

RECEIVED: November 8, 2023

REVISED: April 7, 2024

ACCEPTED: April 27, 2024

PUBLISHED: June 26, 2024

DISCO-DJ I: a differentiable Einstein-Boltzmann solver for cosmology

Oliver Hahn ^{a,b}, Florian List ^a and Natalia Porqueres ^c

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

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

^cDepartment of Physics, University of Oxford,
Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, U.K.

E-mail: oliver.hahn@univie.ac.at, florian.list@univie.ac.at,
natalia.porqueres@physics.ox.ac.uk

ABSTRACT: We present the Einstein-Boltzmann module of the DISCO-DJ (**D**ifferentiable **S**imulations for **C**osmology — **D**one with **J**AX) software package. This module implements a fully differentiable solver for the linearised cosmological Einstein-Boltzmann equations in the JAX framework, and allows computing Jacobian matrices of all solver output with respect to all input parameters using automatic differentiation. This implies that along with the solution for a given set of parameters, the tangent hyperplane in parameter space is known as well, which is a key ingredient for cosmological inference and forecasting problems as well as for many other applications. We discuss our implementation and demonstrate that our solver agrees at the per-mille level with the existing non-differentiable solvers CAMB and CLASS, including massive neutrinos and a dark energy fluid with parameterised equation of state. We illustrate the dependence of various summary statistics in large-scale structure cosmology on model parameters using the differentiable solver, and finally demonstrate how it can be easily used for Fisher forecasting, with a forecast for Euclid as an example. Since the implementation is significantly shorter and more modular than existing solvers, we believe it will be more straightforward to extend our solver to include additional physics, such as additional dark energy and dark matter models, modified gravity, or other non-standard physics in the future.

KEYWORDS: cosmological parameters from LSS, cosmological perturbation theory, baryon acoustic oscillations, power spectrum

ARXIV EPRINT: [2311.03291](https://arxiv.org/abs/2311.03291)

Contents

1	Introduction	1
2	Implementation	3
2.1	Overview	3
2.2	Automatic differentiation	4
2.3	Stiff ODEs and implicit solvers	5
2.4	Approximation methods	6
2.5	Thermal history/recombination solver	8
3	Validation and performance	9
3.1	Choice of fiducial cosmological model	9
3.2	Comparison with CAMB and CLASS	9
3.3	Gradient of the power spectrum w.r.t. cosmological parameters	12
4	Applications	14
4.1	Redshift-space correlators for large-scale structure studies	14
4.1.1	Linear redshift space distortions	14
4.1.2	Correlation function multipoles	16
4.1.3	Two-dimensional power spectrum	17
4.2	Fisher-forecasting with differentiable solvers	19
5	Summary and conclusions	22
A	Governing equations	24
B	Comparison with Class and Camb at early times	27

1 Introduction

Describing the evolution of perturbations of Einstein’s field equations in a cosmological setting is an almost century-old quest and the foundation of inhomogeneous cosmology as much as that of all of extragalactic astrophysics. These perturbations, leading to deviations from the homogeneous Friedmann-Lemaître-Robertson-Walker (FLRW) spacetimes and the growth of structure, are described by the Einstein-Boltzmann equations, which are a set of coupled partial differential equations that describe the evolution of the energy density of the various components of the Universe along with the perturbations to the FLRW metric. First derivations of the equations date back many decades [1, 2] and subsequent revisions provide the foundation of all of modern inhomogeneous cosmology [3–5].

The Einstein-Boltzmann equations are a cornerstone of modern cosmology, and are the basis for the computation of the cosmic microwave background (CMB) anisotropies, the matter power spectrum, and other observables. The matter power spectrum, or its

constituents, the individual transfer functions of the various matter components, is the starting point for all predictions of large-scale structure observables, such as the galaxy n -point spectra, weak lensing spectra, the halo mass function (in analytic models), and in particular also the initial conditions for all cosmological simulations. Building on the *linear* Einstein-Boltzmann equations, the *non-linear* evolution of structure is then typically described in the non-relativistic collisionless limit by the Vlasov-Poisson system of equations, which is the basis for all non-linear perturbation theory and all N -body simulations (e.g. [4, 6, 7]).

Several famous software packages have been developed that integrate the linearised Einstein-Boltzmann equations. The most commonly used (historically) include the COSMICS/LINGER code [5], followed by a version specifically optimised for CMB predictions, the CMBFAST code [8], which included the line-of-sight approximation which allowed a dramatically faster computation of CMB spectra. The increasing requirements of the WMAP [9] and Planck [10] CMB missions for efficient codes have led to the highly optimised and accurate CAMB¹ [11] and CLASS² [12, 13] codes. A recent addition is PyCOSMO³ [14] which relies on automatic translation of symbolic expressions to compiled code. Both CAMB and CLASS feature a large number of optimizations (e.g. approximate asymptotic forms of the equations valid in various regimes, robust high-order implicit integrators), are highly optimised, and are written in compiled functional languages — FORTRAN (CAMB) and C (CLASS) — but both now have convenient PYTHON interfaces.

Both packages share a common minimal set of implemented equations (see [5, 15]), which model the evolution of cold dark matter (CDM), baryons, photons, massless and massive neutrinos, a dark energy fluid, as well as metric perturbations. While already the first codes included these energy components, CAMB and CLASS have by now been extended to include a plethora of additional models, including various dark energy and modified gravity models [16–18], and other variations such as scalar field dark matter [19, 20], early dark energy [21], and many more.

In this article, we present the Einstein-Boltzmann (EB) module of a new software package, DISCO-DJ (**D**ifferentiable **S**imulations for **CO**smology — **D**one with **JAX**), that implements a fully differentiable Einstein-Boltzmann solver in the JAX framework [22]. This solver is based on the DIFFRAX PYTHON package [23], which (among others) implements a high-order implicit Runge-Kutta solver for stiff ordinary differential equations (ODEs). All code (including the DIFFRAX package) is built on top of the JAX framework, which provides automatic differentiation of PYTHON code, and thus allows computing the Jacobian matrix of the Einstein-Boltzmann system as well as functions of its results with respect to all input parameters. This is a key ingredient for many applications relying on gradients, including cosmological inference and the computation of the Fisher matrix computation for forecasting or optimal data compression.

In addition to differentiability, the JAX framework is also optimised for execution on graphics processing units (GPUs), which allows for a significant speed-up of the computation of the Einstein-Boltzmann system, in particular when batching over multiple cosmological models, which is an embarrassingly parallel calculation.

¹<https://camb.info>.

²https://lesgourg.github.io/class_public/class.html.

³<https://cosmology.ethz.ch/research/software-lab/PyCosmo.html>.

The current version of the EB module of DISCO-DJ (referred to as DISCO-EB in what follows) has a very lean and modular code base and is thus easy to extend and modify beyond its current minimal implementation from a programming point of view. In particular, we hope that in the future additional physics modules, such as dark energy models, modified gravity, or other non-standard physics will be added by the community. Eventually, we expect that also other deliverables beyond what is shown in this paper will be added, such as e.g. the CMB anisotropies or the lensing potential. Further, the EB module will be connected to a fast N -body solver, which we are also developing as part of the DISCO-DJ package and which allows the inclusion of non-linear effects in the computation of the matter power spectrum and other key statistics and predictions for large-scale structure cosmology. We will present this in a future publication.

The structure of this article is as follows: in section 2 we discuss the implementation of the Einstein-Boltzmann solver in DISCO-DJ, giving attention to differentiable solvers, implicit integrators for stiff ODEs, approximation schemes, and the thermal history solver. In section 3 we validate the accuracy of the solver against the commonly used CAMB and CLASS codes. We then discuss various applications of a differentiable Einstein-Boltzmann solver, including the computation of the dependence of various summary statistics on cosmological parameters, and the computation of the Fisher information matrix in section 4. Finally, we summarise and conclude in section 5. An overview of the equations implemented in the Einstein-Boltzmann module of DISCO-DJ is given in appendix A.

2 Implementation

In this section, we discuss aspects of the implementation of the linear Einstein-Boltzmann module of DISCO-DJ.

2.1 Overview

The Einstein-Boltzmann equations are a set of non-linear coupled partial differential equations that describe the evolution of the energy density of the various components of the Universe along with the perturbations to the FLRW metric. A fully self-consistent non-linear solution all the way to late times and small scales is however so far computationally untractable (but see [24] for a second-order EB solver). Instead, in cosmology, one typically splits the evolution into a linear and a non-linear part. The linear part is described by the linearised Einstein-Boltzmann equations. The non-linear part is typically described in the non-relativistic collisionless limit by the Vlasov-Poisson system of equations, which is solved with N -body methods (e.g. [7]), using input from the linear Einstein-Boltzmann solver.

In this work, we focus on the linear Einstein-Boltzmann equations including CDM, baryons, photons, massive and massless neutrinos, and clustering dark energy. The governing equations are stated in full in appendix A for reference. The same equations are also implemented in the popular CAMB and CLASS solvers. Those solvers also include additional physical models (e.g. early dark energy, dark radiation, etc.) that we do not include (yet).

The new aspect of our work is that we implement the Einstein-Boltzmann equations in the JAX framework, which allows us to compute the Jacobian matrix of the Einstein-Boltzmann

system (which is crucial input for implicit ODE solvers) as well as of any function of its results with respect to all parameters. For this reason, in this article, we first discuss the aspect of automatic differentiation in section 2.2 below. The linearised Einstein-Boltzmann equations are so-called ‘stiff’ coupled ODEs, which are most efficiently solved by implicit integrators, an aspect we discuss in section 2.3. Solvers such as CAMB and CLASS employ a number of numerical approximation schemes in order to reduce the large number of equations needed to be solved and optimise the time-to-result for a given precision, which is critical for the use in data analysis. We discuss these aspects in section 2.4. Finally, we also discuss our choice of thermal history solver and its performance in section 2.5.

2.2 Automatic differentiation

The autodiff paradigm [25] has seen rapidly increasing interest recently due to the machine learning revolution. Traditionally, the computer-aided computation of derivatives could be subdivided into *symbolic* and *numerical* differentiation methods. Symbolic differentiation involves algorithmically determining the derivative of a function as a new symbolic expression by applying differentiation rules such as the chain rule, the product rule, etc. This approach is implemented in software packages such as the popular MATHEMATICA program [26] and the PYTHON module SYMPY [27]. While these programs are often able to find surprisingly simple expressions for derivatives that would be hard to compute by hand, more complex functions may result in lengthy expressions that are far from computationally optimal when being evaluated with concrete numerical values. On the other hand, numerical differentiation relies on finite difference approximations of derivatives. This introduces a hyperparameter, namely the step size, upon which the results can depend sensitively. In particular, numerical approximations of higher derivatives are oftentimes simply too noisy to be useful in practice.

More recently, the autodiff paradigm (also known as *automatic* differentiation) has become a popular third contender among computational differentiation methods and is now the de-facto standard in the field of machine learning. Automatic differentiation exploits the fact that programs perform the evaluation of any mathematical function of some inputs as a nested sequence of simple functions such as summation, multiplication, exponentiation, etc. The derivative of each of these building blocks is known, so when the program sequentially propagates the inputs through these blocks, the derivative information can be passed on alongside the computed values simply by using the chain rule. This particular flavour of autodiff is known as forward-mode differentiation. Alternatively, if the chain rule is instead performed starting at the outputs of the function and propagating backwards to the inputs, one obtains reverse-mode differentiation. Which of the two is computationally cheaper mainly depends on the dimensionality of the inputs and outputs (backward-mode for many inputs \mapsto few outputs, e.g. when computing a scalar loss function w.r.t. the millions of parameters of a neural network, forward-mode for few inputs \mapsto many outputs, e.g. when differentiating a field quantity w.r.t. a cosmological parameter). Keeping track of the derivatives with autodiff only entails a constant computational overhead for each of the function’s building blocks.

In our case, the quantity of interest, namely the matter power spectrum P , is obtained as the solution of a system of differential equations, specifically the cosmological Einstein-Boltzmann equations. Thus, the computation of gradients such as $\partial P / \partial \Omega_m$ requires us to

differentiate through the solution of the differential equations. This can be done in two different ways: (1) differentiate iteratively through the time steps performed by the numerical solver (“discretise-then-optimize” approach) or (2) set up a continuous “adjoint equation” for $\partial P/\partial\Omega_m$ that then must be solved backwards in time, also with a numerical method (“optimize-then-discretise” approach).⁴ Generally, the optimize-then-discretise approach yields less accurate derivatives because the numerical discretisation of the backwards-in-time equations required for the computation of the derivatives is not synchronised with that for the forward-in-time equations for the quantities themselves [28]; however, it is more memory-efficient than its discretise-then-optimize counterpart.

In this work, we rely on Google’s array-based autodiff PYTHON library JAX [22]. JAX conveniently allows writing differentiable programs in high-level PYTHON that are then compiled just-in-time (JIT) to execute efficiently on graphics or tensor processing units (GPUs, TPUs), making use of the fast XLA (Accelerated Linear Algebra) compiler. Syntactically, much of JAX follows the syntax of NUMPY, making it easy for users with PYTHON experience to switch from plain NUMPY (which runs on central processing units, i.e. CPUs, and is neither compiled nor differentiable) to JAX. It supports forward- and reverse-mode differentiation, the computation of higher-order derivatives such as Hessian matrices, holomorphic differentiation, and the vectorised computation of derivatives, among other useful features. The JAX ecosystem is steadily expanding, and there exist JAX-based packages for machine learning (e.g. [29]), for physics simulations (e.g. [30]), probabilistic programming (e.g. [31, 32]), and many other applications, which can be integrated seamlessly with custom JAX code. Importantly, we use the JAX-based DIFFRAX library for solving differential equations [23], which supports different flavours of “discretise-then-optimize” and “optimize-then-discretise” derivatives. The current version of DISCO-EB supports both forward (via DIFFRAX’s `DirectAdjoint`) and backward (via DIFFRAX’s `BacksolveAdjoint`) differentiation. Other differentiable frameworks exist; notably e.g. BOLT⁵ implements a forward-differentiable Einstein-Boltzmann solver in JULIA (but is unpublished at the time of writing).

In the past few years, JAX has garnered considerable interest in the cosmology community. For example, ref. [33] presented a differentiable cosmology library using JAX, ref. [34] introduced a JAX-based inference framework for cosmology, and the N -body code by ref. [35] is fully differentiable thanks to its implementation in JAX. It has also been used in the context of estimating weak lensing shear [36] and for modelling the evolution of dark matter haloes [37].

2.3 Stiff ODEs and implicit solvers

The Einstein-Boltzmann equations, cf. appendix A, are a stiff system of ODEs due to the multiple time scales involved in the various physical processes, which makes them ill-suited for explicit time integration schemes [38]. Various approximations can be employed (cf. [13], and section 2.4) to reduce the stiffness of the system (see also [39]), but these approximations are generally not valid for the entire domain of integration. State-of-the-art solvers therefore use a combined approach of approximation schemes and a powerful high-order implicit time

⁴Instead of the power spectrum itself, one could similarly differentiate a *loss function* that gauges the match between the computed power spectrum P and an observation P_{obs} w.r.t. the cosmological parameters, which is a typical scenario in the context of parameter inference.

⁵<https://github.com/xzackli/Bolt.jl>.

integration scheme [13, 40]. A main draw-back of implicit solvers is that the Jacobian matrix of the system needs to be computed and inverted at each time step, which is computationally expensive, particularly so for large systems such as the cosmological Einstein-Boltzmann equations.

A main advantage of the autodiff paradigm is that the Jacobian matrix of the system can be computed automatically without specifying it explicitly or having to recourse to finite-difference approximations. For DISCO-EB, we adopt the DIFFRAX package [23] for this purpose, which comes with various high-order implicit and explicit integration and is also built on top of the JAX framework. Specifically, we adopt the highest order implicit solver implemented in DIFFRAX: the 7-stage KVAERNO5 solver [41], a singly diagonally implicit Runge-Kutta (ESDIRK) method of order 5/4.

Due to the large number of equations, it is in principle critical to exploit the sparse structure of the Jacobian matrix for efficient implicit solvers — such as done by CLASS which uses sparse storage for the Jacobian. However, sparse matrix operations in JAX are still in an experimental stage and thus, DIFFRAX does not (yet) support sparse matrices, which is the current performance bottleneck of DISCO-EB.

After linearisation of the Einstein-Boltzmann system [5], the equations for individual wave numbers k decouple in Fourier space and can thus be solved individually. The resulting Jacobian for multiple wave numbers thus has a block-diagonal structure. Since such sparse Jacobian matrices can currently not be handled, an optimised joint integration of multiple wave numbers, which leads to a block diagonal Jacobian (since the different modes do not couple), is currently not possible. Storing the full Jacobian for multiple modes is computationally prohibitive. While we are, therefore, currently restricted to integrating mode-by-mode in an embarrassingly parallel way, there is plausibly room for future performance improvements along such lines. In applications where the Einstein-Boltzmann system needs to be solved for multiple sets of cosmological parameters, e.g. in the context of inference tasks, it is straightforward to batch the computation over those sets in order to increase the computational speed, harnessing the highly parallel architecture of GPUs. Currently, we use however the JAX `vmap` function to vectorise the computation of multiple modes on the GPU.

2.4 Approximation methods

In order to speed up the integration of the Einstein-Boltzmann system, for CLASS and CAMB, it has been critical to employ various approximation schemes [13, 40] in order to reduce the number of degrees of freedom of the high-dimensional system and to reduce the stiffness of the system. The situation is somewhat different for GPU-JAX-based solvers. While the Jacobian matrix of the system can be computed automatically (which is, in fact, more efficient than a finite difference calculation), the computation of the Jacobian matrix is still computationally expensive, and it is still desirable to reduce the number of degrees of freedom of the system. However, adaptively switching to a reduced set of variables based on a predefined criterion introduces additional algorithmic complexity, which is not desirable in the context of GPU-based solvers. Specifically, we found that neither of the commonly used approximation schemes yielded a $\mathcal{O}(1)$ performance improvement, and, while we implemented the tight coupling approximation [5, 13] (valid when the baryons are fully ionised and tightly coupled

to the photons), and all the approximations used in CLASS [13] (notably the ultra-relativistic fluid approximation, applying to massless neutrinos at late times, as well as the radiation streaming approximation, applying to massless neutrinos and photons at even later times) in DISCO-EB, we do not currently use any of them.

A main performance bottleneck lies in the computation of the Jacobian matrix, which scales quadratically with the number of equations (see appendix A). The number of equations scales with the number ℓ_{max} of the multipole where the photon and neutrino distribution functions are truncated, and the number of massive neutrino mass bins. Specifically we have

$$8 + 2 \times \ell_{\gamma, \text{max}} + \ell_{\nu \text{UR}, \text{max}} + n_q \times \ell_{\nu \text{massive}, \text{max}} \quad (2.1)$$

equations. The ‘8’ entails the equations for metric+CDM+baryon+dark energy perturbations, $\ell_{\gamma, \text{max}}$ is the maximum multipole moment of the photon distribution function, $\ell_{\nu \text{UR}, \text{max}}$ is the maximum multipole moment of the massless neutrino distribution function, n_q is the number of neutrino momentum bins, and $\ell_{\nu \text{massive}, \text{max}}$ is the maximum multipole moment of the massive neutrino distribution function.

A key optimisation, therefore, lies in a clever choice of neutrino momentum bins so that n_q can be chosen as small as possible. Specifically, [42] found that a 3-5 bin Gauss-Laguerre inspired integral approximation used by CAMB (see appendix A of [42]) works very well in this context. The neutrino momentum bins are chosen so that the following integral is well approximated by a sum over the bins

$$\frac{1}{4} \int_0^\infty dq \frac{q^4 e^q}{(1 + e^q)^2} v^w \nu_\ell \approx \sum_i K_i v^w \nu_\ell, \quad (2.2)$$

where q is the neutrino momentum, v is the (relativistic) neutrino velocity, ν_ℓ is a rescaled function of the ℓ -th multipole moment of the neutrino distribution function, $w = -1, 0, 1$ (depending on which moment of the distribution function is to be computed, cf. eqs. (55) of [5]), and K_i are the kernel weights approximating the integral. The authors of [42] propose the following very accurate binning schemes:

3-point:	$q = (0.913201, 3.37517, 7.79184),$ $K = (0.0687359, 3.31435, 2.29911),$
4-point:	$q = (0.7, 2.62814, 5.90428, 12),$ $K = (0.0200251, 1.84539, 3.52736, 0.289427)$
5-point:	$q = (0.583165, 2.0, 4.0, 7.26582, 13.0),$ $K = (0.0081201, 0.689407, 2.8063, 2.05156, 0.12681)$

The truncation of the neutrino Boltzmann hierarchy $\ell_{\nu, \text{max}} = \text{lnumax}$ then determines the number of equations that need to be solved as the product of lnumax and the number of neutrino momentum bins. When neutrinos are sufficiently non-relativistic, lnumax can, in principle, be reduced to a very small number (2 or 3, cf. [42]).

For our first proof-of-concept version of DISCO-EB, we decided to leave out all such further approximations, but we expect that a possible factor of 2–3 in speed could be gained by a careful evaluation of the various approximation schemes on integration of the ODEs




Figure 1. Performance of the simplified recombination solver ‘RECFASTmini’ in DISCO-DJ in comparison with the more accurate RECFAST and HYREC solvers used in CLASS, and the original MB95 model [5] (also available in DISCO-DJ). We show the ionization fraction x_{e-} as a function of the scale factor a for the baseline cosmology computed using the simplified model based on the original RECFAST (blue), the much older MB95 model (orange dashed), as well the most recent RECFAST (red dashed) and HYREC (green) modules in CLASS. The main difference is due to details in the treatment of the combined hydrogen-helium recombination. None of these differences (except MB95) have a measurable effect late-time on matter components.

on GPUs (specifically optimizing how to best change the number of equations integrating directly in the integrator). We leave this for future work.

2.5 Thermal history/recombination solver

Currently, we adopt a simplified thermal history solver in DISCO-DJ, that is based on the original RECFAST model [43]. The original implementation of the thermal history in COSMICS (section 5.8 of [5], which is based on [44] with addition Helium) is not accurate enough to obtain sub-per-cent-level agreement with CAMB or CLASS. All state-of-the-art Einstein-Boltzmann solvers employ at least RECFAST [43], or the newer COSMOREC [45] and HYREC [46] (which has been further developed into HYREC2 [47]) solvers. We found however that, since we are predominantly interested in the matter power spectrum and late times (rather than a CMB spectrum), adopting the original RECFAST model is accurate enough for our purposes. While the thermal history solver is thus not the main source of error in the late time matter power spectrum, we are planning to eventually add a more accurate recombination solver in DISCO-DJ. The current version implements the original RECFAST model, that we dub ‘RECFASTmini’, as well as the much older MB95 model [5].

In figure 1, we compare the ionization history of the baseline cosmology computed using the simplified models in DISCO-EB with the RECFAST and HYREC solvers interfaced in CLASS. The main difference is due to the different treatment of the helium recombination, with a longer phase of He^+ recombination in RECFAST. As we will show below, for late times $z \gtrsim 1000$, the simplified ‘RECFASTmini’ model is sufficient to reach agreement of late time spectra with CLASS or CAMB to below one per cent.

Parameter	Value	Meaning
Ω_k	0.0	curvature density parameter (constraint)
Ω_b	0.0488911	baryon density parameter
Ω_m	0.3099	matter density parameter
h	0.67742	Hubble parameter
n_s	0.96822	primordial spectral index
A_s	2.1064×10^{-9}	amplitude at $k_p = 0.05 \text{ Mpc}^{-1}$
N_{eff}	3.046	effective number of ultrarelativistic species
m_ν	0.06 eV	mass of neutrino
w_0	-0.99	dark energy equation of state parameter ⁶
w_a	0.0	dark energy equation of state parameter
c_a^2	1.0	dark energy sound speed
Y_{He}	0.248	primordial Helium fraction
T_{CMB}	2.7255 K	CMB temperature today

Table 1. Baseline flat PLANCK2018 EE+BA0+SN cosmological parameters used for all tests unless otherwise indicated. The dark energy density Ω_{DE} is chosen so that all densities add up to critical and thus $\Omega_k = 0$ serves as a constraint, not a parameter.

3 Validation and performance

We next validate the performance of our solver against the CAMB and CLASS solvers. First, we compare the precision of our solver for a given set of cosmological input parameters. Second, we demonstrate that our solver is able to compute the Jacobian matrix of the system with respect to all input parameters, and validate the correctness of the Jacobian matrix by comparing it to finite-difference approximations.

3.1 Choice of fiducial cosmological model

For all comparisons, unless otherwise indicated, we adopt a flat (extended) Λ CDM cosmology, with the cosmological parameter set shown in table 1, which is consistent with the PLANCK 2018 EE+BA0+SN constraints [48]. Specifically, we consider the standard extension with inclusion of one massive neutrino, along with a dark energy fluid with parameterised (Chevallier-Polarski-Linder) equation of state $w = p/\rho = w_0 + w_a(1-a)$. We adopt $w_0 = -0.99$ instead of -1 since derivatives w.r.t. w_0 are defined only in the one-sided limit at $w = -1$, a complication we thus avoid (although defining one-sided derivatives is, in principle, also supported by JAX).

3.2 Comparison with CAMB and CLASS

In this subsection, we validate the performance of the DISCO-DJ Einstein-Boltzmann solver against the well-established and commonly used CAMB⁷ [11], specifically version 1.5.0, and CLASS⁸ [12], version 3.2.0, software packages. We use both codes via their PYTHON interfaces.

⁶Note that we modified this from -1 as solutions are only one-sidedly differentiable at $w_0 = -1$.

⁷Available from <https://camb.info>.

⁸Available from https://lesgourg.github.io/class_public/class.html.

```

AccuracyBoost = 3
lAccuracyBoost = 3
DoLateRadTruncation = False
MassiveNuMethod= 'Nu_int'
Transfer.high_precision = True
Transfer.accurate_massive_neutrino_transfers = True
Reion.Reionization = False

```

Table 2. Parameters used for the reference high-precision CAMB runs.

```

l_max_g = 64
l_max_pol_g = 64
l_max_ur = 64
l_max_ncdm = 64
use_ppf = no
radiation_streaming_approximation = 3
ncdm_fluid_approximation = 3
ur_fluid_approximation = 3
reio_parametrization = reio_none

```

Table 3. Parameters used for the reference high-precision CLASS runs. Note that for DISCO-EB, we used only $\ell_{\max} = 32$.

Since DISCO-DJ is aimed at large-scale structure studies, we are mainly interested in the post-recombination regime. We use high accuracy settings for all three codes, listed in table 2 and table 3, for CAMB and CLASS respectively, to boost the accuracy of the integration results. For DISCO-EB, we adopted a maximum ℓ of the hierarchy truncation of 32, as well as the 5-point neutrino momentum integration. Note that we found that CLASS enables by default some optimizations which can lead to relatively large differences, notably we had to disable some optimizations for massive neutrinos and the dark energy fluid.

In all tests, we employ the cosmological parameters listed in table 1 as our baseline cosmological model and compute the evolution from adiabatic (isentropic) initial conditions using DISCO-EB, CAMB, and CLASS. We compare the resulting power spectra for baryons, CDM, the massive neutrino, as well as the combined total matter (CDM+baryon+massive neutrino) and the ‘ $b + c$ ’ (baryon+CDM) spectra to those obtained from CAMB and CLASS at $z = 0$ in figure 2.

In all cases, we find excellent agreement between the three codes, with relative differences of at most a few per mille for all quantities of interest, with the sole exception of the massive neutrino perturbations. Generally the difference is similar w.r.t. both CLASS and CAMB, with possibly slightly better agreement between DISCO-EB and CLASS for everything but the massive neutrinos spectrum, and slightly better agreement between DISCO-EB and CAMB for the massive neutrinos. We find that the spectra of massive neutrino perturbations differ at more than 1 per cent for $k \gtrsim 1 \text{ Mpc}^{-1}$, where they are suppressed by about 10^{-6} relative to the CDM perturbations due to free streaming. In practice, these differences are




Figure 2. Comparison at $a = 1$ of the performance of our differentiable Einstein-Boltzmann solver vs. CAMB and CLASS. The top panel shows the power spectrum of density perturbations in baryons (solid), CDM (dashed), and the 0.06 eV massive neutrino (dot-dashed). The lower panels indicate the relative errors $P/P_{\text{ref}} - 1$ for CDM, baryons, massive neutrinos, CDM+baryons, and total matter = CDM+baryons+massive neutrino. The grey bands indicate 1 per cent deviation (light grey) and 0.5 per cent deviation (darker grey). All EB solvers were run on high-accuracy settings for this comparison, but we note that it might well be possible to improve agreement further with other parameter choices.

therefore not significant for the matter power spectrum, and entirely attributable to the chosen numerical precision.

A similar comparison at $z = 99$ ($a = 10^{-2}$) (shown in figure 8 in the appendix) yields very similar results with few per-mille differences in all perturbations. Slightly larger differences exist for baryons, presumably due to differences in the recombination solver used. To improve agreement even further at higher z , we would therefore have to improve the quality of the recombination solver, which we leave for future work (see discussion in section 2.5). Again, larger differences exist for massive neutrinos in the free-streaming suppressed regime, which are however of no real importance.

All in all, the agreement is therefore excellent, and we conclude that DISCO-EB is able to reproduce the results of CAMB and CLASS at the per mille accuracy level for all quantities of interest. We caution that the tiny differences between the results of CAMB and CLASS might still be reduced by more optimal parameter choices.

3.3 Gradient of the power spectrum w.r.t. cosmological parameters

Having a fully differentiable Einstein-Boltzmann solver permits taking derivatives with respect to any parameter of the model. In this subsection, we demonstrate this capability by computing the Jacobian matrix of the total matter (CDM+baryon+massive neutrino) density power spectrum at fixed time ($a = 1$) with respect to all cosmological model input parameters. This is achieved by wrapping the power spectrum calculation in a JAX function, which takes the parameters as an argument, and then calls the `jacfwd` method of the JAX library. Specifically, let $\boldsymbol{\vartheta} \in \mathbb{R}^n$ be the vector of cosmological model input parameters, where $n = 12$ for our baseline cosmology (see table 1), and specifically

$$\boldsymbol{\vartheta} = \left(H_0, \Omega_m, \Omega_b, N_{\text{eff}}, m_\nu, T_{\text{CMB}}, Y_{\text{He}}, A_s, n_s, w_0, w_a, c_a^2 \right)^\top. \quad (3.1)$$

In principle, all physical constants and other parameters, including e.g. numerical parameters, could be added here. Let $P(k, a | \boldsymbol{\vartheta})$ be the vector of power spectrum values at fixed time, then we are interested in computing the logarithmic Jacobian

$$\mathbf{J}_{\boldsymbol{\vartheta}} := \nabla_{\boldsymbol{\vartheta}} \log P(k, a | \boldsymbol{\vartheta}). \quad (3.2)$$

The result of the autodiff computation is shown in figure 3 as a solid black line. The respective component of the Jacobian is indicated by a label in each plot. For comparison, we also show the finite-difference approximation of the Jacobian computed with CAMB and CLASS (note that we have adopted $\ell_{\text{max}} = 32$ for CLASS in this case). Specifically, we compute the finite difference approximation of the Jacobian as

$$\mathbf{J}_{\vartheta_i}^{\text{FD}} = \frac{\log P(k, a | \boldsymbol{\vartheta} + \frac{\epsilon}{2} \vartheta_i \mathbf{e}_i) - \log P(k, a | \boldsymbol{\vartheta} - \frac{\epsilon}{2} \vartheta_i \mathbf{e}_i)}{\epsilon \vartheta_i} + \mathcal{O}(\epsilon^2), \quad (3.3)$$

which means that for second-order finite differences two calls to the EB solver are needed per model parameter. We show the results for $\epsilon = 10^{-2}$ in figure 3. The top panel in each case represents the value of the derivative as a function of the wave number, while the bottom panel shows the relative difference between the DISCO-EB autodiff computation and the finite difference solutions using CAMB and CLASS.

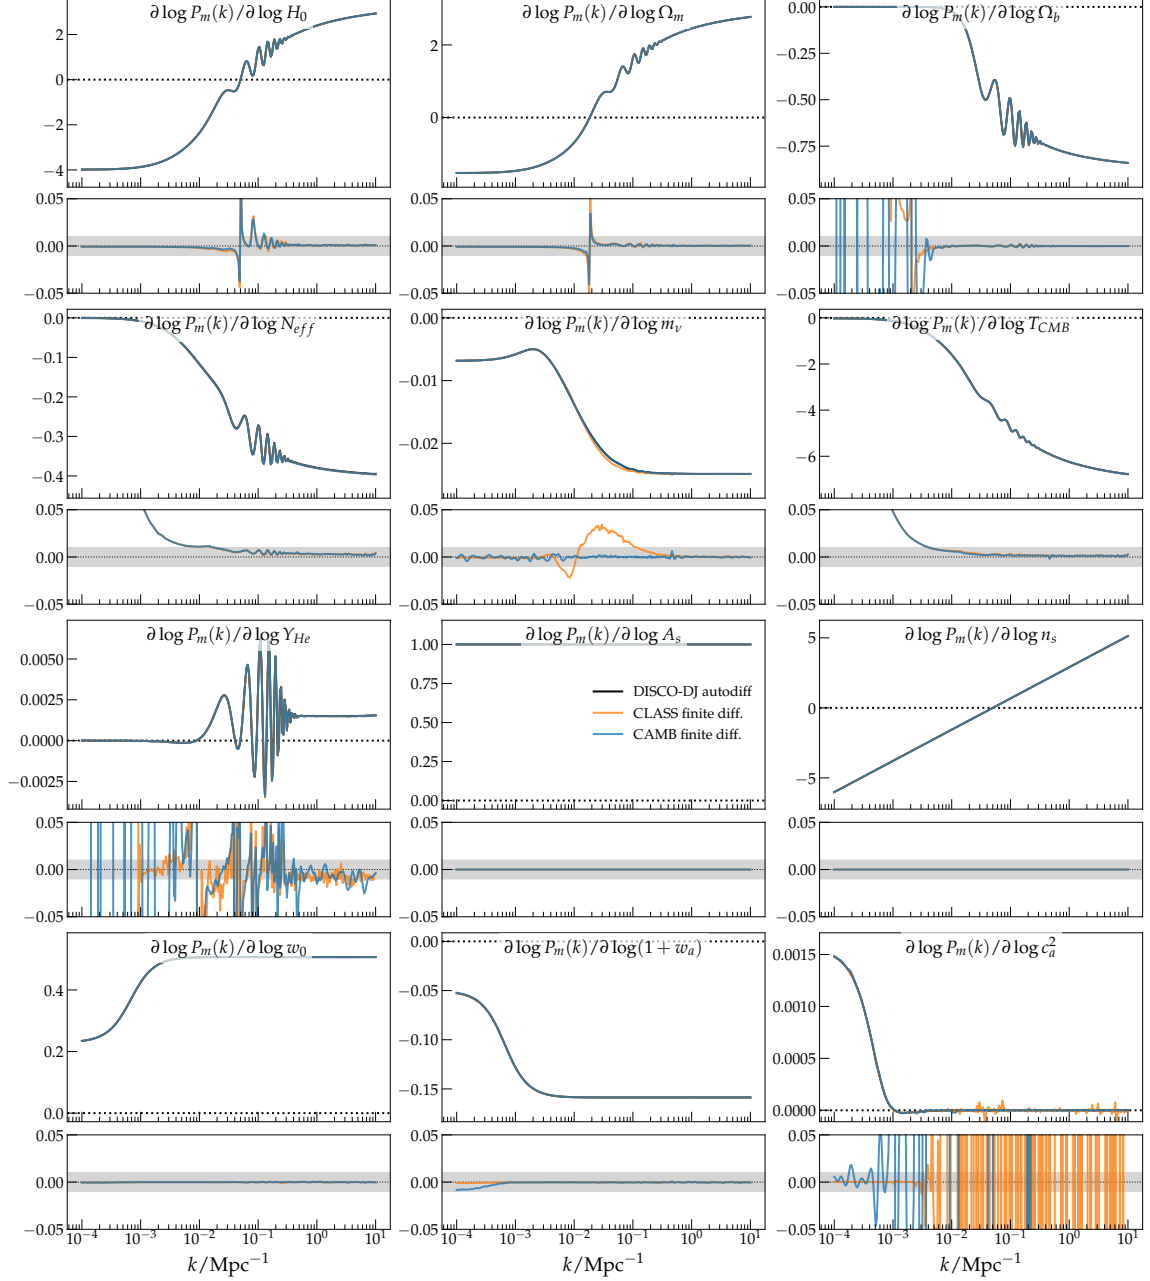


Figure 3. Derivative of the total matter (CDM+baryon+massive neutrino) density power spectrum at fixed time ($a = 1$) with respect to all cosmological model input parameters, evaluated at the fiducial parameter values. We show the result of the DISCO-EB autodiff computation (black) and a second-order finite-difference approximation based on CLASS (orange) and CAMB, computed using eq. (3.3) with $\epsilon = 10^{-2}$. The respective derivative is indicated by a label in each plot. The top panel in each case represents the value of the derivative as a function of the wave number, while the bottom panel shows the relative difference between the DISCO-EB autodiff computation and the finite difference solutions.

We note that the autodiff evaluation and the finite difference approximation agree generally very well to mostly better than one per cent. We found that a good bias-variance

tradeoff requires $\epsilon \approx 10^{-2}$, while for smaller values, the finite difference estimates become very noisy. Nonetheless, where visible, the autodiff version shows fewer signs of numerical artefacts. The only significant difference is found for the neutrino mass derivative w.r.t. the finite difference result for CLASS, which is attributable to the reduced accuracy we used (specifically not setting `ncdm_fluid_approximation` as this leads to prohibitive runtimes). We note that the finite difference approximation is computationally expensive, since it requires two calls to the EB solver per parameter, while the autodiff computation is computationally much cheaper. Due to the large impact of optimizations and specific parameter choices on the performance of the EB solvers, we prefer to refrain from a detailed comparison of the computational performance of the autodiff and finite difference approaches.

Finally, in figure 4, we show full time and space dependence of the derivative of the total matter power spectrum (cdm+baryons+massive neutrino) w.r.t. the model parameters. Each panel shows the colour-coded change in the power spectrum, with red colours indicating an increase and blue a decrease, as a function of scale factor (y -axis) and wave number (x -axis). A white colour indicates that the respective parameter has no impact on the matter power spectrum at that time and scale. For this analysis, we extend this (qualitative) analysis to earlier times ($a = 10^{-5}$) revealing how the BAO feature propagates to increasingly larger scale prior to decoupling at $a \approx 10^{-3}$. Also intuitively, the neutrino mass dependence clearly shows when the neutrino becomes non-relativistic at $z \lesssim 100$, while the signature of the dark energy EOS parameters w_0 , w_a and c_a^2 is restricted to very late times $z \lesssim 1$, with c_a^2 clearly impacting only the sound horizon in the dark energy fluid.

4 Applications

In this section, we demonstrate a few applications of the differentiable Einstein-Boltzmann solver. We note that these are not meant to be exhaustive in any way. Specifically, two main applications of a differentiable solver will be for parameter inference and the construction of emulators, aspects that we postpone to future work. Nonetheless, the cases we highlight here are meant to demonstrate the applicability and integration of the solver for a few select applications relevant to large-scale structure cosmology.

4.1 Redshift-space correlators for large-scale structure studies

The matter power spectrum itself is not directly observable but can be constrained e.g. through the clustering of galaxies. Galaxies are, however, biased tracers of the underlying matter distribution (cf. [49] for a review) and their three-dimensional distribution can generally only be determined in redshift space [50]. As a next proof-of-concept test, we demonstrate a differentiable calculation of the redshift-space power spectrum and the correlation function multipoles of biased tracers, as this is a common input for large-scale structure studies.

4.1.1 Linear redshift space distortions

The proper motions v_{\parallel} of galaxies parallel to the line-of-sight (LOS) contribute to the measured redshift

$$1 + z_{\text{obs}} = (1 + z_{\text{cosmo}}) (1 + v_{\parallel}/c) \quad (4.1)$$




Figure 4. Time and spatial dependence of the logarithmic derivative of the total matter power spectrum w.r.t. the cosmological parameters. Each panel shows the colour-coded change in the power spectrum in arbitrary normalisation, with red indicating an increase and blue a decrease w.r.t. an arbitrary baseline, as a function of scale factor (y -axis) and wave number (x -axis). Compare with figure 3 which represents a slice at fixed time $a = 1$.

and lead to the redshift space distortion (RSD) effect. At linear order and in the Newtonian limit, the redshift-space power spectrum can be computed from the matter power spectrum using the Kaiser formula [50]. For linearly biased tracers with density contrast δ_g it states that $\delta_g^{\text{Kaiser}} := b\delta_m - \mu^2 \frac{\theta_m}{\mathcal{H}}$ and so

$$P_{\text{lin}}(k, \mu; z) = \left[b(z) \delta_m(k, z) - \frac{\mu^2}{\mathcal{H}(z)} \theta_m(k, z) \right]^2 \quad (4.2)$$

where b is the linear bias coefficient (which we fix to $b = 2$ here), $\mu := \cos \alpha$ for an angle α between \mathbf{k} and the line-of-sight, θ_m is the matter velocity divergence, and $\mathcal{H} = a'/a$ is the conformal-time expansion rate. Further, one needs to assume an irrotational velocity field ($\mathbf{v} = \nabla \nabla^{-2} \theta$) and the absence of velocity bias (i.e. galaxy velocity $\theta_g = \theta_m$) for this relation to hold. Commonly, the growth rate f is used to express θ_m in terms of δ_m , but since we have both fields from the EB solver, it is more convenient to work directly with the velocity divergence. The fields δ_m and θ_m are readily obtained from DISCO-EB.

4.1.2 Correlation function multipoles

Redshift space distortions on large scales are conveniently expressed in terms of a multipole expansion of the power spectrum. These redshift-space power spectrum multipoles are commonly defined as the following expansion in Legendre polynomials [51, 52]

$$P_\ell(k; z) := \frac{2\ell + 1}{2} \int_{-1}^1 d\mu P(k, \mu; z) L_\ell(\mu), \quad (4.3)$$

where $\mu = \cos \alpha$ is the cosine of the angle α between the wave vector \mathbf{k} and the line-of-sight direction \mathbf{n} , $P(k, \mu; z)$ is the anisotropic power spectrum of the tracer, and L_ℓ is the Legendre polynomial of order $\ell = 0, 1, 2, \dots$. Then the redshift-space power spectrum multipoles for the Kaiser formula (eq. (4.2), [50]) are given by

$$P_\ell(k; z) = \frac{2\ell + 1}{2} \int_{-1}^1 d\mu \left(b(z) \delta_m(k; z) - \frac{\mu^2}{\mathcal{H}(z)} \theta_m(k; z) \right)^2 L_\ell(\mu) \quad (4.4a)$$

and therefore with $P_{\delta\delta} := \delta_m \delta_m^* = P_m$, $P_{\delta\theta} := \delta_m \theta_m^*$, and $P_{\theta\theta} := \theta_m \theta_m^*$ the only non-zero multipoles for the linear Kaiser model are

$$P_0(k; z) = b^2 P_{\delta\delta} - \frac{2b}{3\mathcal{H}} P_{\delta\theta} + \frac{1}{5\mathcal{H}^2} P_{\theta\theta} \quad (4.4b)$$

$$P_2(k; z) = -\frac{4b}{3\mathcal{H}} P_{\delta\theta} + \frac{4}{7\mathcal{H}^2} P_{\theta\theta} \quad (4.4c)$$

$$P_4(k; z) = \frac{8}{35\mathcal{H}^2} P_{\theta\theta}. \quad (4.4d)$$

Finally, the correlation function multipoles are defined in terms of the power spectrum multipoles via the integral

$$\xi_\ell(r; z) = \frac{i^\ell}{2\pi^2} \int_0^\infty dk k^2 P_\ell(k; z) j_\ell(kr), \quad (4.5)$$

where j_ℓ is the spherical Bessel function of order ℓ . These integrals are typically efficiently computed using the FFTLOG method [52, 53], which can be implemented in JAX in a few lines, and thus allows differentiable computation of the correlation function.

We show the resulting derivatives of the monopole, quadrupole, and hexadecapole w.r.t. the baseline cosmological parameters in figure 5. The respective derivative is indicated by a label in each plot. We have omitted the dark energy parameters here as they predominantly affect the growth history and thus only the normalisation of the correlation function multipoles. Clearly visible are the BAO features in many derivatives (but not all). The calculation in

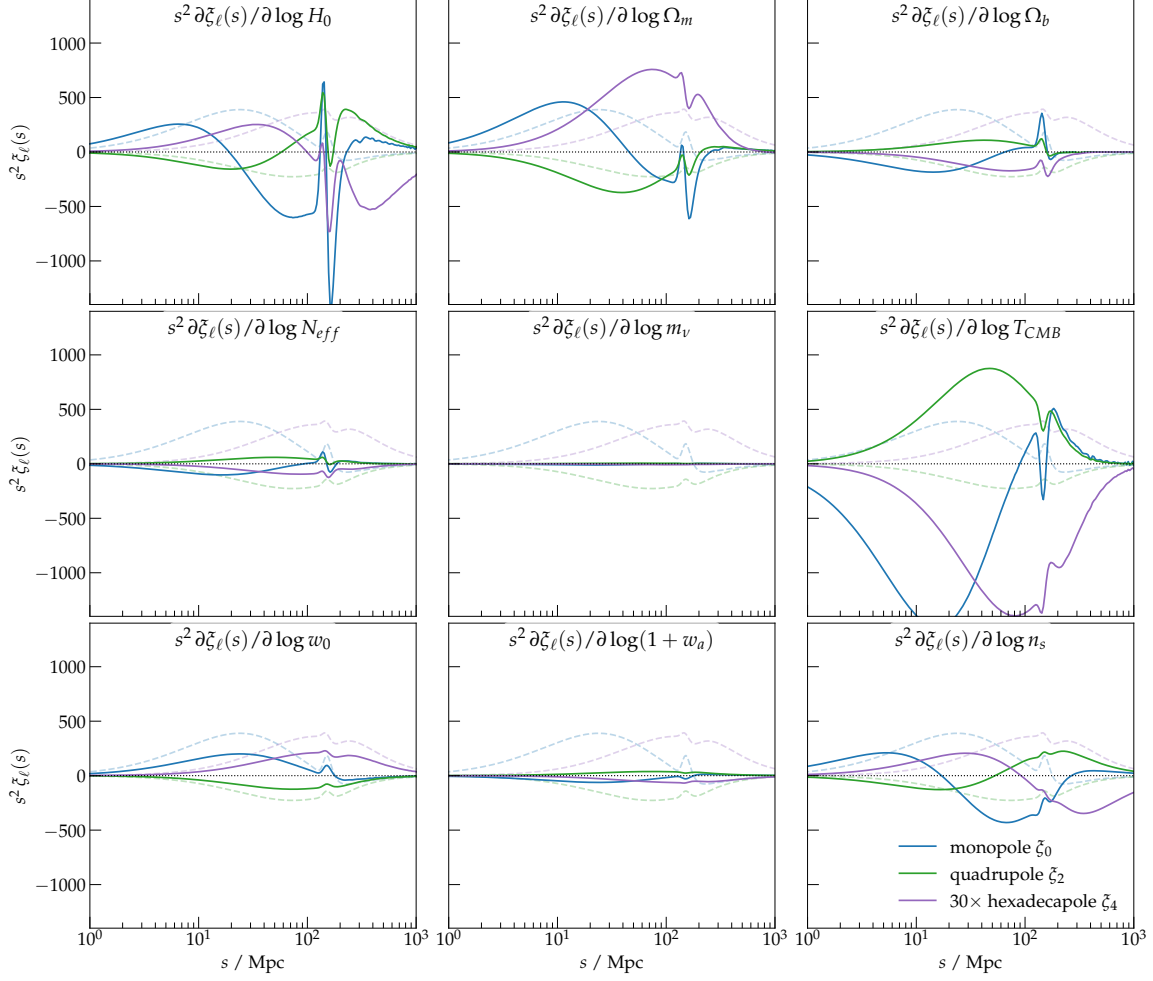


Figure 5. Derivative of the redshift-space correlation function multipoles w.r.t. the cosmological parameters. Each panel shows the derivative as a function of scale r of the monopole $\xi_0(r)$ (blue), the quadrupole $\xi_2(r)$ (green), and the hexadecapole $\xi_4(r)$ (purple), the respective correlation function multipoles (without derivative applied) are shown as light dashed lines. We have multiplied the hexadecapole ξ_4 by a factor of 30 for visual clarity. In this case the derivative is carried through the FFTLOG method and the Einstein-Boltzmann solver.

this case was carried out using 512 modes geometrically distributed between $k = 10^{-5} \text{ Mpc}^{-1}$ and $k = 100 \text{ Mpc}^{-1}$ since the FFTLOG transform is relatively noisy with fewer sampling points, and even in this case, some ringing can be observed at $r \gtrsim 200 \text{ Mpc}$. This could likely be improved by additional tweaks to the FFTLOG method, which goes beyond our minimal FFTLOG implementation, and which we leave for future work.

4.1.3 Two-dimensional power spectrum

In a more realistic analysis, in addition to the Kaiser effect, typically also the Finger-of-God effect and the Alcock-Paczynski effect [54] are considered.

The Alcock-Paczynski (AP) effect. It arises from the fact that the observables redshift and angle on the sky need to be converted to length scales assuming a fiducial cosmology

which might not be the true cosmology. The AP effect can be expressed as a modification of the observed wave vector k_{obs} and the observed angle μ_{obs} in terms of the true wave vector k and angle μ as

$$k_{\text{obs}}(k, \mu, z) := \frac{k}{q_{\perp}} \left[1 + \mu^2 \left(\frac{q_{\perp}^2}{q_{\parallel}^2} - 1 \right) \right]^{\frac{1}{2}}, \quad \mu_{\text{obs}}(\mu, z) := \mu \frac{q_{\perp}}{q_{\parallel}} \left[1 + \mu^2 \left(\frac{q_{\perp}^2}{q_{\parallel}^2} - 1 \right) \right]^{-\frac{1}{2}} \quad (4.6)$$

and volume deformation $V_{\text{obs}} := q_{\perp}^2 q_{\parallel}$. The deformation parameters are

$$q_{\perp}(z) := \frac{D_A(z)}{D_A^{\text{fid}}(z)}, \quad q_{\parallel}(z) := \frac{\mathcal{H}^{\text{fid}}(z)}{\mathcal{H}(z)} \quad (4.7)$$

where $D_A(z)$ is the angular diameter distance and $\mathcal{H}(z) = a'/a$ is the conformal Hubble function. The superscript ^{fid} indicates the value in the reference (fiducial) cosmology. which leads to a volume deformation $V_{\text{obs}} := q_{\perp}^2 q_{\parallel}$.

The Finger-of-God (FoG) effect. It arises from the fact that galaxies in clusters have a velocity dispersion which leads to a smearing of the observed power spectrum. The effect can be modelled as a Lorentzian damping of the power spectrum in redshift space [55–57]. It amounts to a multiplicative factor of the form

$$F_{\text{FoG}} := \frac{1}{1 + k^2 \mu^2 \sigma_{\theta}^2(z)} \quad \text{where} \quad \sigma_{\theta}^2(z) := \frac{1}{6\pi^2 (\mathcal{H}^{\text{fid}})^2} \int P_{\theta\theta}^{\text{fid}}(k, z) dk \quad (4.8)$$

The observed power spectrum. Putting all together, we can write the observed (linear) power spectrum including all mentioned effects as

$$P_{\text{obs}}(k, \mu; z) = \frac{1}{q_{\perp}^2(z) q_{\parallel}(z)} \frac{P_{\text{lin}}(k_{\text{obs}}(k, \mu, z), \mu_{\text{obs}}(\mu, z); z)}{1 + k_{\text{obs}}^2(k, \mu, z) \mu_{\text{obs}}^2(\mu, z) \sigma_{\theta}^2(z)}. \quad (4.9)$$

Note that the power spectrum of course always also depends on all physical/cosmological parameters $\boldsymbol{\vartheta}$, which we generally have suppressed in this section to avoid cluttering the notation further.

In figure 6, we show the logarithmic gradients of the observed matter power spectrum $\partial \log P_{\text{obs}} / \partial \log \vartheta_i$ with respect to a key subset of model parameters $\boldsymbol{\vartheta}$ as obtained with DISCO-DJ. Wave vectors parallel to the LOS are indicated in the 0° direction, while modes perpendicular are at 90° and 270° . The radial coordinate is $\log k$, which is clipped to be between $k_{\text{min}} = 10^{-4} \text{ Mpc}^{-1}$ (the center point) and $k_{\text{max}} = 10 \text{ Mpc}^{-1}$ (the outer circle). The colour scale indicates a positive dependence of P_{obs} on the parameter (red) or a negative dependence (blue). A white region indicates that the information content due to parameter ϑ_i of the power spectrum vanishes.

A few aspects are notable. First, the BAO feature is clearly visible in the derivatives w.r.t. H_0 , Y_{He} , and the DE EOS parameters w_0 and w_a , indicating that these parameters have a strong impact on the BAO feature (i.e. causing a shift in the BAO scale). Thirdly, the derivatives w.r.t. Ω_m and the bias parameter b reveal a quadrupole pattern in the k - μ plane which is due to their impact on the redshift space distortion effect (Ω_m determines the growth rate $f = d \log D_+ / d \log a$, and b and f directly determine the RSD effect in the Kaiser formula).

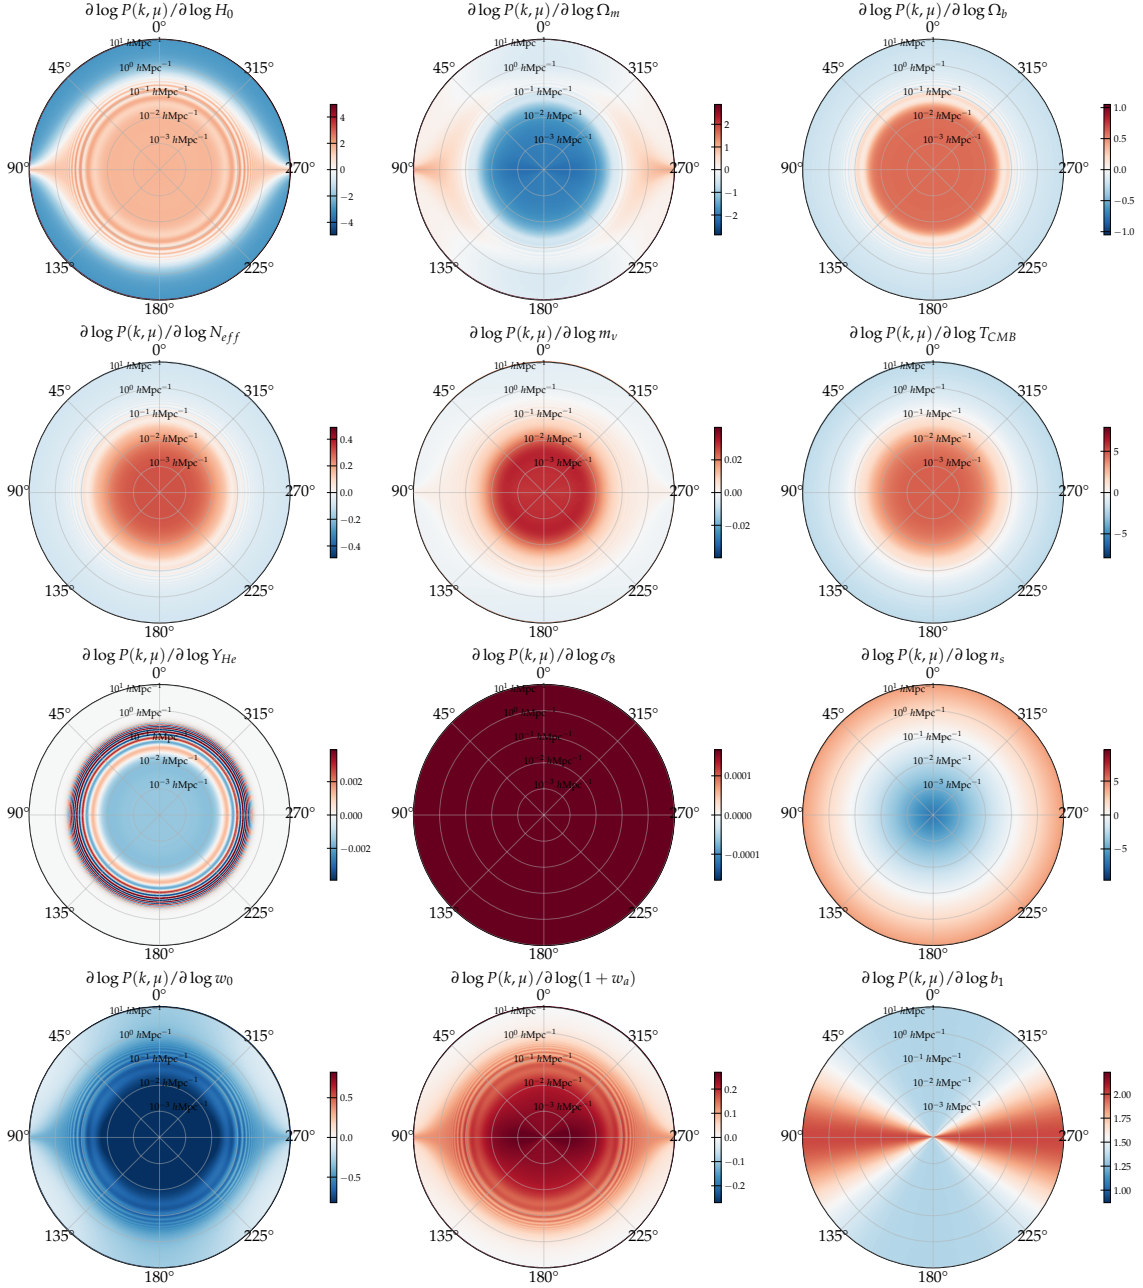


Figure 6. Logarithmic derivatives of the two-dimensional redshift-space matter power spectrum w.r.t. model parameters ϑ , $\partial \log P_{\text{obs}}(k, \mu) / \partial \log \vartheta_i$. The angular coordinate in each plot is the angle between a mode \mathbf{k} and the line-of-sight, while the radial coordinate is k , which is clipped to be between $k_{\text{min}} = 10^{-4} \text{ Mpc}^{-1}$ (the center point) and $k_{\text{max}} = 10 \text{ Mpc}^{-1}$ (the outer circle). The colour scale indicates a relative increase (red) vs. decrease (blue).

4.2 Fisher-forecasting with differentiable solvers

As a proof-of-concept, we demonstrate an application to compute the Fisher information matrix for forecasting. Fisher forecasts are widely used in cosmology for experimental design [58, 59], and they require computing derivatives of the likelihood with respect to the

cosmological parameters. Typically, these derivatives are computed with finite differences, which requires careful tuning for convergence [60]. Autodiff makes the computation of the Fisher matrix easier and avoids the convergence problem.

We demonstrate the computation of a Fisher forecast for the ‘optimistic’ spectroscopic Euclid survey with 4 bins as described in [56, 57]. For a Gaussian likelihood, the elements of the Fisher matrix for a summary statistic mean $y(\ell)$ are

$$F_{ij} = \sum_{\ell} \frac{\partial y(\ell)^{\top}}{\partial \vartheta_i} C^{-1}(\ell) \frac{\partial y(\ell)}{\partial \vartheta_j} \quad (4.10)$$

where $C(\ell)$ is the covariance matrix computed at the fiducial cosmology. Following [56], we account for effective bias, anisotropies due to redshift space distortions, and redshift uncertainty in the observed two-dimensional power spectra, given by eq. (4.9). When considering a tomographic survey, the bias parameter $b_1(z)$ and redshift error, $\sigma_z(z)$ are of course redshift dependent. Specifically, one has that redshift errors grow roughly like $\sigma_z = (1+z)\sigma_{z,0}$.

Linear covariance and survey properties. In our simplified proof-of-concept analysis, we assumed the linear approximation of the covariance matrix [56, 57]:

$$C(\mathbf{k}, \mathbf{k}') \approx \frac{2(2\pi)^2}{V_s(z)} P_{\text{obs}}^2(\|\mathbf{k}\|, \mu; z) \delta_D(\mathbf{k} - \mathbf{k}'), \quad (4.11)$$

where δ_D is the Dirac delta distribution and $V_s(z)$ is the survey volume for the redshift bin at z . We assume a galaxy density for this bin of

$$n(z) = \frac{d^2 N(\Omega, z)}{d\Omega dz} \frac{A_{\text{survey}}}{V_s(z)} \Delta z. \quad (4.12)$$

The survey area in the sky is $A_{\text{survey}} = 15000 \text{ deg}^2$, and the rest of the parameters are set as in table 3 in ref. [56] (equivalent to table 3 in [57]).

Observed power spectrum. In the standard Euclid forecasting [56, 57], the observed power spectrum \hat{P}_{obs} is the linear spectrum P_{obs} from eq. (4.9) modified in three additional ways: (1) the damping of the BAO feature due to non-linearities is mimicked by a blending of the linear spectrum with a filtered version with BAO wiggles removed. Specifically, as in [57], we apply a 3rd order Savitzky-Golay filter of width $\Delta \log k \approx 1.53$ to obtain the filtered $P_{\text{obs,nw}}$. This is then blended with the original spectrum to yield

$$\tilde{P}_{\text{obs}}(k, \mu; z) := P_{\text{obs}}(k, \mu; z) e^{-g(k, \mu; z)k^2} + P_{\text{obs,nw}}(k, \mu; z) (1 - e^{-g(k, \mu; z)k^2}) \quad (4.13)$$

with the transition scale

$$g(k, \mu; z) := \frac{1}{6\pi^2} \left[1 - \mu^2 + \mu^2 (1 + f^{\text{fid}}(k, z))^2 \right] \int dk P_{\delta\delta}(k, z). \quad (4.14)$$

Then, (2) the result is multiplied by a Gaussian dispersion kernel (cf. [61, 62]) modelling the impact of redshift measurement errors, and (3) receives an additional shot-noise contribution

(note that we follow here in notation [57], rather than [56] who absorb the $1/n(z)$ contribution into an effective survey volume)

$$\hat{P}_{\text{obs}}(k, \mu; z) := \tilde{P}_{\text{obs}}(k, \mu; z) G_z(k, \mu; z) + \left[\frac{1}{n(z)} + \delta P_s(z) \right] \quad (4.15)$$

where $\delta P_s(z) = \delta P_s^{\text{bin}}$ is a shot-noise nuisance parameter for each bin (with fiducial value 0), and the redshift error kernel is given by

$$G_z(k, \mu; z) := \exp \left(-\frac{c^2 k^2}{(\mathcal{H}^{\text{fid}})^2} \mu^2 \sigma_{z,0}^2 \right), \quad (4.16)$$

where $\sigma_{z,0}(1+z) = \sigma_z$ is the measurement error on the redshift (we adopt $\sigma_{z,0} = 0.001$ as the fiducial value for the spectroscopic Euclid survey). Similarly to [57] we allow for an additional nuisance parameter correcting the bias, i.e. we replace in eq. (4.9)

$$b(z) \rightarrow b(z) + \frac{\sigma_8(0)}{\sigma_8(z)} \delta b(z), \quad (4.17)$$

where $b(z) = b^{\text{bin}}$ is the fiducial mean bias per bin and $\delta b(z) = \delta b^{\text{bin}}$ is an additional nuisance parameter for which the trivial dependence on changes in the normalisation is corrected by the scaling with σ_8 .

Fisher matrix. The elements of the Fisher matrix are then given by [56, 58, 63]

$$F_{ij}^{\text{bin}} := \frac{1}{8\pi^2} \int_{-1}^1 d\mu \int_{k_{\text{min}}}^{k_{\text{max}}} dk k^2 \frac{\partial \log \hat{P}_{\text{obs}}}{\partial \vartheta_i} \frac{\partial \log \hat{P}_{\text{obs}}}{\partial \vartheta_j} V_s(z_{\text{bin}}) \quad (4.18)$$

for each tomographic bin. In the k -space integral, we employ Euclid’s ‘optimistic’ cuts on the wave number k included in the Fisher analysis as $k_{\text{min}} = 10^{-3} h^{\text{fid}} \text{Mpc}^{-1}$ and $k_{\text{max}} = 0.3 h^{\text{fid}} \text{Mpc}^{-1}$. We assume the four tomographic bins described in table 3 of [56]. The total Fisher matrix is then computed by summing over the different tomographic bins

$$F_{ij} = \sum_{\text{bin}} F_{ij}^{\text{bin}}, \quad (4.19)$$

where we computed the power spectra for each tomographic bin at the bin’s mean redshift. For this Fisher forecast, we use the Euclid fiducial values of cosmological parameters from [57]. Figure 7 shows the resulting Fisher contours. The DISCO-EB Fisher contours (in green) agree well with the forecast of [64], whose results we reproduce here, showing that DISCO-DJ can provide reliable Fisher matrices. We see a discrepancy in expected constraints on n_s in comparison to the official forecast, which we have not managed to track down and that are likely due to the minor differences in the modelling.

In addition, the Fisher matrix provided by DISCO-DJ is itself differentiable, which allows for optimisation of the figure of merit and also for survey optimisation (see also [33]). The source code for our forecasting will be made available along with the main source code of DISCO-EB.




Figure 7. 1– and 2– σ Fisher contours predicted for the ‘optimistic’ Euclid spectroscopic galaxy clustering survey using a simplified model based on DISCO-EB (green) in comparison to comparable predictions from the official Euclid forecast [64].

5 Summary and conclusions

In this article, we have presented the Einstein-Boltzmann module of DISCO-DJ (**D**ifferentiable **S**imulations for **C**osmology — **D**one with **J**AX), which implements a differentiable Einstein-Boltzmann solver for cosmological perturbations and will be part of a comprehensive differentiable package for large-scale structure cosmology. DISCO-DJ is based on the JAX library, which provides automatic differentiation of PYTHON and NUMPY code. Automatic differentiation is a powerful tool for the development of differentiable predictions, as it allows the efficient and accurate computation of derivatives of a model with respect to any input parameter.

In this article, we have introduced the DISCO-DJ Einstein-Boltzmann solver and demonstrated its performance by comparing it to the well-established CAMB and CLASS solvers. Notably, we find that all three solvers typically agree to better than a few per mille when run

with high precision settings. We have also demonstrated the accuracy of the computation of the Jacobian matrix of the total matter density power spectrum at fixed time with respect to all cosmological model input parameters by comparing it to finite-difference approximations obtained with CAMB and CLASS. We have not (yet) provided a detailed comparison of the computational performance of the autodiff and finite difference approaches, since this would require an optimization of parameters and settings for each code given a desired target accuracy, which is beyond the scope of this first paper.

In section 4, we have demonstrated the power of differentiable solvers by computing the Jacobian matrix of the total matter density power spectrum at fixed time with respect to all cosmological model input parameters. We have also shown the full time and space dependence of the derivative of the total matter power spectrum w.r.t. the model parameters. We have validated the correctness of the Jacobian matrix by comparing it to finite-difference approximations. We have demonstrated the power of differentiable solvers by computing the Jacobian matrix of the total matter density power spectrum at fixed time with respect to all cosmological model input parameters. We have also shown the full time and space dependence of the derivative of the total matter power spectrum w.r.t. the model parameters. We have validated the correctness of the Jacobian matrix by comparing it to finite-difference approximations.

The output of an Einstein-Boltzmann solver is typically the input for predictions of the summary statistic of an observable. To demonstrate how easy it is to compute such predictions with differentiable solvers, we have computed the redshift-space power spectrum and the correlation function multipoles of biased tracers. We have also demonstrated the computation of a Fisher forecast for the Euclid spectroscopic survey. While state-of-the-art modelling would have to take into account also non-linearities and other effects, these results serve as proof-of-concept use cases for our differentiable Einstein-Boltzmann solver.

By visualisation of parameter gradients of a model, we see value also of differentiable solvers for educational purposes as they are capable of directly showing the impact of changes in the (cosmological) model parameters on the observables.

Beyond the applications shown here, we believe that a differentiable Einstein-Boltzmann solver will be useful for more efficient likelihood maximization and sampling, the development of likelihood-free inference methods, and in particular the development of differentiable non-linear models, such as differentiable N -body solvers, or differentiable non-linear perturbation theory, based on e.g. the effective field theory of LSS (see e.g. [65] for a review) or high-order Lagrangian perturbation theory (e.g. [66, 67]). We expect that they also allow the development of more efficient emulators based on Einstein-Boltzmann solvers [68–71], as they allow inclusion of information from the derivatives of the power spectrum with respect to the cosmological model parameters when building the emulator and should thus require much fewer training points. Naturally, the gradient information can be made part of the emulator itself (such as [72]), which would allow for usage of gradient information in the inference process.

DISCO-DJ is a work in progress, and we plan to extend it in several directions. Most importantly, this module will serve as the starting point for a comprehensive differentiable package for large-scale structure cosmology, which will include a differentiable n -th order Lagrangian perturbation theory (LPT) module, as well as an N -body code. Several papers

in this direction are currently in preparation. We also invite the community to extend the Einstein-Boltzmann solver to include non-standard cosmological models, such as modified gravity and alternative dark energy and dark matter models.

We make the DISCO-DJ code publicly available upon publication of this paper under the URL <https://github.com/ohahn/DISCO-EB>.

Acknowledgments

The authors thank Raul Angulo, Mateja Gosenca, Julien Lesgourgues, Antony Lewis, Thomas Montandon, Cristiano Porciani, and Cornelius Rampf for discussions and/or comments on an earlier draft of this paper, as well as our anonymous referee for insightful comments. We thank in particular also Julien Lesgourgues and Santiago Casas for discussions about and the sharing of their most recent Fisher forecasts for Euclid with us. The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC). The authors declare no conflicts of interest or external support in the preparation of this manuscript.

A Governing equations

In the DISCO-DJ Einstein-Boltzmann solver, we currently implement the linearised equations for perturbations of the metric, as well as CDM, baryons, relativistic species, massive neutrinos, and a dark energy fluid, which we briefly summarise here. Specifically we have implemented the following set of equations, in the *synchronous* gauge, that we reproduce here for reference. We largely follow [5] and use their notation, except that we denote derivatives w.r.t. conformal time by a dash rather than a dot.

Metric perturbations. In the synchronous gauge, coordinates are defined by freely falling observers, and scalar metric perturbations are described by two scalar potentials η and h , cf. [5]. Their evolution is governed by the Einstein equations, which in Fourier space read

$$k^2\eta - \frac{1}{2}\mathcal{H}h' = 4\pi Ga^2 \delta T_0^0 \quad (\text{A.1a})$$

$$k^2\eta' = 4\pi Ga^2(\bar{\rho} + \bar{P})\theta \quad (\text{A.1b})$$

$$h'' + 2\mathcal{H}h' - 2k^2\eta = -8\pi Ga^2\delta T_i^i \quad (\text{A.1c})$$

$$h'' + 6\eta'' + 2\mathcal{H}(h' + 6\eta') - 2k^2\eta = -24\pi Ga^2(\bar{\rho} + \bar{P})\sigma. \quad (\text{A.1d})$$

These equations are overdetermined so that e.g. h' is not a dynamical degree of freedom.

CDM. As coordinates are defined by freely falling observers, the equations of motion of collisionless cold dark matter are particularly simple:

$$\delta_c' = -\frac{h'}{2} \quad (\text{A.2a})$$

$$\theta_c = \sigma_c = 0. \quad (\text{A.2b})$$

Baryons. In contrast, baryons have a finite sound speed and couple to the photons via Thomson scattering while free electrons are present. The temperature and ionization state are determined at the background level by the thermal history solver. The baryon equations of motion (EOMs) are then given by

$$\delta'_b = -\theta_b - \frac{h'}{2} \quad (\text{A.3a})$$

$$\theta'_b = -\mathcal{H}\theta_b + c_s^2 k^2 \delta_b + \frac{4}{3} \frac{\bar{\rho}_\gamma}{\bar{\rho}_b} a n_e \sigma_T (\theta_\gamma - \theta_b) \quad (\text{A.3b})$$

where n_e is the homogeneous mean electron density, σ_T is the Thomson scattering cross section, and c_s the adiabatic sound speed of the baryon fluid. The latter is given by

$$c_s^2 = \frac{k_B T_b}{\mu} \left(1 - \frac{1}{3} \frac{d \log T_b}{d \log a} \right) \quad (\text{A.3c})$$

where

$$T'_b = -2\mathcal{H}T_b + \frac{8}{3} \frac{\mu}{m_e} \frac{\bar{\rho}_\gamma}{\bar{\rho}_b} a n_e \sigma_T (T_\gamma - T_b) \quad (\text{A.3d})$$

is the evolution of the baryon temperature, which explicitly depends on the mean molecular weight μ and electron density, which are computed by the recombination solver. The ratio $\bar{\rho}_\gamma/\bar{\rho}_b$ denotes the photon-to-baryon ratio.

Photons. Photons follow Bose-Einstein statistics with two polarisation states, and are coupled to the baryons via Thomson scattering off free electrons. Their EOMs are given by a Boltzmann hierarchy of moments. The first three moments (corresponding to density, velocity, and shear perturbations) are given by

$$\delta'_\gamma = -\frac{4}{3}\theta_\gamma - \frac{2}{3}h' \quad (\text{A.4a})$$

$$\theta'_\gamma = k^2 \left(\frac{1}{4}\delta_\gamma - \sigma_\gamma \right) + a n_e \sigma_T (\theta_b - \theta_\gamma) \quad (\text{A.4b})$$

$$2\sigma'_\gamma = F_\gamma^{(2)'} = \frac{8}{15}\theta_\gamma - \frac{3}{5}kF_\gamma^{(3)} + \frac{4}{15}h' + \frac{8}{5}\eta' - \frac{9}{5}a n_e \sigma_T \sigma_\gamma + \frac{1}{10}a n_e \sigma_T (G_\gamma^{(0)} + G_\gamma^{(2)}) \quad (\text{A.4c})$$

The higher Boltzmann moments for $\ell_{\max} > \ell \geq 3$ are given by

$$F_\gamma^{(\ell)'} = \frac{k}{2\ell+1} \left[\ell F_\gamma^{(\ell-1)} - (\ell+1) F_\gamma^{(\ell+1)} \right] - a n_e \sigma_T F_\gamma^{(\ell)} \quad (\text{A.4d})$$

and the hierarchy is truncated by hand as in [5] by setting

$$F_\gamma^{(\ell_{\max})'} \approx k F_\gamma^{(\ell_{\max}-1)} - \frac{\ell_{\max}+1}{\tau} F_\gamma^{(\ell_{\max})} - a n_e \sigma_T F_\gamma^{(\ell_{\max})}. \quad (\text{A.4e})$$

The polarisation moments for $0 \leq \ell < \ell_{\max}$ are given by

$$G_\gamma^{(\ell)'} = \frac{k}{2\ell+1} \left[\ell G_\gamma^{(\ell-1)} - (\ell+1) G_\gamma^{(\ell+1)} \right] + a n_e \sigma_T \left[-G_\gamma^{(\ell)} + \frac{1}{2} \left(F_\gamma^{(2)} + F_\gamma^{(0)} + G_\gamma^{(2)} \right) \left(\delta_{\ell 0} + \frac{1}{5} \delta_{\ell 2} \right) \right], \quad (\text{A.4f})$$

(where $\delta_{\ell j}$ is a Kronecker symbol). The respective hierarchy truncation is given by

$$G_{\gamma}^{(\ell_{\max})'} \approx k G_{\gamma}^{(\ell_{\max}-1)} - \frac{\ell_{\max} + 1}{\tau} G_{\gamma}^{(\ell_{\max})} - a n_e \sigma_T G_{\gamma}^{(\ell_{\max})}. \quad (\text{A.4g})$$

Massless neutrinos. Massless neutrinos follow Fermi-Dirac statistics, with EOMs given by a Boltzmann hierarchy of moments, of which the first three are

$$\delta'_{\nu} = -\frac{4}{3}\theta_{\nu} - \frac{2}{3}h' \quad (\text{A.5a})$$

$$\theta'_{\nu} = k^2 \left(\frac{1}{4}\delta_{\nu} - \sigma_{\nu} \right) \quad (\text{A.5b})$$

$$2\sigma'_{\nu} = F_{\nu}^{(2)'} = \frac{8}{15}\theta_{\nu} - \frac{3}{5}kF_{\nu}^{(3)} + \frac{4}{15}h' + \frac{8}{5}\eta'. \quad (\text{A.5c})$$

The higher Boltzmann moments for $3 \leq \ell < \ell_{\max}$ are given by

$$F_{\nu}^{(\ell)'} = \frac{k}{2\ell + 1} \left[\ell F_{\nu}^{(\ell-1)} - (\ell + 1)F_{\nu}^{(\ell+1)} \right] \quad (\text{A.5d})$$

with a truncation given by

$$F_{\nu}^{(\ell_{\max})'} \approx k F_{\nu}^{(\ell_{\max}-1)} - \frac{\ell_{\max} + 1}{\tau} F_{\nu}^{(\ell_{\max})}. \quad (\text{A.5e})$$

Massive neutrinos. Finally, for massive neutrinos, a momentum-dependent Boltzmann hierarchy has to be evolved. This is done by discretising the momentum space into n bins with momenta $q \in \{q_0, \dots, q_n\}$, and evolving the hierarchy for each bin. The moments of the neutrino distribution function $\psi_q^{(\ell)}$ are then given for each momentum bin, i.e.

$$\psi_q^{(0)'} = -\frac{qk}{\epsilon}\psi_q^{(1)} + \frac{1}{6}h' \frac{d \log f_0}{d \log q}, \quad (\text{A.6a})$$

$$\psi_q^{(1)'} = \frac{qk}{3\epsilon} \left(\psi_q^{(0)} - 2\psi_q^{(2)} \right), \quad (\text{A.6b})$$

$$\psi_q^{(2)'} = \frac{qk}{5\epsilon} \left(2\psi_q^{(1)} - 3\psi_q^{(3)} \right) - \left(\frac{1}{15}h' + \frac{2}{5}\eta' \right) \frac{d \log f_0}{d \log q}, \quad (\text{A.6c})$$

where $\epsilon^2 := q^2 + a^2 m_{\nu}^2$. The higher Boltzmann moments for $3 \leq \ell < \ell_{\max}$ are given by

$$\psi_q^{(\ell)'} = \frac{qk}{(2\ell + 1)\epsilon} \left(\ell \psi_q^{(\ell-1)} - (\ell + 1)\psi_q^{(\ell+1)} \right), \quad (\text{A.6d})$$

and the hierarchy is truncated as

$$\psi_q^{(\ell_{\max})'} \approx \frac{qk}{\epsilon} \psi_q^{(\ell_{\max}-1)} - \frac{\ell_{\max} + 1}{\tau} \psi_q^{(\ell_{\max})}. \quad (\text{A.6e})$$

The massive neutrino perturbations entering the field equations are then given in terms of the lowest moments as

$$\delta\rho_\nu = \frac{4\pi}{a^4} \int dq q^2 \epsilon f_0(q) \psi_0(q) , \quad (\text{A.6f})$$

$$\delta P_\nu = \frac{4\pi}{3a^4} \int dq \frac{q^4}{\epsilon} f_0(q) \psi_1(q) , \quad (\text{A.6g})$$

$$(\bar{\rho}_\nu + \bar{P}_\nu)\theta_\nu = \frac{4\pi k}{a^4} \int dq q^3 f_0(q) \psi_1(q) , \quad (\text{A.6h})$$

$$(\bar{\rho}_\nu + \bar{P}_\nu)\sigma_\nu = \frac{8\pi}{3a^4} \int dq \frac{q^4}{\epsilon} f_0(q) \psi_2(q) , \quad (\text{A.6i})$$

where $f_0(q)$ is the background neutrino distribution function, which is assumed to be Fermi-Dirac with zero chemical potential. The momentum bins are chosen in order to approximate well these integrals over the Fermi-Dirac distribution, cf. [42] and section 2.4.

Finally, in the presence of massive neutrinos, the effective degrees of freedom of the ultrarelativistic species are modified. We follow [73] (their eq. 84) and set the effective number of relativistic species given N_{eff} as

$$N_{\nu,\text{rel}} = N_{\text{eff}} - T_{\nu,\text{massive}}^4 \left(\frac{4}{11} \right)^{-4/3} N_{\nu,\text{massive}} , \quad (\text{A.7})$$

where in our case $N_{\nu,\text{massive}} = 1$ for one single species of massive neutrino currently. The massive neutrino temperature $T_{\nu,\text{massive}}$ is by default set to the temperature of the massless neutrinos $T_{\nu,\text{massive}} = (4/11)^{1/3} T_{\text{CMB}}$ so that in this case $N_{\nu,\text{rel}} = N_{\text{eff}} - N_{\nu,\text{massive}}$. Note that $T_{\nu,\text{massive}}$ corresponds to CLASS' `T_ncdm` parameter.

Quintessence/dark energy. Finally, we allow for a dark energy perturbation, which is implemented following [74]. We evolve the following perturbation variables

$$\delta'_Q = -(1+w)(\theta_Q + \tfrac{1}{2}h') - 3\mathcal{H}(c_a^2 - w)\delta_Q - 9\mathcal{H}(1+w)(c_a^2 - c_w^2)\frac{\theta_Q}{k^2} \quad (\text{A.8a})$$

$$\theta'_Q = -\mathcal{H}(1 - 3c_a^2)\theta_Q + \frac{c_a^2}{1+w}k^2\delta_Q , \quad (\text{A.8b})$$

where we assume the usual (Chevallier-Polarski-Linder [75, 76]) parameterisation of the equation of state parameter $w := w_0 + w_a(1 - a)$, c_a^2 is the (constant) adiabatic sound speed of the dark energy perturbation, and $c_w^2 := w - w'/(3\mathcal{H}(1+w))$.

B Comparison with Class and Camb at early times

In figure 8, we show the same comparison as in figure 2, but for $z = 99$ ($a = 0.01$). The agreement is very good, with the exception of the massive neutrino perturbations where differences are large for wave numbers suppressed by free-streaming damping. Those scales are however dynamically unimportant. Minor differences are seen in the baryon perturbations, likely caused by our simplified thermal history solver. Finally, we see a difference between the CAMB and CLASS cdm+baryon+massive neutrino spectrum for $k \lesssim 10^{-2} \text{ Mpc}^{-1}$, which is arguably due to differences in how massive neutrinos are weighted in the total matter output spectrum in CLASS compared to CAMB and DISCO-EB.

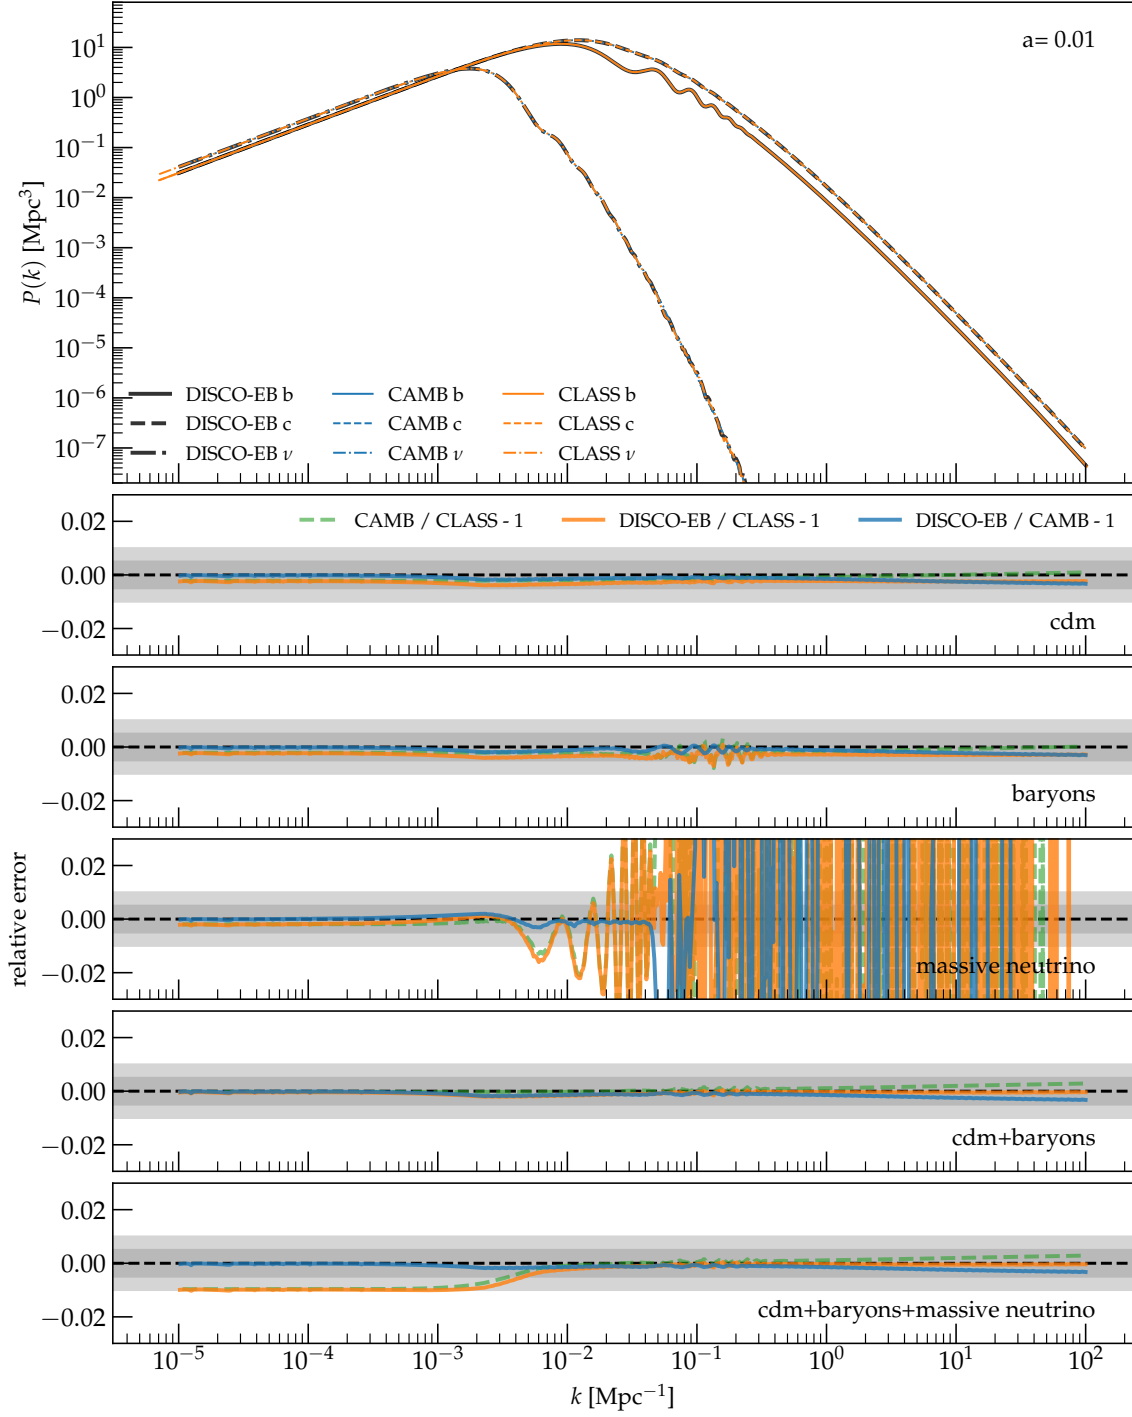


Figure 8. Same as figure 2, but for $z = 99$ ($a = 0.01$).

References

- [1] E.M. Lifshitz, *On the gravitational stability of the expanding universe*, *Zh. Eksp. Teor. Fiz.* **16** (1946) 587.
- [2] E.M. Lifshitz and I.M. Khalatnikov, *Investigations in relativistic cosmology*, *Adv. Phys.* **12** (1963) 185 [[INSPIRE](#)].
- [3] S. Weinberg, *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*, Wiley (1972).
- [4] P. Peebles, *The Large-scale Structure of the Universe*, *Princeton Series in Physics*, Princeton University Press (1980).
- [5] C.-P. Ma and E. Bertschinger, *Cosmological perturbation theory in the synchronous and conformal Newtonian gauges*, *Astrophys. J.* **455** (1995) 7 [[astro-ph/9506072](#)] [[INSPIRE](#)].
- [6] F. Bernardeau, S. Colombi, E. Gaztanaga and R. Scoccimarro, *Large scale structure of the universe and cosmological perturbation theory*, *Phys. Rept.* **367** (2002) 1 [[astro-ph/0112551](#)] [[INSPIRE](#)].
- [7] R.E. Angulo and O. Hahn, *Large-scale dark matter simulations*, [arXiv:2112.05165](#) [[DOI:10.1007/s41115-021-00013-z](#)] [[INSPIRE](#)].
- [8] U. Seljak and M. Zaldarriaga, *A Line of sight integration approach to cosmic microwave background anisotropies*, *Astrophys. J.* **469** (1996) 437 [[astro-ph/9603033](#)] [[INSPIRE](#)].
- [9] WMAP collaboration, *The Microwave Anisotropy Probe (MAP) mission*, *Astrophys. J.* **583** (2003) 1 [[astro-ph/0301158](#)] [[INSPIRE](#)].
- [10] PLANCK collaboration, *Planck Early Results. I. The Planck mission*, *Astron. Astrophys.* **536** (2011) A1 [[arXiv:1101.2022](#)] [[INSPIRE](#)].
- [11] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[INSPIRE](#)].
- [12] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [arXiv:1104.2932](#) [[INSPIRE](#)].
- [13] D. Blas, J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS). Part II. Approximation schemes*, *JCAP* **07** (2011) 034 [[arXiv:1104.2933](#)] [[INSPIRE](#)].
- [14] A. Refregier, L. Gamper, A. Amara and L. Heisenberg, *PyCosmo: An Integrated Cosmological Boltzmann Solver*, *Astron. Comput.* **25** (2018) 38 [[arXiv:1708.05177](#)] [[INSPIRE](#)].
- [15] M. Bucher, K. Moodley and N. Turok, *The General primordial cosmic perturbation*, *Phys. Rev. D* **62** (2000) 083508 [[astro-ph/9904231](#)] [[INSPIRE](#)].
- [16] B. Hu, M. Raveri, N. Frusciante and A. Silvestri, *Effective Field Theory of Cosmic Acceleration: an implementation in CAMB*, *Phys. Rev. D* **89** (2014) 103530 [[arXiv:1312.5742](#)] [[INSPIRE](#)].
- [17] M. Zumalacárregui et al., *hi_class: Horndeski in the Cosmic Linear Anisotropy Solving System*, *JCAP* **08** (2017) 019 [[arXiv:1605.06102](#)] [[INSPIRE](#)].
- [18] A. Zucca, L. Pogosian, A. Silvestri and G.-B. Zhao, *MGCAMB with massive neutrinos and dynamical dark energy*, *JCAP* **05** (2019) 001 [[arXiv:1901.05956](#)] [[INSPIRE](#)].
- [19] R. Hlozek, D. Grin, D.J.E. Marsh and P.G. Ferreira, *A search for ultralight axions using precision cosmological data*, *Phys. Rev. D* **91** (2015) 103512 [[arXiv:1410.2896](#)] [[INSPIRE](#)].

- [20] H. Foidl and T. Rindler-Daller, *Cosmological structure formation in complex scalar field dark matter versus real ultralight axions: A comparative study using class*, *Phys. Rev. D* **105** (2022) 123534 [[arXiv:2203.09396](#)] [[INSPIRE](#)].
- [21] J.C. Hill, E. McDonough, M.W. Toomey and S. Alexander, *Early dark energy does not restore cosmological concordance*, *Phys. Rev. D* **102** (2020) 043507 [[arXiv:2003.07355](#)] [[INSPIRE](#)].
- [22] Google Inc., *Jax: High-performance array computing*, <https://jax.readthedocs.io>.
- [23] P. Kidger, *On Neural Differential Equations*, [arXiv:2202.02435](#).
- [24] G.W. Pettinari et al., *The intrinsic bispectrum of the Cosmic Microwave Background*, *JCAP* **04** (2013) 003 [[arXiv:1302.0832](#)] [[INSPIRE](#)].
- [25] R.E. Wengert, *A simple automatic derivative evaluation program*, *Commun. ACM* **7** (1964) 463.
- [26] Wolfram Inc., *Mathematica*, Version 13.3.
- [27] A. Meurer et al., *SymPy: symbolic computing in Python*, *PeerJ Comput. Sci.* **3** (2017) e103 [[INSPIRE](#)].
- [28] D. Onken and L. Ruthotto, *Discretize-Optimize vs. Optimize-Discretize for Time-Series Regression and Continuous Normalizing Flows*, [arXiv:2005.13420](#).
- [29] J. Heek et al., *Flax: A neural network library and ecosystem for JAX*, (2023).
- [30] M.C. Kaymak et al., *End-to-end differentiable reactive molecular dynamics simulations using JAX*, in *International Conference on High Performance Computing*, Denver, U.S.A. (2023), Springer (2023), pg. 202.
- [31] D. Phan, N. Pradhan and M. Jankowiak, *Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro*, [arXiv:1912.11554](#) [[INSPIRE](#)].
- [32] E. Bingham et al., *Pyro: Deep Universal Probabilistic Programming*, *J. Mach. Learn. Res.* **20** (2019) 1 [[arXiv:1810.09538](#)].
- [33] J.-E. Campagne et al., *JAX-COSMO: An End-to-End Differentiable and GPU Accelerated Cosmology Library*, *Open J. Astrophys.* **6** (2023) 1 [[arXiv:2302.05163](#)] [[INSPIRE](#)].
- [34] D. Piras and A. Spurio Mancini, *CosmoPower-JAX: high-dimensional Bayesian inference with differentiable cosmological emulators*, *Open J. Astrophys.* **6** (2023) 20 [[arXiv:2305.06347](#)] [[INSPIRE](#)].
- [35] Y. Li et al., *pmwd: A Differentiable Cosmological Particle-Mesh N-body Library*, [arXiv:2211.09958](#) [[INSPIRE](#)].
- [36] X. Li et al., *A differentiable perturbation-based weak lensing shear estimator*, *Mon. Not. Roy. Astron. Soc.* **527** (2023) 10388 [[arXiv:2309.06506](#)] [[INSPIRE](#)].
- [37] D. Stevanovich, A.P. Hearin and D. Nagai, *A differentiable model of the evolution of dark matter halo concentration*, *Mon. Not. Roy. Astron. Soc.* **526** (2023) 1528 [[arXiv:2309.07854](#)] [[INSPIRE](#)].
- [38] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, *Springer Series in Computational Mathematics*, Springer (2010).
- [39] S. Nadkarni-Ghosh and A. Refregier, *The Einstein-Boltzmann equations revisited*, *Mon. Not. Roy. Astron. Soc.* **471** (2017) 2391 [[arXiv:1612.06697](#)] [[INSPIRE](#)].
- [40] J. Lesgourgues and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics*, *JCAP* **09** (2011) 032 [[arXiv:1104.2935](#)] [[INSPIRE](#)].

- [41] A. Kværnø, *Singly diagonally implicit Runge-Kutta methods with an explicit first stage*, *BIT Numer. Math.* **44** (2004) 489.
- [42] C. Howlett, A. Lewis, A. Hall and A. Challinor, *CMB power spectrum parameter degeneracies in the era of precision cosmology*, *JCAP* **04** (2012) 027 [[arXiv:1201.3654](#)] [[INSPIRE](#)].
- [43] S. Seager, D.D. Sasselov and D. Scott, *A new calculation of the recombination epoch*, *Astrophys. J. Lett.* **523** (1999) L1 [[astro-ph/9909275](#)] [[INSPIRE](#)].
- [44] P.J.E. Peebles, *Recombination of the Primeval Plasma*, *Astrophys. J.* **153** (1968) 1 [[INSPIRE](#)].
- [45] J. Chluba and R.M. Thomas, *Towards a complete treatment of the cosmological recombination problem*, *Mon. Not. Roy. Astron. Soc.* **412** (2011) 748 [[arXiv:1010.3631](#)] [[INSPIRE](#)].
- [46] Y. Ali-Haïmoud and C.M. Hirata, *HyRec: A fast and highly accurate primordial hydrogen and helium recombination code*, *Phys. Rev. D* **83** (2011) 043513 [[arXiv:1011.3758](#)] [[INSPIRE](#)].
- [47] N. Lee and Y. Ali-Haïmoud, *HYREC-2: a highly accurate sub-millisecond recombination code*, *Phys. Rev. D* **102** (2020) 083517 [[arXiv:2007.14114](#)] [[INSPIRE](#)].
- [48] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [Erratum *ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [49] V. Desjacques, D. Jeong and F. Schmidt, *Large-Scale Galaxy Bias*, *Phys. Rept.* **733** (2018) 1 [[arXiv:1611.09787](#)] [[INSPIRE](#)].
- [50] N. Kaiser, *Clustering in real space and in redshift space*, *Mon. Not. Roy. Astron. Soc.* **227** (1987) 1 [[INSPIRE](#)].
- [51] S. Cole, K.B. Fisher and D.H. Weinberg, *Fourier analysis of redshift space distortions and the determination of Omega*, *Mon. Not. Roy. Astron. Soc.* **267** (1994) 785 [[astro-ph/9308003](#)] [[INSPIRE](#)].
- [52] A.J.S. Hamilton, *Uncorrelated modes of the nonlinear power spectrum*, *Mon. Not. Roy. Astron. Soc.* **312** (2000) 257 [[astro-ph/9905191](#)] [[INSPIRE](#)].
- [53] J.D. Talman, *Numerical Fourier and Bessel transforms in logarithmic variables*, *J. Comput. Phys.* **29** (1978) 35.
- [54] C. Alcock and B. Paczynski, *An evolution free test for non-zero cosmological constant*, *Nature* **281** (1979) 358 [[INSPIRE](#)].
- [55] 2dFGRS collaboration, *The 2dF Galaxy Redshift Survey: Spherical harmonics analysis of fluctuations in the final catalogue*, *Mon. Not. Roy. Astron. Soc.* **353** (2004) 1201 [[astro-ph/0406513](#)] [[INSPIRE](#)].
- [56] EUCLID collaboration, *Euclid preparation. VII. Forecast validation for Euclid cosmological probes*, *Astron. Astrophys.* **642** (2020) A191 [[arXiv:1910.09273](#)] [[INSPIRE](#)].
- [57] EUCLID collaboration, *Euclid: Validation of the MontePython forecasting tools*, *Astron. Astrophys.* **682** (2024) A90 [[arXiv:2303.09451](#)] [[INSPIRE](#)].
- [58] M. Tegmark, *Measuring cosmological parameters with galaxy surveys*, *Phys. Rev. Lett.* **79** (1997) 3806 [[astro-ph/9706198](#)] [[INSPIRE](#)].
- [59] EUCLID THEORY WORKING GROUP collaboration, *Cosmology and fundamental physics with the Euclid satellite*, *Living Rev. Rel.* **16** (2013) 6 [[arXiv:1206.1225](#)] [[INSPIRE](#)].
- [60] N. Bhandari, C.D. Leonard, M.M. Rau and R. Mandelbaum, *Fisher Matrix Stability*, [arXiv:2101.00298](#) [[INSPIRE](#)].

- [61] J.A. Peacock and S.J. Dodds, *Reconstructing the linear power spectrum of cosmological mass fluctuations*, *Mon. Not. Roy. Astron. Soc.* **267** (1994) 1020 [[astro-ph/9311057](#)] [[INSPIRE](#)].
- [62] J. Chaves-Montero, R.E. Angulo and C. Hernández-Monteagudo, *The effect of photometric redshift uncertainties on galaxy clustering and baryonic acoustic oscillations*, *Mon. Not. Roy. Astron. Soc.* **477** (2018) 3892 [[arXiv:1610.09688](#)] [[INSPIRE](#)].
- [63] H.-J. Seo and D.J. Eisenstein, *Probing dark energy with baryonic acoustic oscillations from future large galaxy redshift surveys*, *Astrophys. J.* **598** (2003) 720 [[astro-ph/0307460](#)] [[INSPIRE](#)].
- [64] EUCLID collaboration, *Euclid preparation. Sensitivity to neutrino parameters*, [arXiv:2405.06047](#) [[INSPIRE](#)].
- [65] M.M. Ivanov, *Effective Field Theory for Large-Scale Structure*, [arXiv:2212.08488](#) [[INSPIRE](#)].
- [66] C. Rampf and O. Hahn, *Shell-crossing in a Λ CDM Universe*, *Mon. Not. Roy. Astron. Soc.* **501** (2021) L71 [[arXiv:2010.12584](#)] [[INSPIRE](#)].
- [67] F. Schmidt, *An n -th order Lagrangian Forward Model for Large-Scale Structure*, *JCAP* **04** (2021) 033 [[arXiv:2012.09837](#)] [[INSPIRE](#)].
- [68] G. Aricò, R.E. Angulo and M. Zennaro, *Accelerating Large-Scale-Structure data analyses by emulating Boltzmann solvers and Lagrangian Perturbation Theory*, [arXiv:2104.14568](#) [[DOI:10.12688/openreseurope.14310.2](#)] [[INSPIRE](#)].
- [69] M. Bonici, L. Biggio, C. Carbone and L. Guzzo, *Fast emulation of two-point angular statistics for photometric galaxy surveys*, [arXiv:2206.14208](#) [[INSPIRE](#)].
- [70] A. Spurio Mancini et al., *CosmoPower: emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys*, *Mon. Not. Roy. Astron. Soc.* **511** (2022) 1771 [[arXiv:2106.03846](#)] [[INSPIRE](#)].
- [71] J. Donald-McCann, F. Beutler, K. Koyama and M. Karamanis, *matryoshka: halo model emulator for the galaxy power spectrum*, *Mon. Not. Roy. Astron. Soc.* **511** (2022) 3768 [[arXiv:2109.15236](#)] [[INSPIRE](#)].
- [72] M. Bonici, F. Bianchini and J. Ruiz-Zapatero, *Capse.jl: efficient and auto-differentiable CMB power spectra emulation*, [arXiv:2307.14339](#) [[DOI:10.21105/astro.2307.14339](#)] [[INSPIRE](#)].
- [73] LSST DARK ENERGY SCIENCE collaboration, *Core Cosmology Library: Precision Cosmological Predictions for LSST*, *Astrophys. J. Suppl.* **242** (2019) 2 [[arXiv:1812.05995](#)] [[INSPIRE](#)].
- [74] G. Ballesteros and J. Lesgourgues, *Dark energy with non-adiabatic sound speed: initial conditions and detectability*, *JCAP* **10** (2010) 014 [[arXiv:1004.5509](#)] [[INSPIRE](#)].
- [75] M. Chevallier and D. Polarski, *Accelerating universes with scaling dark matter*, *Int. J. Mod. Phys. D* **10** (2001) 213 [[gr-qc/0009008](#)] [[INSPIRE](#)].
- [76] E.V. Linder, *Exploring the expansion history of the universe*, *Phys. Rev. Lett.* **90** (2003) 091301 [[astro-ph/0208512](#)] [[INSPIRE](#)].