history using the **alpha_box100** model. *Left*: The distribution of halos in a_{50} from the predicted and true accretion histories. We also show the **delta_box100** (dotted light blue line) distribution as a comparison. *Right*: Predicted vs true formation times, with hexagonal binning for the halos. 95
- 5.12. As in Fig. 5.7, but for models based in the **box300** simulation. *Left*: Predictions from the two neural networks. *Right*: The excursion-set theory predictions. 96
- 5.13. Predicted vs true halo masses, as in Fig. 5.4 but for models based in **box300**. *Top*: The two neural network predictions – **delta_box300** (left) and **evals_box300** (right). *Bottom*: The excursion-set theory predictions – SC (left) and EC (right). 97

5.14.	Predicted vs true formation times, as in Fig. 5.8, but for models based in <code>box300</code> . <i>Top</i> : The two neural network predictions – <code>delta_box_300</code> (left) and <code>delta_box_300</code> (right). <i>Bottom</i> : The excursion-set theory predictions – SC (left) and EC (right).	98
B.1.	The median accretion history predicted by the two machine learning models (dotted lines), compared against the true median accretion history of the test set (solid red line). We also show the excursion-set predictions from both spherical collapse (dashed teal line) and two ellipsoidal collapse models (dot-dashed lines). <i>Left</i> : Low- a_{50} tail. <i>Centre</i> : Majority subset containing halos with $0.4 \leq a_{50} \leq 0.7$. <i>Right</i> : High- a_{50} tail. We also indicate the fraction of halos in each subset (f).112	112
C.1.	Predicted vs true halo mass (left) and formation time (right) for Bo17 using <code>box100</code> . We group the halos in hexagonal bins to better illustrate the underlying correlation.	113
C.2.	Same as Fig. C.1, but for models based in <code>box300</code>	114

List of Tables

3.1. Simulation parameters for box100 and box300 . The only difference between the two is the box size, which leads to a different mass resolution.	61
4.1. Details of the simulations used to test the EFTofLSS in this chapter. For each simulation, we have eight runs, such that $\varphi_1 = \pi$, and φ_2 is uniformly drawn from $[0, 2\pi)$	70
5.1. Comparison of model performance for box100 in predicting the accretion histories. We use the coefficient of determination (R^2) as our metric. The highest value is bold for emphasis.	92
6.1. Comparison of model performance for box100 . The rows contain our metrics, and the columns represent the six models. The highlighted values are the highest values in each row.	103
6.2. As in Table 6.1, but for models based in box300 . The excursion-set models perform better than their box100 counterparts.	103