Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **531,** 1351–1370 (2024) https://doi.org/10.1093/mnras/stae1201
Advance Access publication 2024 May 9

High-accuracy emulators for observables in ΛCDM, $N_{\rm eff}$, Σm_ν, and w cosmologies

Boris Bolliet [1,2★] Alessio Spurio Mancini [3★] J. Colin Hill [4] Mathew Madhavacheril [5] Hidde T. Jense,[6] Erminia Calabrese[6] and Jo Dunkley[7,8]

[1]*DAMTP, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK*
[2]*Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*
[3]*Mullard Space Science Laboratory, University College London, Dorking, RH5 6NT, UK*
[4]*Department of Physics, Columbia University, New York, NY 10027, USA*
[5]*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*
[6]*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, UK*
[7]*Department of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA*
[8]*Department of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544, USA*

Accepted 2024 April 16. Received 2024 March 19; in original form 2023 August 14

ABSTRACT

We use the emulation framework `CosmoPower` to construct and publicly release neural network emulators of cosmological observables, including the cosmic microwave background (CMB) temperature and polarization power spectra, matter power spectrum, distance-redshift relation, baryon acoustic oscillation (BAO) and redshift-space distortion (RSD) observables, and derived parameters. We train our emulators on Einstein–Boltzmann calculations obtained with high-precision numerical convergence settings, for a wide range of cosmological models including ΛCDM, wCDM, ΛCDM + $N_{\rm eff}$, and ΛCDM + Σm_ν. Our CMB emulators are accurate to better than 0.5 per cent out to $\ell = 10^4$, which is sufficient for Stage-IV data analysis, and our $P(k)$ emulators reach the same accuracy level out to $k = 50$ Mpc^{-1}, which is sufficient for Stage-III data analysis. We release the emulators via an online repository (CosmoPower Organisation), which will be continually updated with additional extended cosmological models. Our emulators accelerate cosmological data analysis by orders of magnitude, enabling cosmological parameter extraction analyses, using current survey data, to be performed on a laptop. We validate our emulators by comparing them to `class` and `camb` and by reproducing cosmological parameter constraints derived from *Planck* TT, TE, EE, and CMB lensing data, as well as from the Atacama Cosmology Telescope Data Release 4 CMB data, Dark Energy Survey Year-1 galaxy lensing and clustering data, and Baryon Oscillation Spectroscopic Survey Data Release 12 BAO and RSD data.

Key words: methods: data analysis – methods: statistical – cosmic background radiation – large-scale structure of Universe.

1 INTRODUCTION

Next-generation cosmic microwave background (CMB) surveys, such as the Simons Observatory (SO; Simons Observatory 2019), CMB-S4 (Abazajian et al. 2019; CMBS4 2022), and CMB-HD (Sehgal et al. 2019), will provide high-signal-to-noise measurements of the CMB temperature and polarization anisotropy down to small scales, characterized by their power spectra. The final seasons of the Atacama Cosmology Telescope (ACT; Henderson et al. 2016) and of the South Pole Telescope (SPT; Benson et al. 2014) will also make progress in that direction in the forthcoming years. In addition, current and next-generation galaxy surveys, including the Dark Energy Survey (DES; Dark Energy Survey Collaboration 2016), the Kilo-Degree Survey (KiDS; Kuijken et al. 2015), the Hyper Suprime-Cam Survey (HSC; Aihara et al. 2018), the Vera C. Rubin Observatory (Ivezić et al. 2019), *Euclid* (Laureijs et al. 2011), *Roman* (Spergel et al. 2015; Akeson et al. 2019), or *SPHEREx* (Doré et al. 2014), probe the matter power spectrum via weak gravitational lensing and galaxy clustering and spectroscopic galaxy surveys such as the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013) and the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016) also probe the baryon acoustic oscillation (BAO) feature and cosmological velocities via redshift-space distortions (RSD).

Einstein–Boltzmann codes such as `camb` (Lewis, Challinor & Lasenby 2000) and `class` (Blas, Lesgourgues & Tram 2011; Lesgourgues 2011a) are routinely used to accurately compute linear-theory cosmological power spectra, as well as the background cosmic evolution (e.g. the distance-redshift relation). These codes also implement various prescriptions for the calculation of non-linear corrections to the matter power spectrum and quantities derived therefrom, e.g. Halofit (Smith et al. 2003) and HMCode (Mead et al. 2015, 2016, 2021). The computation of these cosmological quantities through a Boltzmann code represents a significant computational requirement in Markov Chain Monte Carlo (MCMC) methods, traditionally used to extract constraints on cosmological

★ E-mail: boris.bolliet@gmail.com(BB); a.spuriomancini@ucl.ac.uk(AS)

Table 1. Parameter ranges used to generate the LHC of cosmological parameters used to compute the training data (see Section 2 for details). Our emulators should never be used outside these ranges. For the matter power spectrum emulators, the LHC is supplemented by $z_{\rm pk}$, the redshift at which the power spectrum is evaluated, which is varied between 0 and 5. For n_s, the values in parentheses refer to the prior bounds that are used in the emulators for extensions, which are slightly broader than for the ΛCDM emulators (outside the parentheses).

Parameter	Min. Value	Max. Value
$\ln 10^{10} A_{\rm s}$	2.5	3.5
$\Omega_{\rm cdm} h^2$	0.08	0.20
$\Omega_{\rm b} h^2$	0.01933	0.02533
H_0 [km s^{-1} Mpc^{-1}]	39.99	100.01
$n_{\rm s}$	0.8812(0.8)	1.0492(1.2)
τ	0.02	0.12
m_ν (eV)	0.	2.
w	−2.	−0.33
$N_{\rm eff}$	1.5	5.5

parameters. MCMC methods require $\sim 10^4$–10^6 calls to a Boltzmann code as they achieve convergence through a similar number of likelihood evaluations. The problem is exacerbated if one requires high numerical accuracy in the computation of the cosmological quantities obtained from Boltzmann codes, as needed for upcoming surveys (e.g. McCarthy, Hill & Madhavacheril 2022).

Several groups have developed emulators of cosmological quantities in order to accelerate the MCMC calculations by bypassing the call to a Boltzmann code with a faster algorithm (e.g. Auld et al. 2007; Fendt & Wandelt 2007; Albers et al. 2019; Aricò, Angulo & Zennaro 2021; Günther et al. 2022; Mootoovaloo et al. 2022; Bonici, Bianchini & Ruiz-Zapatero 2023). Recently, Spurio Mancini et al. (2022) developed `CosmoPower`, an emulation framework for cosmological quantities based on TensorFlow neural networks (Abadi et al. 2015) and with an architecture similar to `Speculator` (Alsing et al. 2020). The accuracy of the `CosmoPower` emulators presented in Spurio Mancini et al. (2022) was tested on CMB and large-scale structure (LSS) analysis of current and future data. More recently, Balkenhol et al. (2022) constructed `CosmoPower` emulators for the analysis of the SPT data, covering angular scales up to $\ell = 3000$, trained on high-precision `camb` spectra and covering LCDM, LCDM+Neff, and LCDM + Alens models.[1] These were subsequently used in the `candl` package (Balkenhol et al. 2024).

Here we use `CosmoPower` to provide high-accuracy emulators of CMB temperature (emulator acronym: TT), polarization (TE and EE), and lensing potential power spectra (PP), up to $\ell_{\rm max} = 10^4$, as well as linear and non-linear matter power spectra (PKL and PKNL) up to $k = 50\,{\rm Mpc}^{-1}$ in the redshift range $z \in [0, 5]$, the Hubble parameter (H), angular diameter distance (DA), $\sigma_8(z)$ (S8) for $z \in [0, 20]$, and various derived parameters (DER; see Section 2 for a complete list).

With this suite of emulators, we are able to accurately reproduce marginalized posterior distributions from current CMB and LSS likelihood codes (based on CMB, matter power spectra, and BAO distances) within minutes on a laptop, whereas the original MCMC analyses took hours to days on large computing clusters. To accomplish this, we use the `cobaya` MCMC sampler and its implemented likelihoods (Torrado & Lewis 2021). Our `CosmoPower` emulators are wrapped into `cobaya` via a simple wrapper that will be available online.[2] We reproduce parameter constraints from the ACT Data Release 4 (Aiola et al. 2020; Choi et al. 2020), *Planck* 2018 temperature, polarization, and lensing potential data (Planck Collaboration 2020a, b), DES-Y1 (Krause et al. 2017; DES Collaboration 2018; Troxel et al. 2018), and BAO distances from BOSS and other surveys (Beutler et al. 2011; Ross et al. 2015; Alam et al. 2017).

Our emulators are explicitly constructed with sufficient accuracy to remain valid for use in upcoming CMB analyses for the foreseeable future, including ACT, SPT, SO, CMB-S4, and CMB-HD, as well as ongoing galaxy surveys.

In addition to the emulators, we release the full pipeline to generate them, which can be used by the community to build emulators for other extensions to the standard cosmological model. The cosmological models covered here include Λ Cold Dark Matter (CDM), wCDM, where w is the dark energy equation of state; ΛCDM with varying $N_{\rm eff}$, where $N_{\rm eff}$ is the effective number of relativistic species at recombination; and ΛCDM with varying neutrino mass Σm_ν. In ΛCDM, ΛCDM + $N_{\rm eff}$, and wCDM, we assume one massive and two massless neutrino states, as in the baseline *Planck* analyses (Planck Collaboration 2020a). In ΛCDM + Σm_ν, we generate emulators for both the case of three species of neutrinos with degenerate mass, and the case of one massive and two massless neutrino states.

This paper is organized as follows. In Section 2, we outline our methodology, while in Section 3, we study and calibrate accuracy settings of Boltzmann solvers by comparing `camb` and `class`. This sets the accuracy target for our emulators, which we validate in Section 4. In Section 5, we illustrate the validity and applicability of our emulators on a variety of likelihoods mentioned above. We conclude in Section 6.

2 METHOD

The general methodology used to construct training and testing sets for our neural network emulators follows that of Spurio Mancini et al. (2022). Here we provide a brief summary and refer the reader to that paper for further details. To compute all of the cosmological quantities on which our emulators are trained and tested we use `class` v2.9.4.[3] We proceed in six steps, as follows.

(i) We specify the cosmological parameters upon which the emulators are built, i.e. the inputs for our neural networks. We choose the following six ΛCDM parameters, as defined in `class`:

(a) ω_b, the physical baryon density;
(b) $\omega_{\rm cdm}$, the physical dark matter density;
(c) H_0, the Hubble parameter today;
(d) τ, the reionization optical depth;
(e) $n_{\rm s}$, the scalar spectral index;
(f) $\ln 10^{10} A_{\rm s}$; where $A_{\rm s}$ is the amplitude of primordial scalar perturbations.

[1] github/.../cosmopower/trained_models/SPT_high_accuracy

[2] The wrapper will be shared in a forthcoming paper (Jense et al. 2024), which will cover in detail how to embed these emulators in future experiments' likelihood pipelines.

[3] We initially attempted to use the latest `class` version, i.e. v3; however, it appears that the numerical accuracy of the code has changed compared to v2.9.4, and it was not clear how to recover sufficiently high accuracy for our purposes at high ℓ (see https://github.com/lesgourg/class_public/issues/494). Thus, here we opted for v2.9.4. We note that the updates made in v3 compared to v2.9.4 are detailed on the repository webpage and that after submission of this manuscript `class` v3 was updated to adress these issues (see footnote 24).

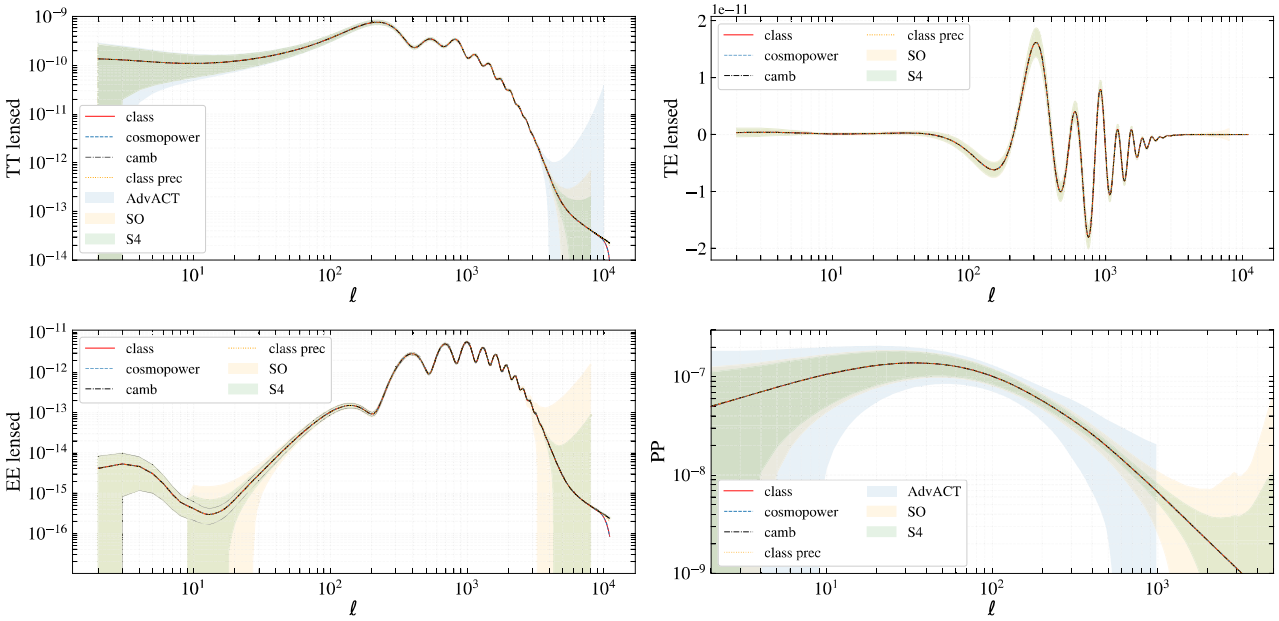

Figure 1. Comparison of lensed CMB power spectra computed with `class` (for default precision as the solid red line and ultra-high precision settings as the dotted yellow line labeled 'class-prec'), `CosmoPower`, and `camb`. We show the dimensionless $\ell(\ell+1)C_\ell/2\pi$ for TT, TE, and EE, and $[\ell(\ell+1)]^2C_\ell^{\phi\phi}/2\pi$ for PP (lensing convergence power spectrum) in a ΛCDM model with one set of cosmological parameters from the *Planck* 2018 results (see text for details). The shaded areas indicate the forecast 1σ uncertainty for Advanced ACT, SO, and CMB-S4 (see Section 4 for details). This figure illustrates the emulated quantities, the relative difference between the spectra are shown in Fig. 4 (in per cent) and 5 (in units of CMB-S4 statistical error bars). On the EE plot, the thin black solid lines indicate the cosmic variance.

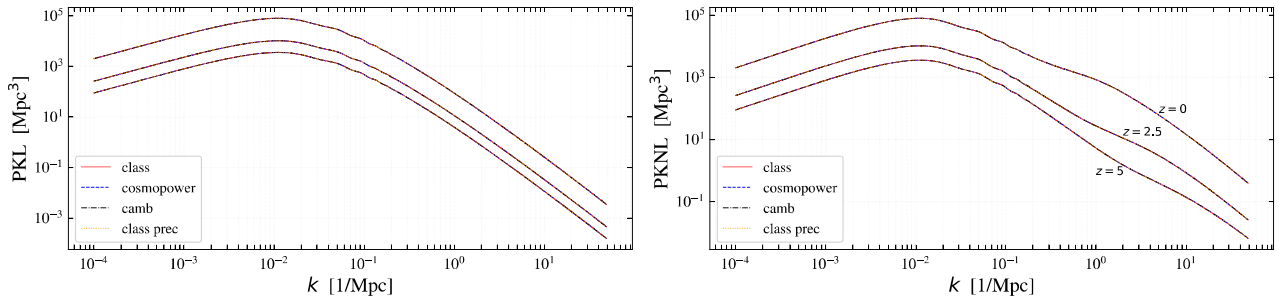

Figure 2. Same settings as Fig. 1 but for matter power spectra. We show the linear matter power spectra on the left and the non-linear spectra on the right (computed according to `HMcode` with $c_{\rm min} = 3.13$ and $\eta_0 = 0.603$, within `class`). In both panels, we show the power at three redshifts, $z = 0$, 2.5, and 5 from top to bottom, spanning the range of the emulators. This figure illustrates the emulated quantities, the relative difference between these spectra are shown in percentages in Fig. 6.

In extensions, we also add w, $N_{\rm eff}$, or Σm_ν, defined in the previous section.

Note that we opt for H_0 rather than for the angular size of the sound horizon at recombination, $100\theta_s$ (which we obtain as a derived parameter), due to the slightly different definitions of the angular acoustic scale implemented in `camb` and `class`.[4] For MCMC analyses, if one needs to sample over $100\theta_s$ rather than H_0, it is straightforward to do so, e.g. using a shooting method with the derived parameter emulators.[5]

(ii) We specify the range of the parameter values over which we train our emulators, reported in Table 1. These ranges should be conservative enough for all physically relevant analyses, assuming the use of data with similar or better constraining power than *Planck* for the CMB, and DES or BOSS for galaxies. We emphasize that the emulators should not be used outside the ranges of Table 1 or for different cosmological models, as there is no guarantee on their accuracy in such cases. For the ΛCDM + $N_{\rm eff}$ model, we let $N_{\rm eff}$ vary between 1.5 and 5.5. In practice, we vary the parameter $N_{\rm ur}$ in `class` between 0.49 and 4.49 and obtain $N_{\rm eff}$ as a derived parameter.[6] In order to generate the data necessary for the matter power spectrum

[4]In `class`, θ_s is defined at the maximum of the visibility function, while in `camb` it is defined when the optical depth equals unity, and these times are slightly different. The θ_s of our emulator is the same as that in `class`.

[5]This adds an extra 20 ms to each MCMC step and would therefore increase slightly the convergence time of the chains. Note that the shooting method is also the standard way to sample over $100\theta_s$ in `class`.

[6]$N_{\rm ur} = N_{\rm eff} - (T_{\rm ncdm}/(4/11)^{1/3})^4$, where $T_{\rm ncdm}$ is in units of $T_{\rm cmb}$ today. This is the convention used by `class`.

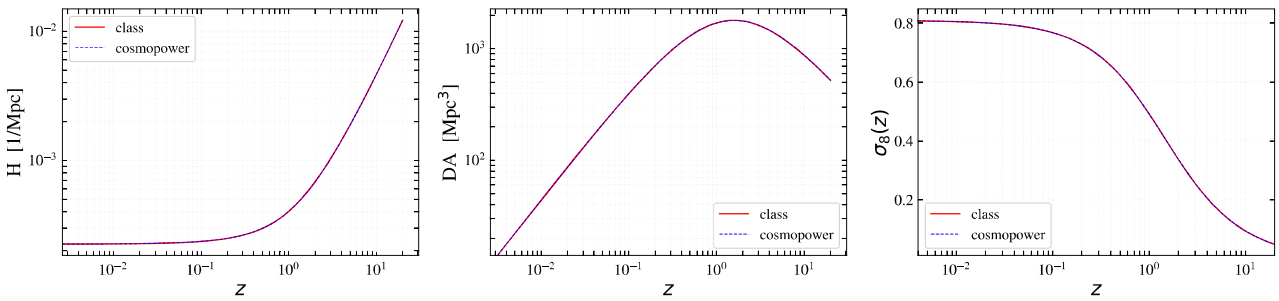

Figure 3. Redshift evolution of the Hubble parameter (left), angular diameter distance (middle), and σ_8 (right) as computed with `class` and `CosmoPower`, between $z = 0$ and 20 (i.e. spanning the redshift range of our emulators). We use the same settings as Fig. 1 and note that in this parameters configuration the relative difference between `class` and `CosmoPower` (not shown here) is less than 0.2 per cent. This figure illustrates the emulated quantities, see Fig. 4 for the relative difference over the full testing set.

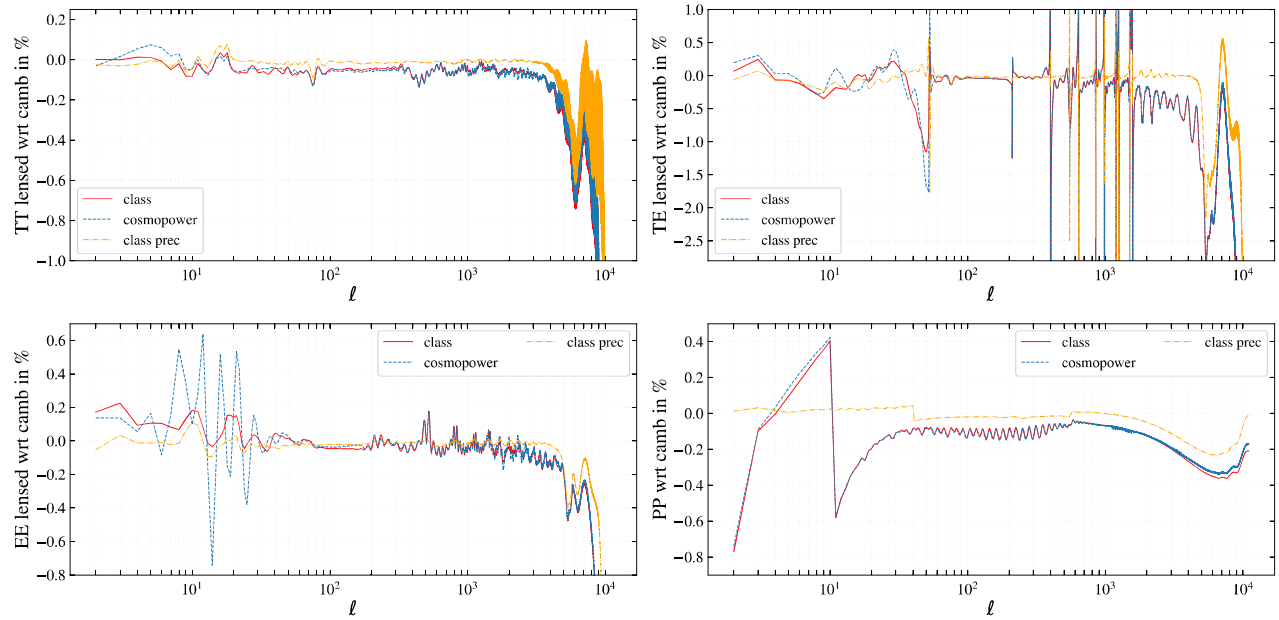

Figure 4. Relative difference of CMB power spectra (top left: TT, top right: TE, bottom left: EE, bottom right: lensing potential, PP) between two different settings of `class` (high precision as the solid red line and ultra-high precision settings as the dotted yellow line labelled 'class-prec'), and `CosmoPower` with respect to `camb` in per cent. We use the same parameters settings as in Fig. 1. Note that the non-negligible fractional deviations seen in a few places for TE occur due to the zero-crossing of the spectrum at those multipoles, and do not reflect inaccuracy in the emulator.

at different redshifts, we add $z_{\rm pk}$, the redshift at which we output the matter power spectra, as an extra varying input parameter to our grid.

(iii) We generate a Latin hypercube (LHC) of the parameter space. We set $N_{\rm S} = 128\,000$ as the number of samples in the LHC,[7] which is the same as the number of spectra we aim at computing. To generate the LHC, we use the python library `pyDOE`.[8]

(iv) We use a computing cluster to calculate the $N_{\rm S}$ training samples in parallel. Note that despite the total number of training samples being similar to that needed for an MCMC, the task of generating training samples is embarrassingly parallel, since the computation for a given set of input parameters does not depend on any other – unlike an MCMC, which is a path-dependent calculation in parameter space. Given this, we find that computing resources are more optimally used by running each computation on a single thread (i.e. setting `OMP_NUM_THREADS = 1` before running `class`). The total amount of disc space taken by the training and testing samples is ≈150 GB. Note that for each sample, in addition to the cosmological observables that we want to emulate, we also save 14 derived parameters, such as σ_8 and $100\theta_{\rm s}$ (see the end of this section for details). For each class of emulators (ΛCDM + extensions) the computation of the data requires ≈48 h on a modern high-performance computing cluster. (The accuracy settings for the Boltzmann codes are given in Section 3.)

(v) We process the data so that the training samples can be used in the `CosmoPower` training pipeline. For all quantities (TT, TE, EE, PP, PKL, PKNL, DA, H, S8, and DER), we select 80 per cent of the samples for training and leave the remaining ones for testing. The TT, EE, PP, PKL, PKNL, H, and DER data consists of positive numbers with large dynamic ranges; to ease training, we take the logarithm

[7]This is roughly the same amount of samples as in a standard cosmological MCMC analysis.
[8]https://pythonhosted.org/pyDOE/

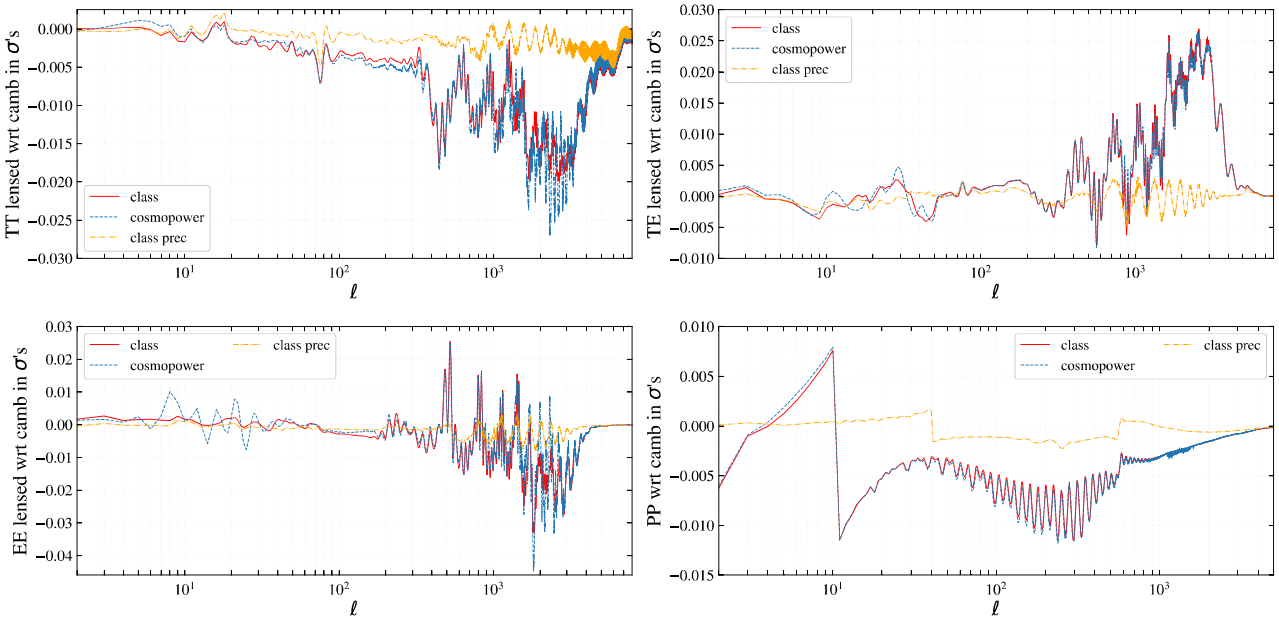

Figure 5. Relative difference between two different settings of `class` (high precision as the solid red line and ultra-high precision settings as the dotted–dashed yellow line labeled 'class-prec') and `CosmoPower` with respect to `camb` in units of the forecast CMB-S4 statistical error bars. Same settings as Fig. 1.

of all these quantities before passing them to `CosmoPower`. The TE power spectra are oscillatory, zero-crossing functions, and thus we cannot directly take the logarithm. In this case, we use Principal Component Analysis (PCA), which reduces the dimensionality of the data set and its dynamic range. We retain only 64 principal components for each spectrum. We verify that adding more does not lead to any significant improvement (Spurio Mancini et al. 2022).

(vi) We generate the emulators using the `CosmoPower` functions and verify their accuracy on the test set. Generating each emulator involves training a neural network model. This operation takes $\mathcal{O}(1\,\mathrm{hr})$ for each emulator and is best performed on a GPU, to fully take advantage of the acceleration provided by the `TensorFlow` library, which `CosmoPower` uses for its neural network implementations. We also stress that for a given cosmological model once the training is done and the emulator generated, this step does not need to be repeated. Note that we emulate TT, TE, EE, PP, PKL, PKNL, DA, H, S8, and DER separately. For all emulators, we use dense neural networks with four hidden layers, each with 512 nodes. When building the emulators for TT, TE, EE, PP, DA, H, S8, and DER, we do not include $z_{\rm pk}$ in the mapping from parameters to data, since $z_{\rm pk}$ has no effect on these quantities. Similarly, we also remove τ from the parameter sets for the PP, DA, H, S8, PKL, and PKNL emulators. The trained TT, TE, EE, PP, PKL, PKNL, DA, H, S8, and DER emulators are stored as `pickle` files. The size of each CMB power spectra emulator file is 25.9 MB, except the TE emulator which is lighter because of the PCA and is 6.3 MB. The PKL and PKNL emulators are 4.2 MB; the H, DA, and S8 emulator files are 13.5 MB; and the DER emulator is 3.2 MB. (The size of each emulator file is proportional to the number of data points that are saved: 11 000 multipoles for the CMB TT and EE spectra, the number of PCA weights for the TE spectra, 500 wavenumbers for PKL and PKNL, and 5000 redshifts for DA, H, and S8.)

For illustration, examples of CMB power spectra are shown in Fig. 1; examples of matter power spectra calculations are shown in Fig. 2; and examples of calculations of $H(z)$, $D_A(z)$, and $\sigma_8(z)$, which are used in the BAO and RSD calculations, are shown in Fig. 3.

2.1 Matter power spectrum

To construct the emulators for the linear power spectrum and its non-linear corrections, we add an extra input to the neural networks, namely, the redshift at which the power spectra are computed, $z_{\rm pk}$. This represents an additional parameter, sampled between 0 and 5. For instance, the LHC used for our ΛCDM emulator has seven dimensions: the six ΛCDM parameters (cf. Table 1) augmented by $z_{\rm pk}$. For each LHC sample we save the matter power spectra at $z_{\rm pk}$ on a logarithmically spaced k-grid between $k_{\rm min} = 10^{-4}$ Mpc^{-1} and $k_{\rm max} = 50$ Mpc^{-1} with 500 points. (See Section 3 for our perturbation and non-linear settings.)

2.2 BAO and redshift-space distortions

Recent galaxy survey data allow us to constrain models based on various BAO distance measures, as well as using RSDs to constrain the parameter combination $f\sigma_8$, where f is the growth rate of cosmic structures. To compute BAO distances and $f\sigma_8$, we need to save the angular diameter distance D_A, the Hubble parameter H, and σ_8 as a function of z, as well as the comoving sound horizon at baryon drag r_d, which is last in the list of derived parameters discussed above. We save $D_A(z)$, $H(z)$, and $\sigma_8(z)$ on a linearly spaced z-grid between $z_{\rm min} = 0$ and $z_{\rm max} = 20$ with 5000 points. The upper bound $z_{\rm max} = 20$ is much higher than what is relevant to galaxy surveys; however we record the high-z distances as they may be useful to other applications, such as studies of reionization or cosmic-dawn 21 cm measurements. With these quantities, we compute BAO distances and $f\sigma_8$ straightforwardly (see e.g. Alam et al. 2017). Note that $f\sigma_8 = -(1$

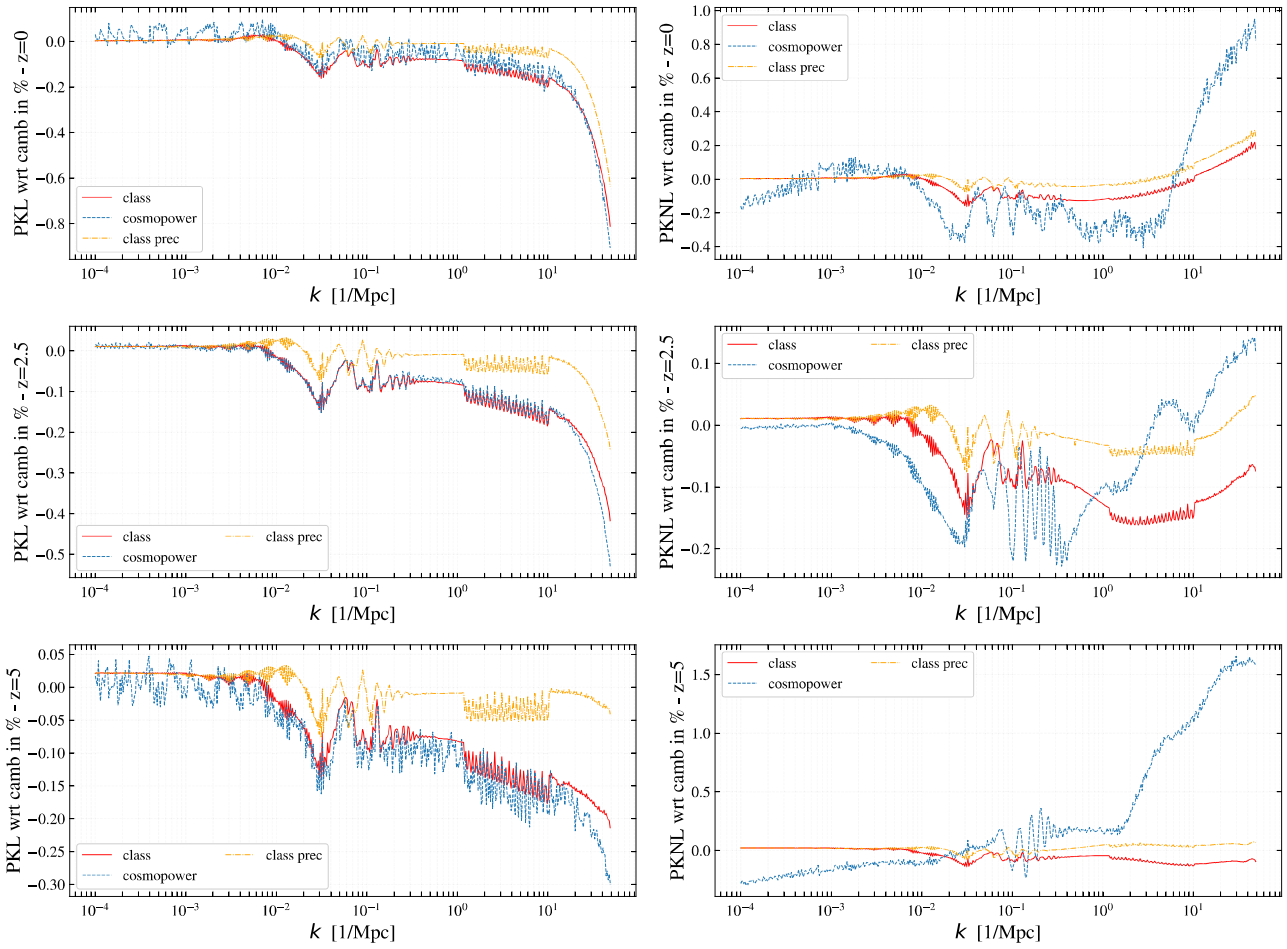

Figure 6. Relative difference of the linear (left) and non-linear (right) matter power spectra between two different settings of `class` (high precision as the solid red line and ultra-high precision settings as the dotted–dashed yellow line labeled 'class-prec') and `CosmoPower` with respect to `camb` in per cent.

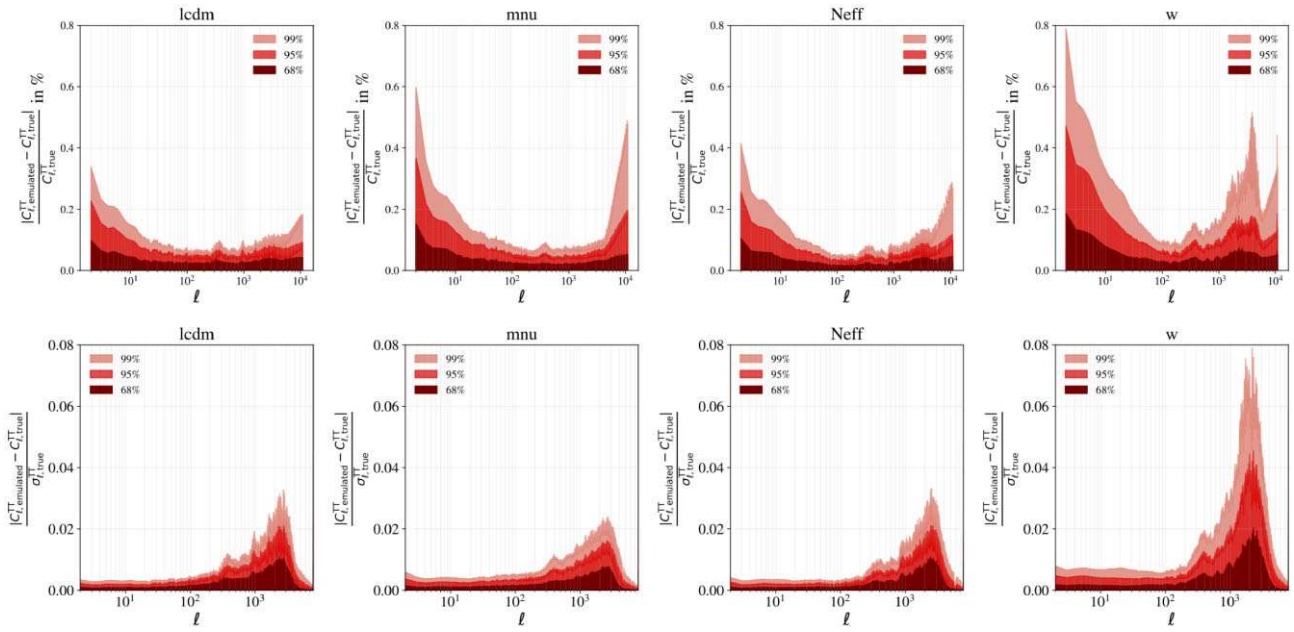

Figure 7. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for CMB TT power spectra in per cent (top) and in units of CMB-S4 sensitivity (bottom). See end of Section 4 for details.

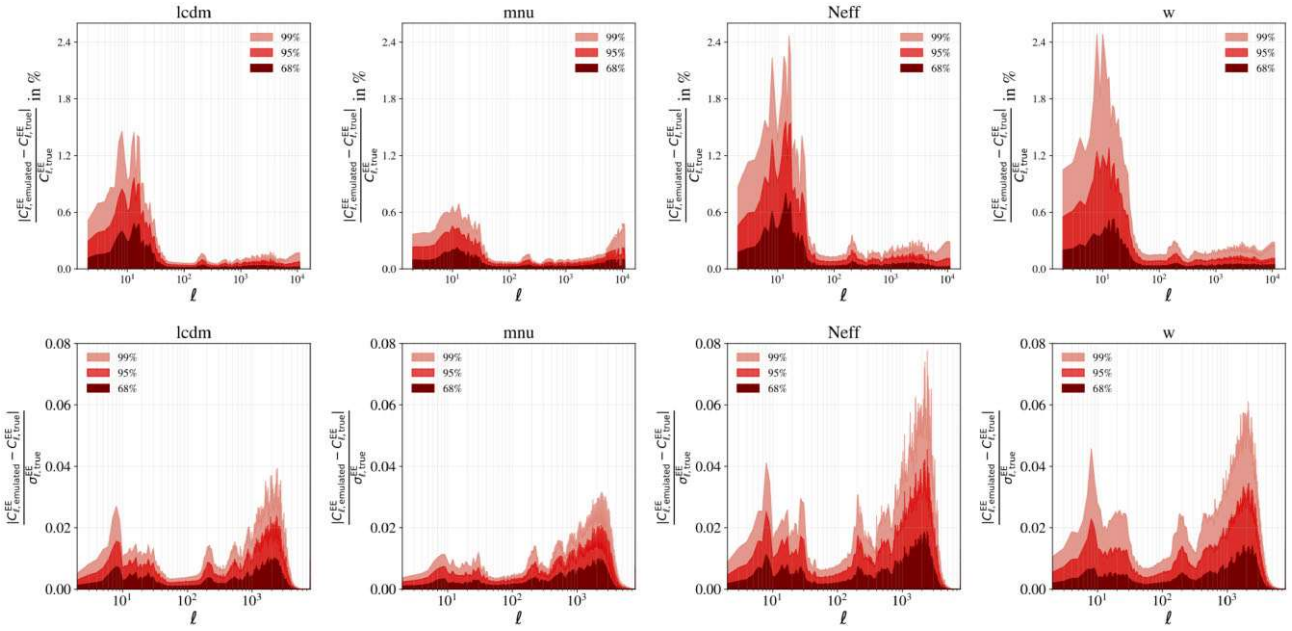

Figure 8. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for CMB EE power spectra in per cent (top) and in units of CMB-S4 sensitivity (bottom). See end of Section 4 for details.

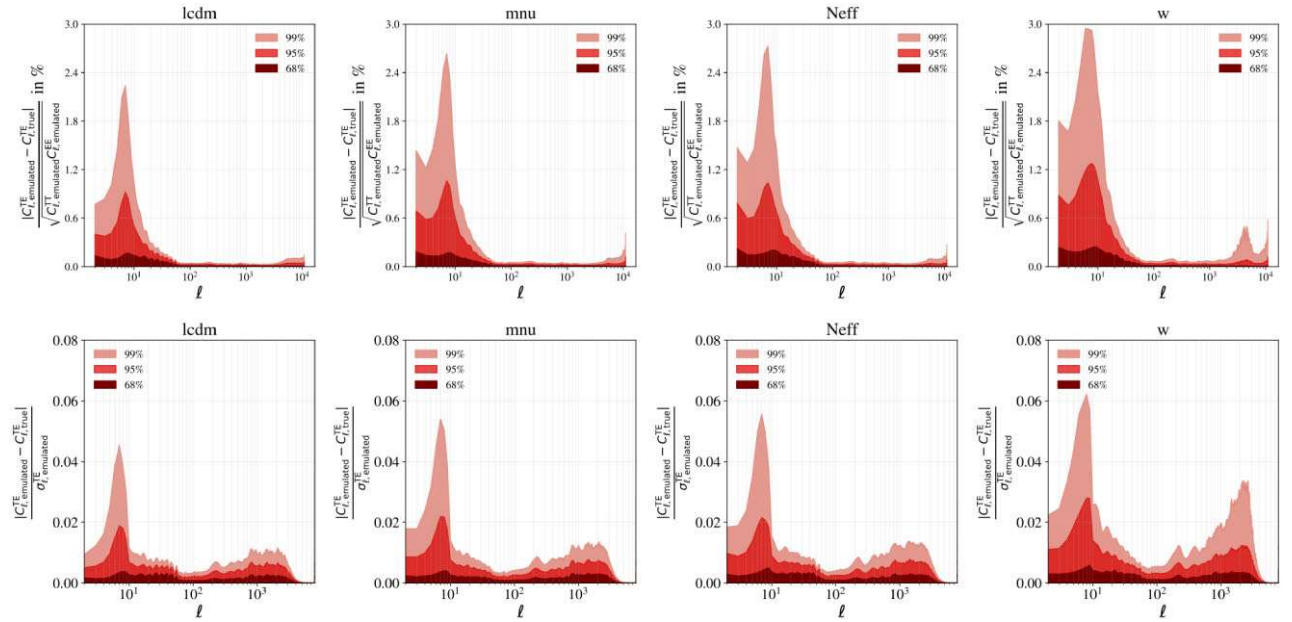

Figure 9. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for CMB TE power spectra in per cent (top) and in units of CMB-S4 sensitivity (bottom). Note that the spikes are simply due to zero crossings of the TE power spectrum. See end of Section 4 for details.

$+\ z)\mathrm{d}\sigma_8(z)/\mathrm{d}z$, where the derivative can be evaluated numerically. This is how we compute $f\sigma_8$, as is done in `class`.

2.3 Derived parameters

We save 14 commonly used derived parameters, namely:

(i) the angular size of the sound horizon at decoupling $100\theta_s$,
(ii) the amplitude of matter clustering σ_8,
(iii) the primordial Helium fraction Y_P,
(iv) the reionization redshift $z_{\rm reio}$,
(v) the number of effective relativistic degrees of freedom in the early Universe $N_{\rm eff}$,
(vi) the conformal time at which the visibility function reaches its maximum (i.e. the recombination time) $\tau_{\rm rec}$,
(vii) its associated redshift $z_{\rm rec}$,
(viii) the comoving sound horizon at recombination $r_{\rm s,\,rec}$,

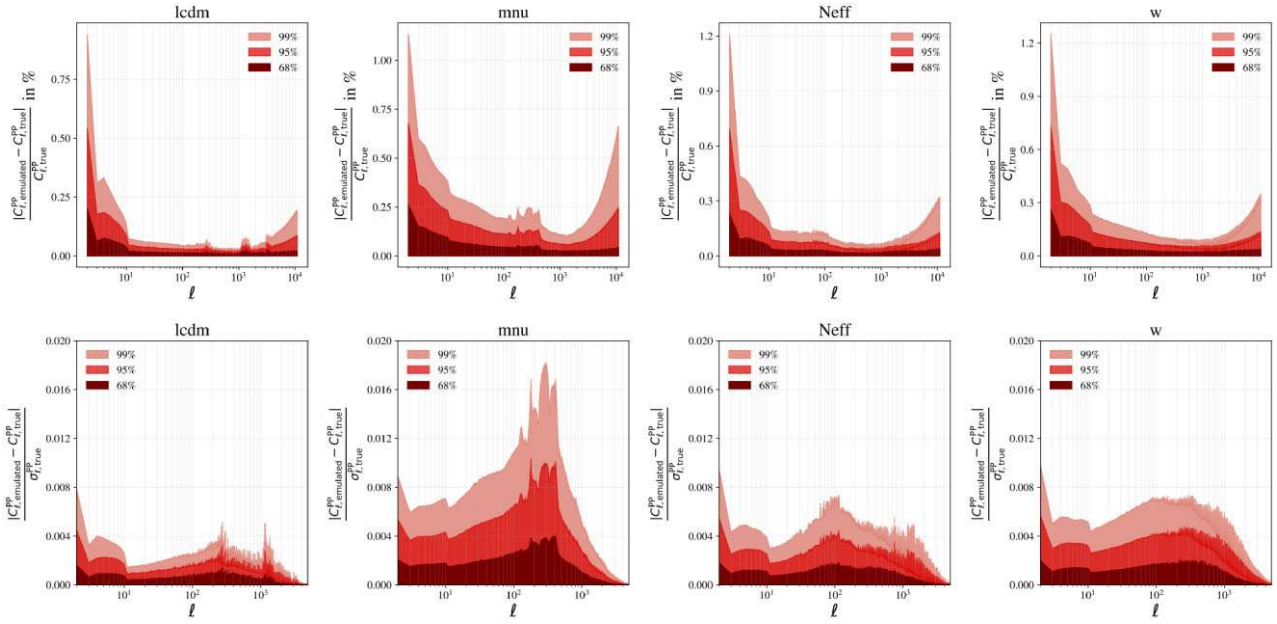

Figure 10. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for CMB lensing convergence power spectra in per cent (top) and in units of CMB-S4 sensitivity (bottom). See end of Section 4 for details.

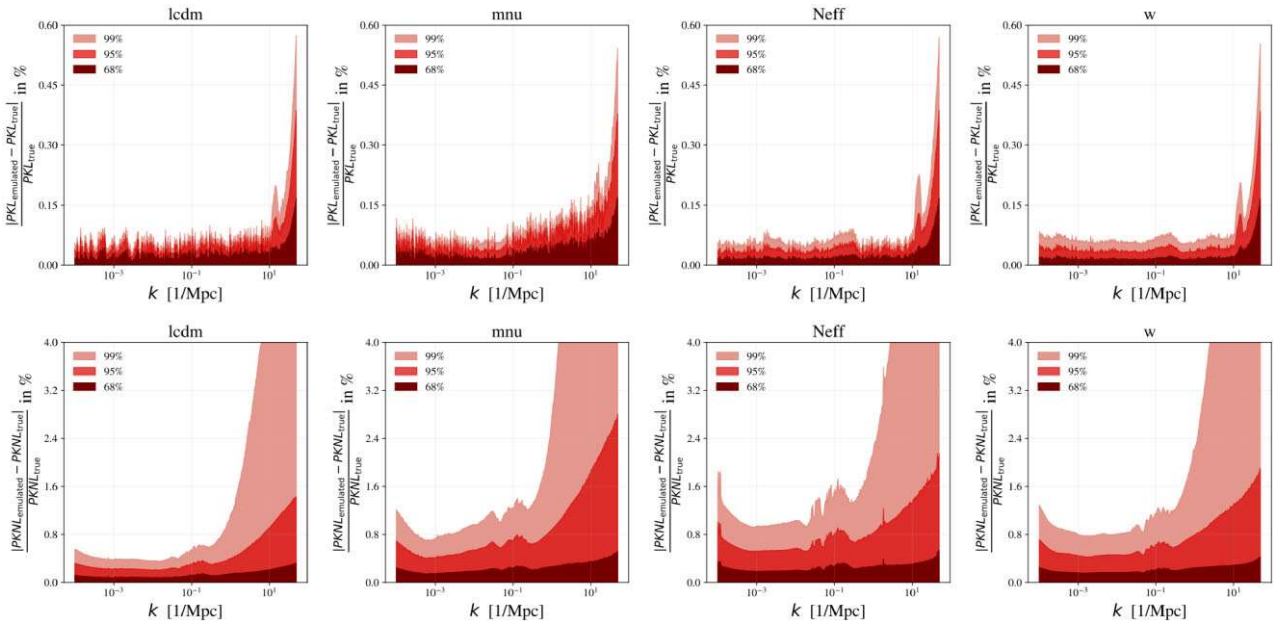

Figure 11. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for the linear (top) and non-linear (bottom) matter power spectra in per cent. The `class` predictions used here are the test data set, made of ≈25 000 spectra. See end of Section 4 for details.

(ix) the conformal angular diameter distance to recombination $r_{\mathrm{a,\,rec}}$,
(x) the conformal time at which the photon optical depth crosses unity $\tau_\star$,
(xi) the redshift $z_\star$ at which the photon optical depth crosses unity,
(xii) its associated comoving sound horizon $r_{\mathrm{s},\star}$,
(xiii) and conformal angular diameter distance $r_{\mathrm{a},\star}$
(xiv) and the comoving sound horizon at baryon drag r_d.

These parameters are saved into a list for each sample in the LHC, with the ordering given above.

3 BOLTZMANN CODE ACCURACY AND SETTINGS

To ensure high precision in our calculations, we carefully study and calibrate the numerical precision parameters in `camb` and `class`.[9] For the maximum multipole, we choose $\ell_{\mathrm{max}} = 11\,000$ in both codes.

[9]See Lesgourgues (2011b) for calibration of the default numerical precision settings of `class`, which are determined to be sufficient for a *Planck*-like experiment, but require to be adjusted for our purpose, as explained in this section.

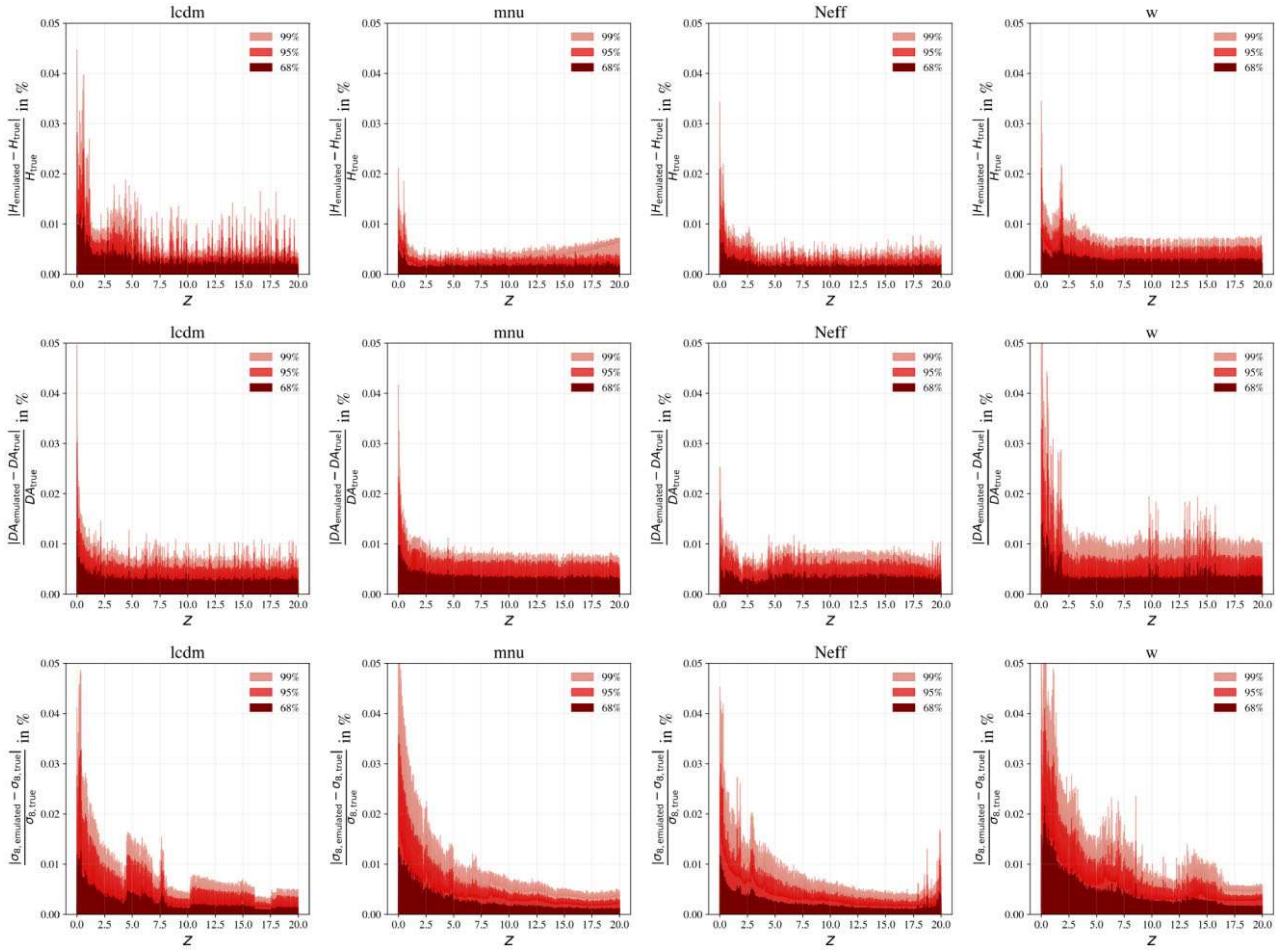

Figure 12. Relative difference between `CosmoPower` and the 'true' high-precision `class` prediction for the redshift evolution of the Hubble parameter (top row), angular diameter distance (middle row), and σ_8 (bottom row) in per cent. See end of Section 4 for details.

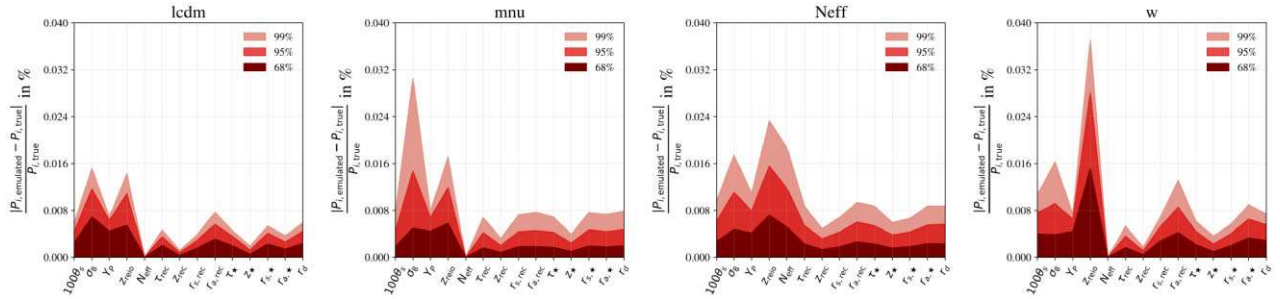

Figure 13. Relative difference between `CosmoPower` emulators and the 'true' high-precision `class` prediction for the derived parameters in per cent. Each unit on the x-axis represents a derived parameter as indicated (see page 3 for the definitions of the fourteen parameters). Note that for $N_{\rm eff}$ the error is null since this parameter is fixed, except in the ΛCDM + $N_{\rm eff}$ model (third panel) where it is varied. In ΛCDM + $N_{\rm eff}$, $N_{\rm eff}$ is not an input parameter but a derived parameter (see Section 2) and hence there is a small error.

We chose $\ell_{\rm max} = 11\,000$ as some forecasts and calculations, e.g. in the context of CMB-HD (Sehgal et al. 2019), require CMB power spectra at such high multipoles. For non-linear corrections to the matter power spectrum, we use `HMcode` (Mead et al. 2016) with fiducial values $c_{\rm min} = 3.13$ and $\eta_0 = 0.603$ in both codes. These values correspond to a dark matter-only scenario and are the same as in the baseline *Planck* 2018 analyses (Planck Collaboration 2020a), justifying our choice. For reionization modelling, we kept the default `class` setting (i.e. the same as in `camb`).

We remark here that in future work it will be useful to train emulators in which these `HMcode` parameters are also varied, or for other non-linear regime prescriptions. Indeed, as shown by, e.g. McCarthy et al. (2022), future CMB power spectrum data will be sensitive to the impact of baryonic physics on the matter power spectrum (via CMB lensing); in addition, LSS statistics are sensitive to the non-linear and baryonic prescriptions at high values of k. With emulators that include these parameters, we will be able to marginalize over this uncertainty.

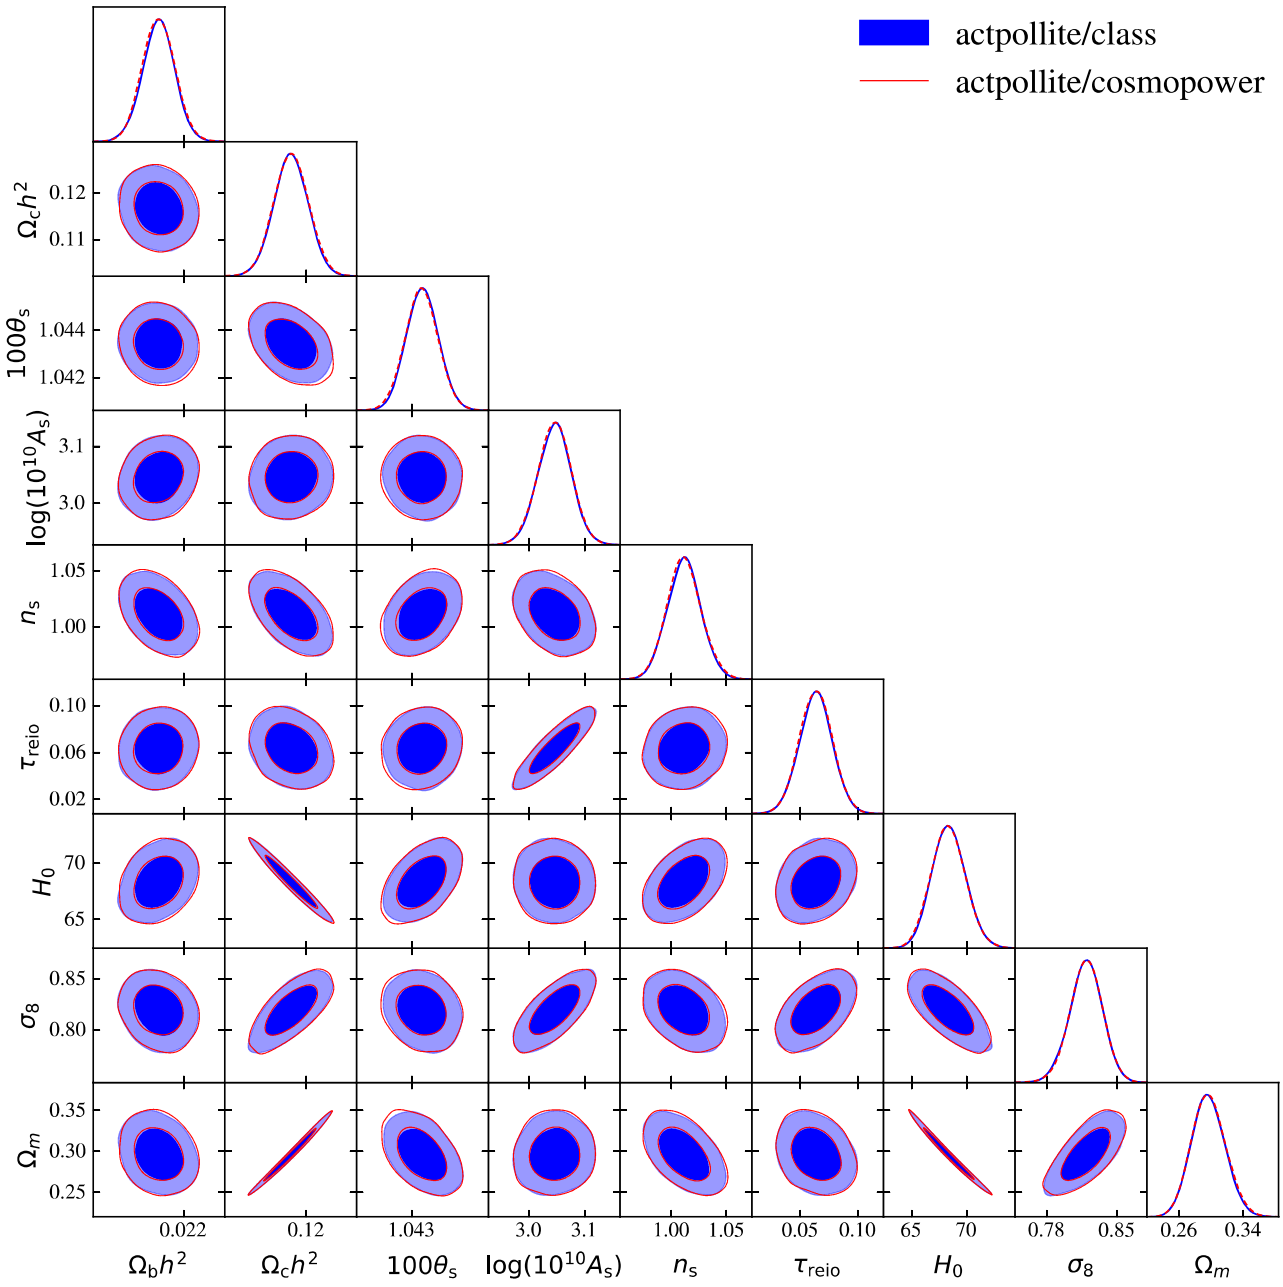

Figure 14. Comparison of 2D marginalized posterior probability distributions for ACT DR4 cosmological parameters in ΛCDM, between `CosmoPower` (empty red contours) and the full `class` calculation (solid blue contours, taken from Hill et al. (2022)). With `CosmoPower`, the derived parameters H_0, σ_8, and Ω_m are added in post-processing of the chains. These results are obtained with the 'actpol_lite' ACT DR4 likelihood (see footnote 16) and Section 5.1 for details.

The TT, TE, EE, and PP spectra computed with `camb`[10] are fully converged when using the following parameters (see e.g. Hill et al. 2022; McCarthy et al. 2022):

`set_matter_power`:

(i) `kmax = 10`
(ii) `k_per_logint = 130`
(iii) `nonlinear = True`

`set_for_lmax`:

(i) `lens_potential_accuracy = 8`
(ii) `lens_margin = 2050`

`set_accuracy`:

(i) `AccuracyBoost = 2.0`
(ii) `lSampleBoost = 2.0`
(iii) `lAccuracyBoost = 2.0`
(iv) `DoLateRadTruncation = False`

We use the `camb` spectra computed with these settings as a high-accuracy reference. This high-accuracy `camb` calculation takes $\mathcal{O}$(1 min) per sample on 16 threads.

[10]We use `camb` v1.3.6.

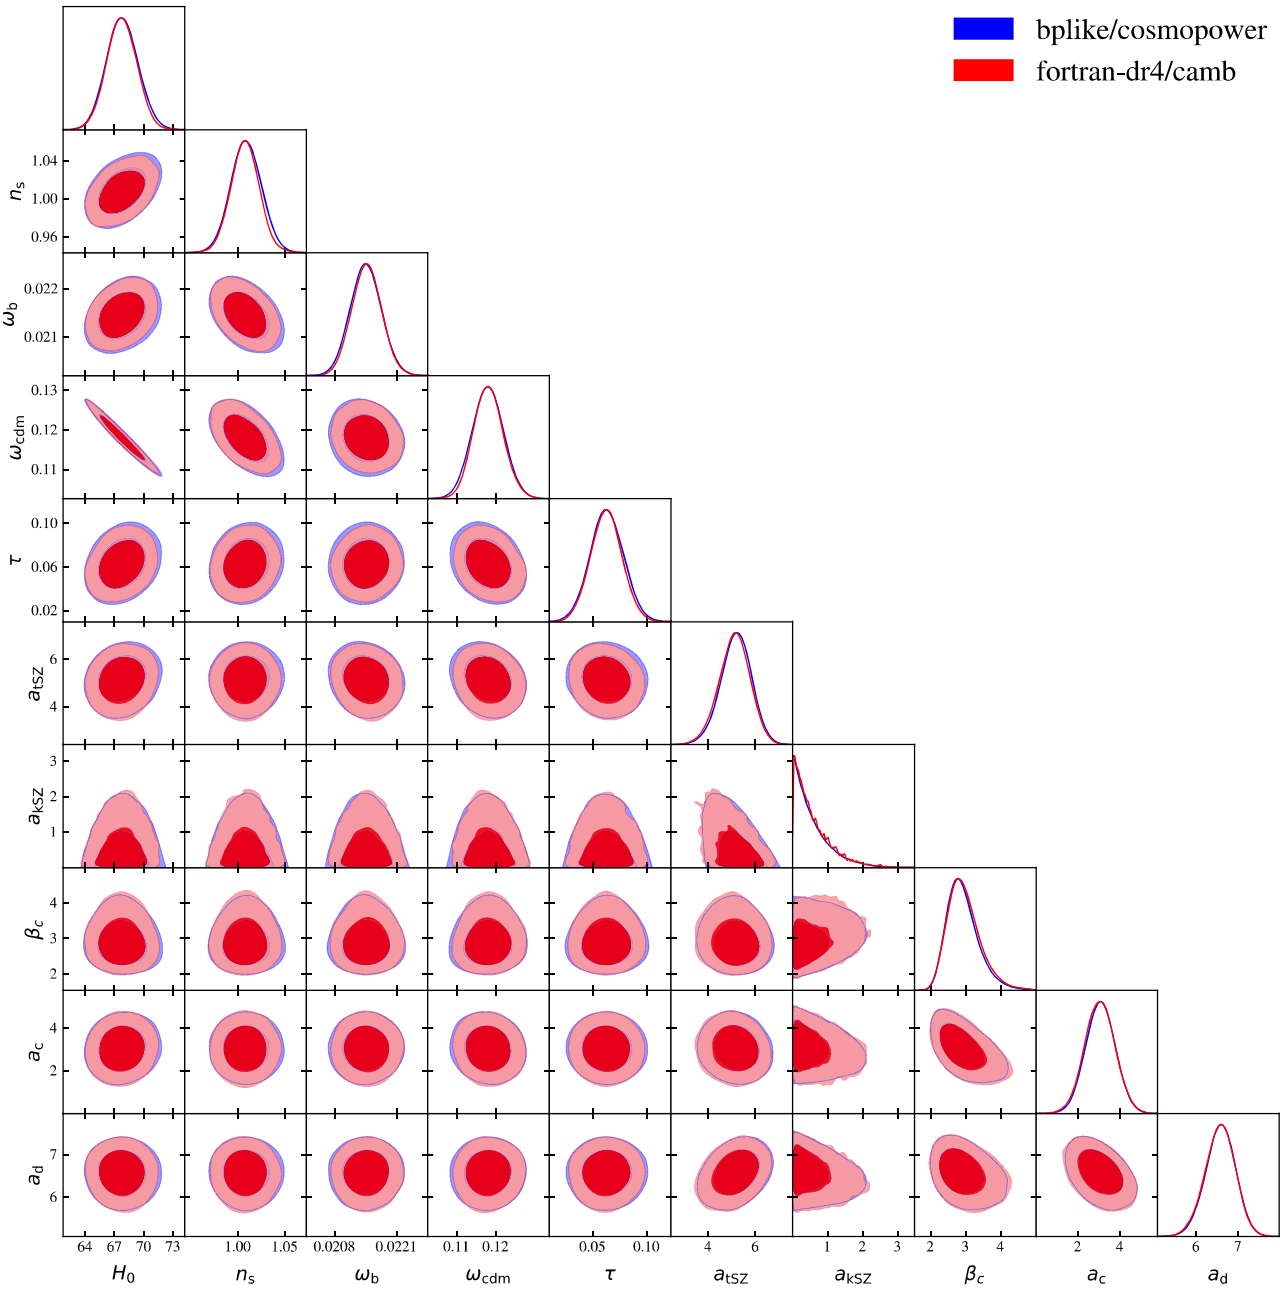

Figure 15. Subset of cosmological and foreground parameters obtained with the ACT DR4 full likelihood (see Choi et al. (2020) for parameters definition). 2D marginalized posterior probability distributions from `CosmoPower` (blue contours) are compared to the ones obtained with the high-precision `camb` calculation (red contours). See Section 5.1 for details.

It is possible to match `class` TT, TE, EE, TE, PKL, and PKNL spectra to `camb` to better than 0.1 per cent precision (for $\ell < 3000$ and $k < 10\,\mathrm{Mpc}^{-1}$), by using the precision parameters in the `cl_ref.pre` file in the `class` GitHub repository (this file ensures convergence at the 0.01 per cent level internally to `class` for TT and EE; Lesgourgues 2011b) and setting `k_max_tau0_over_l_max: 15` to ensure high accuracy at $\ell > 4000$ as discussed below). This is illustrated in Figs 4 and 6, where this calculation is shown as the dotted–dashed yellow lines, labeled 'class prec'. We refer to it as the ultra-high-precision `class` prediction. None the less, this computation takes $\mathcal{O}(1\,\mathrm{h})$ which is prohibitively long to generate $\mathcal{O}(10^5)$ spectra for training the emulators.

By optimizing the tradeoff between precision and computation time, we find that the minimal settings required in order to match the high-precision `camb` calculation with `class` v2.9.4 are obtained by setting the following three parameters:[11]

(i) `accurate_lensing: 1`
(ii) `k_max_tau0_over_l_max: 15.`
(iii) `perturb_sampling_stepsize: 0.05`

[11]Note that these parameters are optimized here, whereas McCarthy et al. (2022) used extremely high-precision `class` parameters without optimizing the tradeoff between precision and computation time.

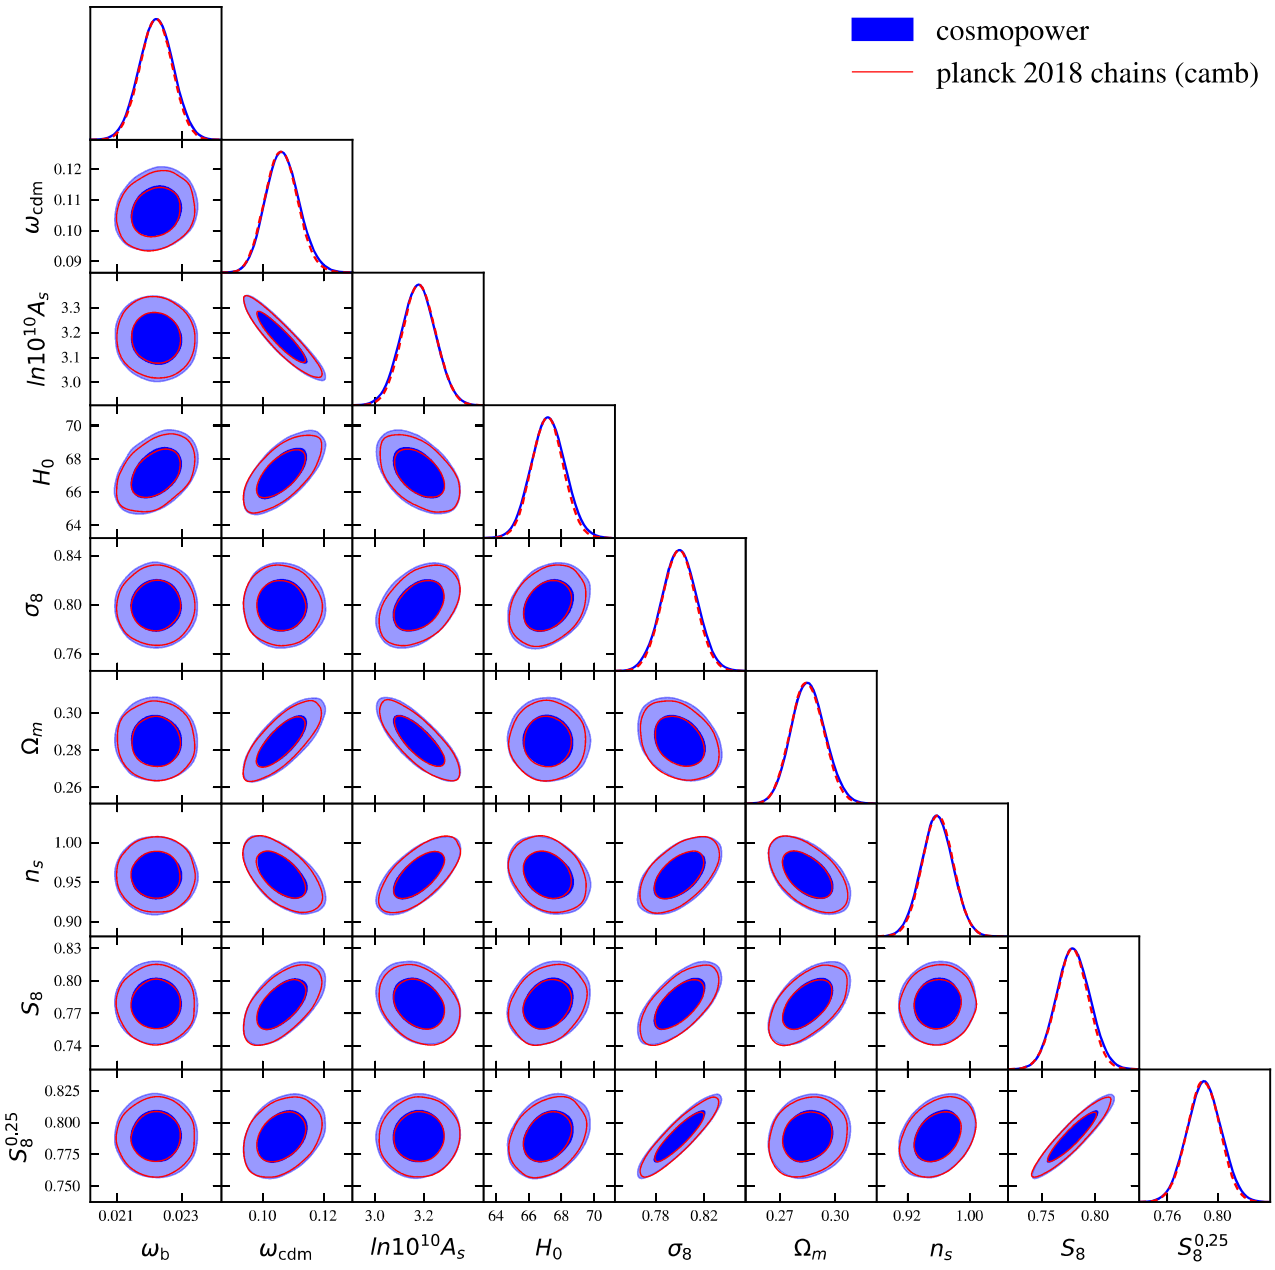

Figure 16. Comparison of 2D marginalized posterior probability distributions between `CosmoPower` (blue contours) and the reference `camb` chains (red contours, dashed line) downloaded from the *Planck* legacy archive (footnote 20) for the *Planck* 2018 CMB lensing likelihood + DES-Y1 + BAO analysis. See Section 5.2 for details.

These are the precision parameters that we use in order to generate our emulator training data. The first parameter ensures that the lensed TT, TE, and EE spectra are converged for $\ell > 3000$. Without it, they have non-physical oscillatory features. The second parameter ensures convergence at high-ℓ for all the spectra, including the unlensed ones. The third parameter is critical to get high-accuracy spectra over the whole multipole range. It is particularly important to get a converged PP spectra for $\ell <$ `l_switch_limber` (default value in `class` v2.9.4: `l_switch_limber = 10`). Note that out of these three `class` parameters, only the second one has an impact on the linear and non-linear matter power spectra. This high-precision `class` calculation takes roughly the same amount of time as the `camb` one, i.e. $\mathcal{O}$(1 min).

It is important to note that the first version of the emulators that we release in this paper are produced based on CMB and matter power spectra computed assuming `RECFAST` recombination model (Seager, Sasselov & Scott 1999; Scott & Moss 2009; Chluba, Fung & Switzer 2012), Big Bang Nucleosynthesis model from `Parthenope` (Consiglio et al. 2018) with a fiducial $N_{\rm eff} = 3.046$ using v1.2 of the code for a neutron lifetime of 880.2 s identical to standard assumptions of *Planck* 2017 papers. Furthermore, for the non-linear modelling of matter transfer functions we use `hmcode` from Mead et al. (2016). Another assumption that we make here is that the CMB monopole temperature is perfectly known and fixed to 2.7255 K, consistent with *COBE/FIRAS* measurements (Fixsen et al.

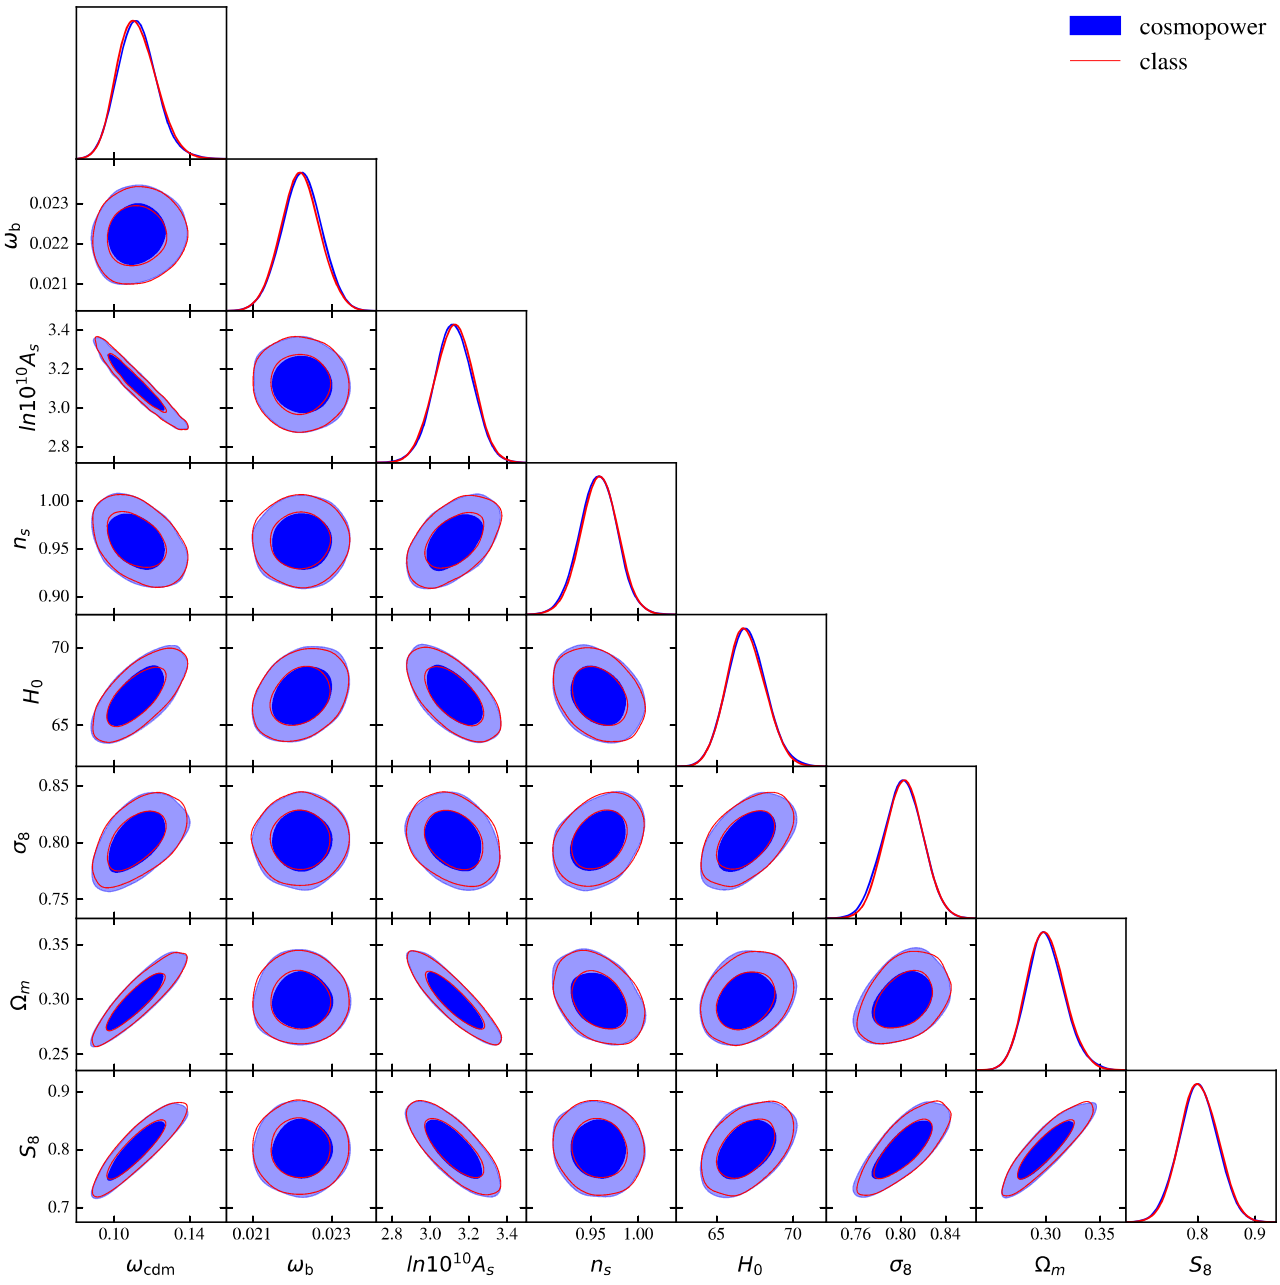

Figure 17. Comparison of 2D marginalized posterior probability distributions between `CosmoPower` (*Planck* 2018 CMB lensing, the BOSS 'full shape' likelihood including RSD information, to illustrate that $f\sigma_8$ can be computed with our emulators. See Section 5.3 for details. This validates the S8 emulator.

1996). We note that these settings, as well as the precision settings specified here, can affect the best-fitting models of future CMB and LSS experiments, and we advise they should be carefully quoted and specified.

4 EMULATORS

Our trained TT, TE, EE, PP, PKL, PKNL, H, DA, S8, and DER emulators in ΛCDM, wCDM, ΛCDM + $N_{\rm eff}$, ΛCDM + Σm_ν are made publicly available on the `CosmoPower` GitHub repository.[12]

[12]https://github.com/cosmopower-organization

To predict the power spectra, they need to be imported in Python, via the `CosmoPower` package. They allow for rapid computation of power spectra and derived parameters, and hence of parameter posterior probability distributions in MCMC analyses (see Section 5). Loading the emulators takes ≈0.1 s. Computing one set of TT, TE, EE, PKL, PKNL power spectra and derived parameters takes ≈60 ms, compared to $\mathcal{O}$(1 min) with `class` or `camb` with high-precision settings, i.e. we achieve a factor of 1000 speed-up. The GitHub repository also contains a set of notebooks showing how to use the emulators.

In Figs 1–6, we show the emulator's predictions for one set of cosmological parameters, namely, the central values of the right column of table I in Planck Collaboration

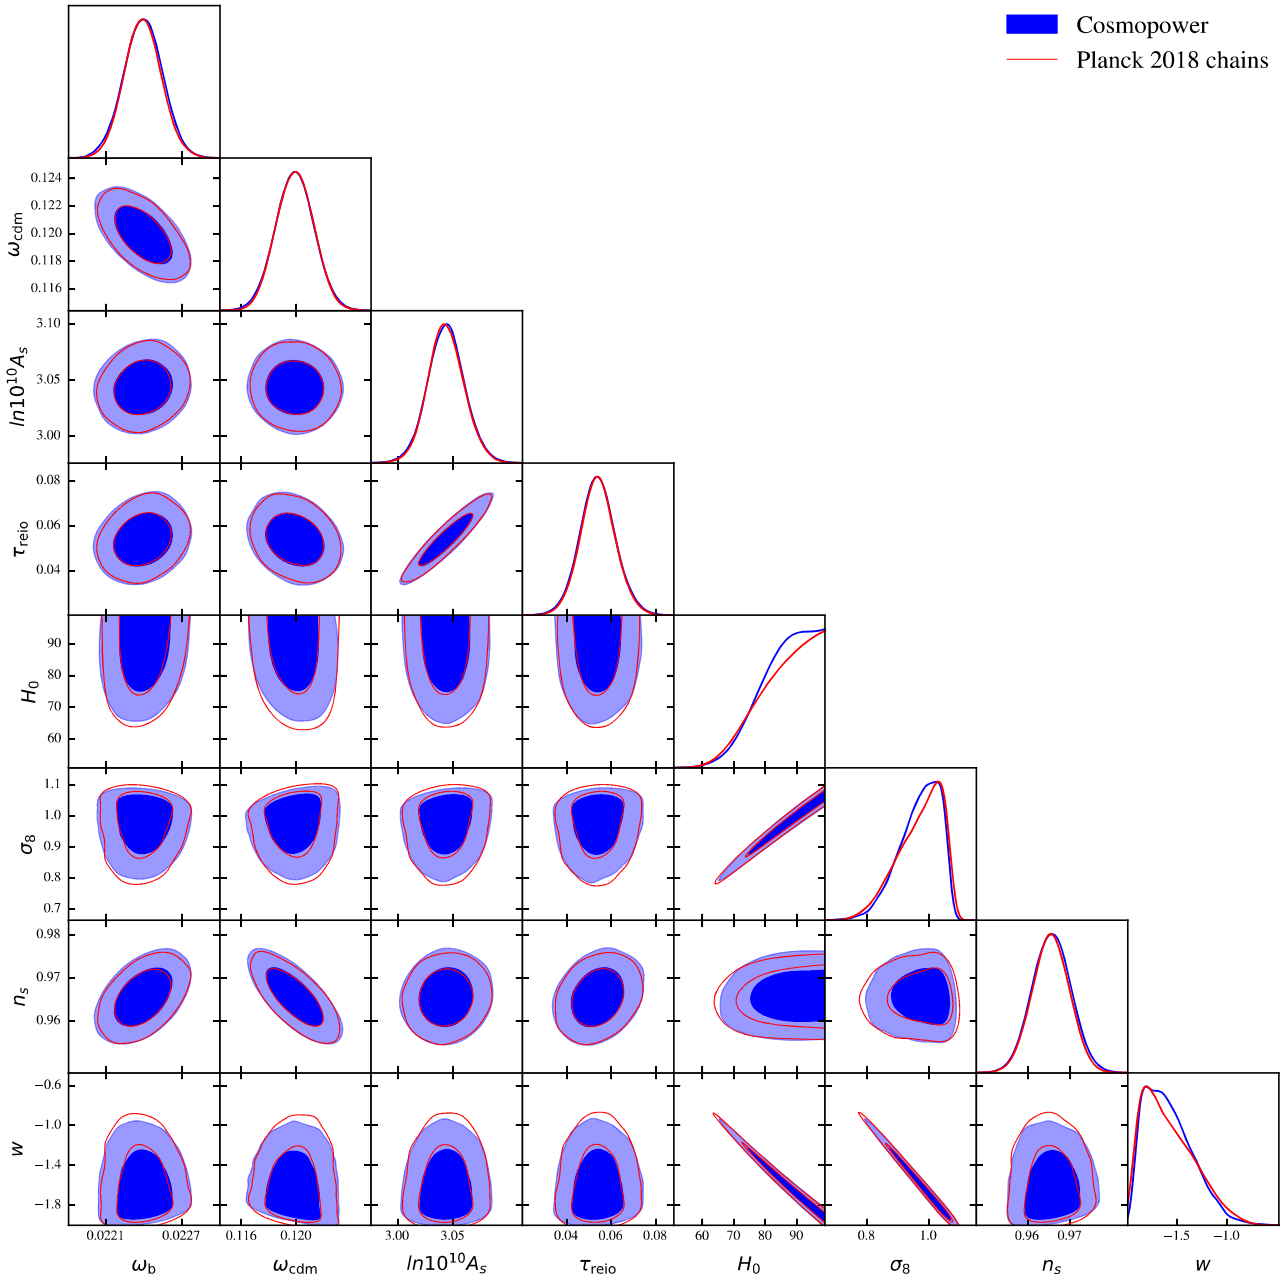

Figure 18. Comparison of 2D marginalized posterior probability distributions between `CosmoPower` (blue contours) and the full calculation (red contours) for the *Planck* 2018 CMB likelihood `HM_TTTEEE+lowl + lowE` computed using `camb` in wCDM and downloaded from the *Planck* legacy archive (footnote 20). See Section 5.4 for details. Note that our emulators cover the range $w > -2$ (which we used as a hard prior bound for our`CosmoPower` chains), as this is the lower bound adopted in the DES Year 1 analysis (Krause et al. 2017), although the *Planck* chains allow for $w > -3$.

(2020a), with one massive neutrino with $m_\nu = 0.06\,\mathrm{eV}$ and in ΛCDM.

In Fig. 1, we compare the power spectra from `class`, `camb`, and `CosmoPower` explicitly, The `class` (i.e. the high-precision calculation discussed in Section 3), 'class prec' (i.e. the ultra-high-precision calculation discussed in Section 3), high-precision `camb`, and `CosmoPower` power spectra are indistinguishable for $\ell < 10\,000$. For $10\,000 < \ell < 11\,000$, some differences can be seen in the TT and EE power spectra: the `class` and `CosmoPower` predictions tend to fall off compared to `camb`. Since this is less significant for the 'class prec' prediction, we attribute this to numerical precision settings, and do not investigate further as this is probably irrelevant for any CMB experiment in the foreseeable future. The shaded areas indicate the sensitivity of Advanced ACT (Henderson et al. 2016), SO (Simons Observatory 2019), and CMB-S4 (CMBS4 2022) forecast sensitivity. We refer to section 4 of Bolliet et al. (2022) for details on the TT and PP sensitivity curves and the links where they can be downloaded from. The TE sensitivity can be computed from TT and EE[13] (see e.g. Spurio Mancini et al. 2022, for the formula we use).

[13]The CMB-S4 EE noise is from `S4_190604d_2LAT_pol_default_noisecurves_deproj0_SENS0`

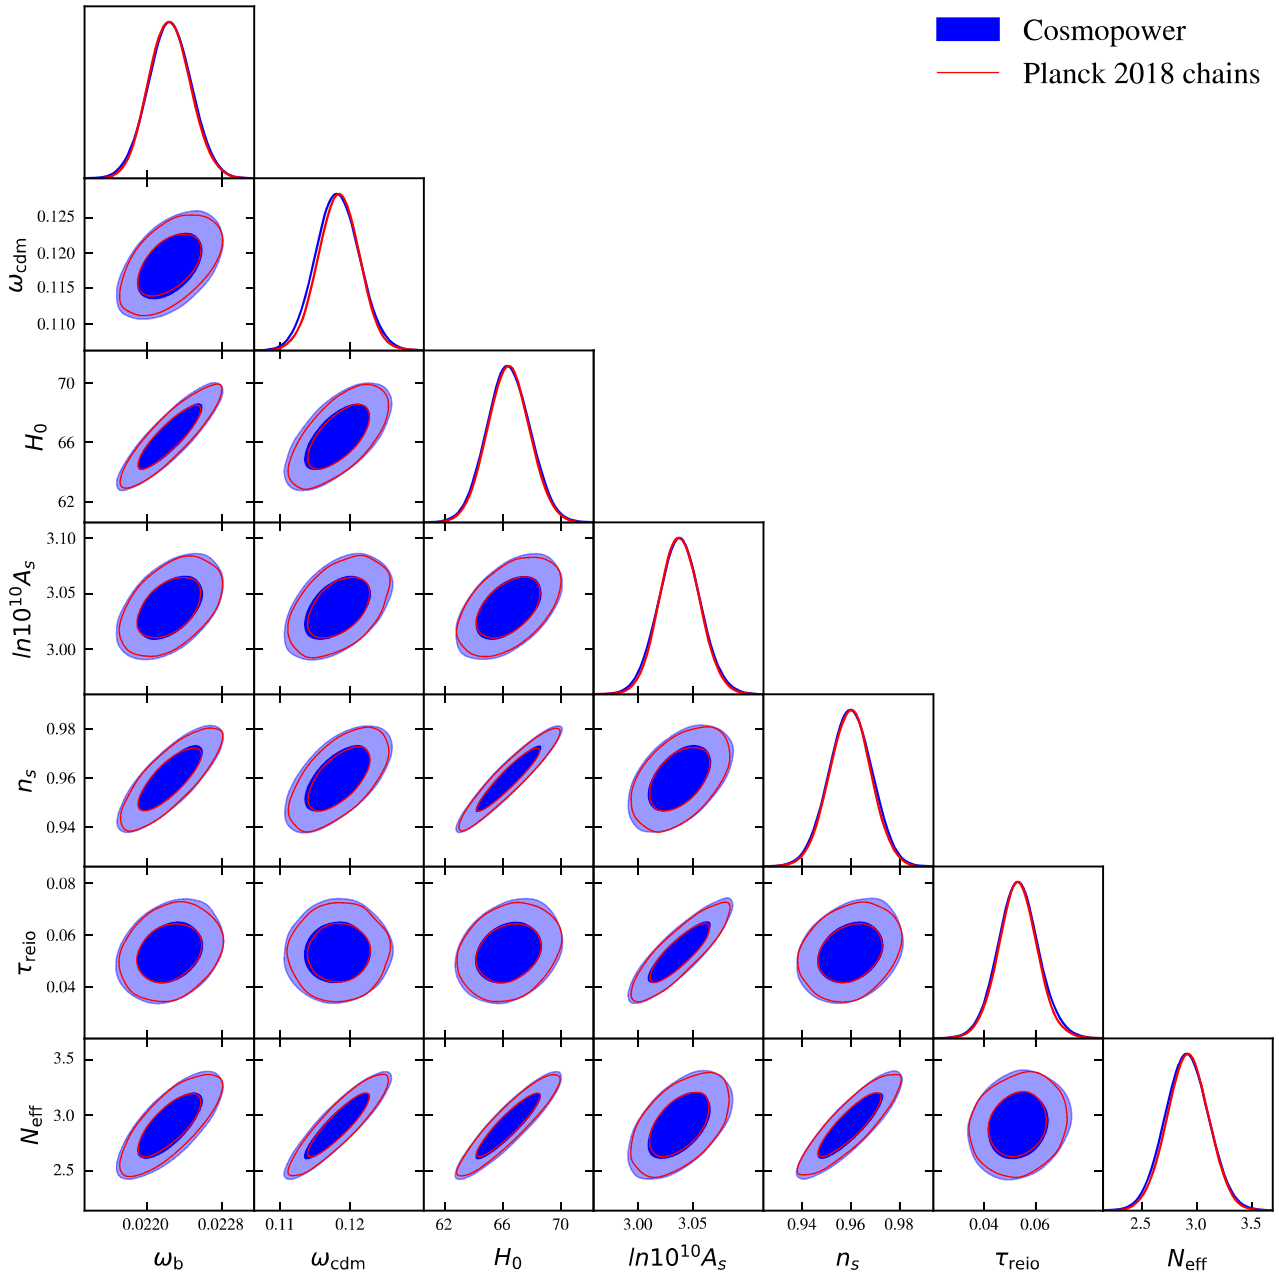

Figure 19. Comparison of 2D marginalized posterior probability distributions between `CosmoPower` (blue contours) and the full calculation (red contours) for the *Planck* 2018 CMB likelihood `HM_TTTEEE+lowl + lowE` (footnote 20) computed using `camb` in ΛCDM + N_{eff} and downloaded from the *Planck* legacy archive (footnote 20). See Section 5.4 for details.

Note that we use $f_{sky} = 0.3$ for Advanced ACT and $f_{sky} = 0.4$ for SO and CMB-S4. Also, note that all noise curves also assume the use of *Planck* data in combination with that from each ground-based experiment (see section 2 of Simons Observatory (2019) for details); the use of *Planck* is important in extending the forecast measurements to low multipoles, where otherwise atmospheric noise in the ground-based data would be very large.

`_mask_16000_ell_EE_BB.txt` and the SO EE noise is from `v3.1.0/SO_LAT_Nell_P_baseline_fsky0p4_ILC_CMB_E.txt`

Fig. 2 shows the same as Fig. 1 but for the linear and non-linear matter power spectra, computed at three redshifts ($z = 0$, 2.5, and 5). In these cases, the `class`, `camb`, and `CosmoPower` predictions are indistinguishable across the full k range. The same is true for the H, DA, and S8 predictions, as can be seen in Fig. 3.

In Fig. 4, we show the relative difference between CMB angular power spectra (TT,TE,EE,PP) and `camb`, in per cent. The spikes in TE are where the spectra cross zero and are not problematic. As discussed before, the agreement between the ultra-high precision `class` prediction and `camb` is at the 0.1 per cent level until $\ell \sim 3000$, and degrades at higher ℓ. None the less, the agreement

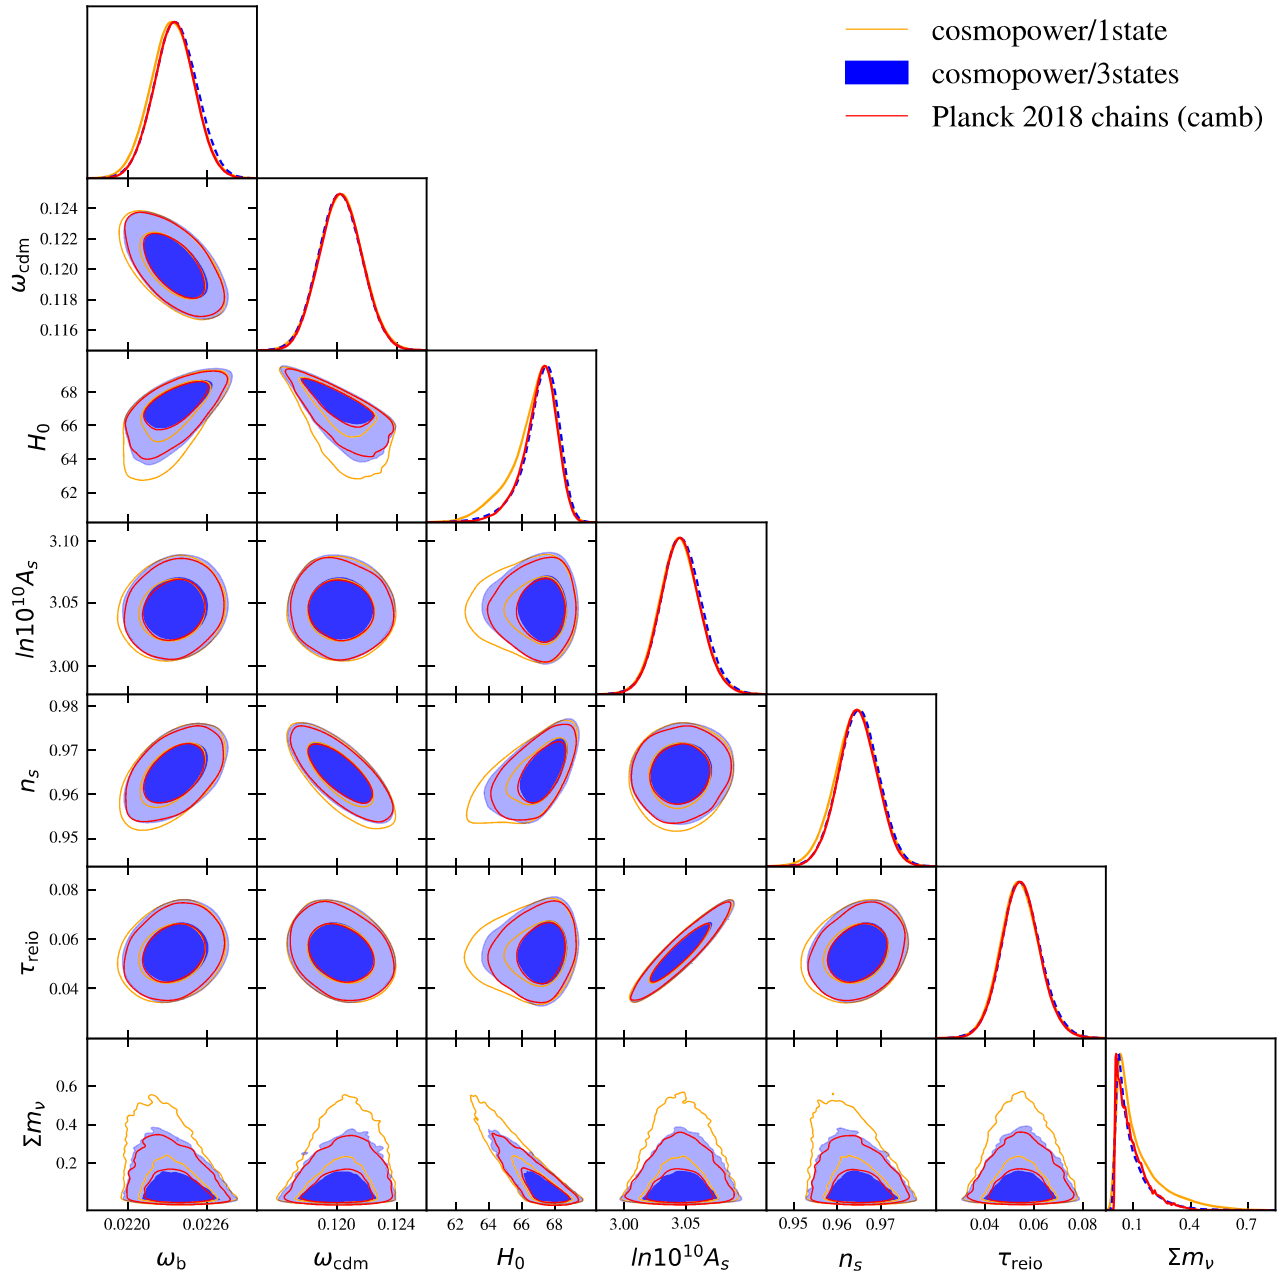

Figure 20. Comparison of 2D marginalized posterior probability distributions between `CosmoPower` (blue contours) and the full calculation for the *Planck* 2018 CMB likelihood `HM_TTTEEE+lowl + lowE` computed using `camb` (red contours) in ΛCDM + Σm_ν and downloaded from the *Planck* legacy archive (footnote 20). The yellow contours are the constraints using the ΛCDM + Σm_ν emulator with one massive and two massless neutrinos (unlike the *Planck* chains), rather than three massive neutrinos (blue and red contours). The model with three degenerate states is a better approximation to the mass splitting results. (We thank Antony Lewis for pointing this out to us.) See Section 5.4 for further details.

between `CosmoPower` and both `class` predictions remains at the 0.1 per cent level across the whole range of multipoles, including $\ell > 3000$, except in two cases. First, in the low-ℓ part ($\ell < 10$) of the lensing convergence power spectrum, we see very small differences ($\lesssim$0.5 per cent) between the ultra-high precision `class` prediction on one hand and the high-precision and `CosmoPower` prediction on the other hand (the agreement between the `CosmoPower` and high-precision `class` calculation remains better than 0.005 per cent). This difference can be reduced further by decreasing the parameter `perturb_sampling_stepsize` (see Section 3), at the cost of a longer runtime. Secondly, in the low-ℓ part ($\ell < 50$) of the EE power spectrum, we see differences between the `CosmoPower` and high-precision `class` prediction at the 0.6 per cent level. If needed, this could be solved by generating more training data.

To assess whether this level of difference is acceptable for CMB Stage-IV analyses, we compare the relative difference between the CMB spectra in terms of CMB-S4 sensitivity in Fig. 5. We see that the `cosmopower` power spectra agree with `camb` to better than 0.03σ across all multipoles and with the high-precision `class`

prediction to better than 0.03σ. Therefore, such level of agreement is sufficient.

In Fig. 6, we show the relative difference between the linear and non-linear matter power spectra with respect to `camb` in per cent, at three redshifts, $z = 0$, 2.5, and 5. In the left-hand panels, we see that the `CosmoPower` linear matter power spectra agree with the high-precision `class` predictions at the 0.05 per cent level across the whole k range. The difference between these and the ultra-high precision `class` prediction is at the 0.3 per cent level. Finally, the difference between the ultra-high precision `class` prediction and `camb` is at the 0.3 per cent level up to $k \approx 20\,\mathrm{Mpc}^{-1}$ and then becomes ≈ 0.8 per cent at higher k. In the right-hand panels, showing the non-linear matter power spectra, we see that differences between `cosmopower` on one hand and `class` and `camb` on the other hand are more important for the non-linear matter power spectra, becoming larger in the non-linear regime, but remain at the 1–1.5 per cent level, including at $z = 5$. The fall-off of the linear $P(k)$ at small scales compared to `camb` is likely to be due to baryonic sound speed effect on non-linear scales that are treated differently in `class` and `camb`.

Since there is no direct measurement of the matter power spectrum, we cannot compare it with an instrumental sensitivity level. However, we can check the accuracy on observables based on the matter power spectrum. The lensing convergence power spectrum (PP, bottom right panel of Fig. 1) is one such example. Other examples include galaxy clustering and galaxy weak lensing observables and to a lesser extent, the lensing of CMB angular power spectra. We perform Stage III posterior inference analyses based on these observables and recover nearly identical constraints to those obtained running the full Boltzmann code calculations (see Section 5), giving us confidence that our emulators are sufficiently precise. We defer to future work a more detailed quantitative analysis, including the release of high-accuracy emulators for the matter power spectrum tested against (forecast) Stage-IV configurations.

The results shown in Figs 1–6, are for one set of cosmological parameters. However, it is important to quantify the precision of our emulators over the whole prior range of parameters (see Table 1), in all the cosmological models. We do so in Figs 7–13.[14] These figures show the relative difference between the emulators, i.e. the `CosmoPower` prediction, and the exact high-precision `class` prediction on the testing data sets,[15] containing ≈25 000 predictions. The results for TT are in Fig. 7, for EE in Fig. 8, for TE in Fig. 9, for PP in Fig. 10, for PKL and PKNL in Fig. 11, for H, DA, and S8 in Fig. 12, and for the derived parameters in Fig. 13. For TT, TE, EE, and PP, we show the difference in units of CMB-S4 sensitivity and in percentages, and for the other emulators, we show the difference in percentages (since there is no simple way to compare with instrumental sensitivity; see comments in the previous paragraph). In all the figures, the three shades of red correspond to 68 per cent, 95 per cent, and 99 per cent of the testing data set, from dark to light, respectively. For instance, for TT (Fig. 7), we find that our emulators predict the power spectra to better than 0.02σ at CMB-S4 precision level, 68 per cent of the time in all the four cosmological models considered here. The agreement for TE and EE is similar. For PP the agreement is at the 0.02σ CMB-S4 precision level, 99 per cent of the time in all the four cosmological models considered here.

The quantity that is the least accurately emulated is the non-linear matter power spectrum, which is reproduced at the $1 - 1.5$ per cent level 95 per cent of the time in all models (see Fig. 11). This level of agreement is satisfactory for application to current survey configurations, especially since non-linear modelling and baryonic uncertainty is at the ≈ 10 per cent level (see e.g. Mead et al. 2021) and even exceeds 30 per cent at $k > 1\,\mathrm{Mpc}^{-1}$ (see e.g. Amon & Efstathiou 2022).

The emulated redshift-dependent $H(z)$, $D_A(z)$, and $\sigma_8(z)$ agree to better than 0.05 per cent between $z = 0$ and 20, for 99 per cent of the testing set in all models (see Fig. 12). We find a similar agreement between emulated derived parameters and the testing data set (see Fig. 13).

5 ACCELERATED LIKELIHOOD ANALYSES

In this section, we run posterior inference analysis and use the emulators presented above in MCMC extractions of cosmological parameters. A full Stage-IV analysis will require optmization of likelihood codes alongside emulators which is beyond the scope of this paper and will be covered in future work. Therefore, here we demonstrate accelerated likelihood analyses for current (Stage II-III) data and reproduce marginalized constraints for ACT DR4 (Section 5.1), *Planck* lensing + DES + BAO (Section 5.2) and *Planck* TT,TE,EE in wCDM, ΛCDM + $N_{\rm eff}$ and ΛCDM + Σm_ν (Section 5.4). We use `cobaya` (Torrado & Lewis 2021) for MCMC sampling. To analyse the resulting chains, we use `GetDist` (Lewis 2019).

We define a `CosmoPower` theory object within `cobaya` such that we are able to implement the `CosmoPower` call completely outside of the likelihoods. Our implementation of the `CosmoPower` theory object is available upon request. In a forthcoming paper, we plan to release an official `CosmoPower` wrapper for both `cobaya` and `cosmosis` (Zuntz et al. 2015).

5.1 Accelerated ACT DR4 power spectrum analysis

We use the ACT Data Release 4 (DR4) 'actpol_lite' likelihood (Aiola et al. 2020; Choi et al. 2020) with `CosmoPower` to reproduce the original DR4 results. This is the foreground-marginalized likelihood publicly available in the `pyactlike` repository.[16] For the `CosmoPower` runs, this likelihood involves the TT, TE, and EE emulators, relying on multipoles up to $\ell \approx 4500$.

In Fig. 14, we show the resulting 2D marginalized posterior probability distribution obtained in a few minutes on a laptop with `CosmoPower` as red empty contours, and compare it with the publicly available chains[17] shown as the filled green contours. The reference `class` chains were obtained in Hill et al. (2022) using high-accuracy settings and took several days to converge. The overlap of the `CosmoPower` and reference contours is nearly perfect. We note that for the `CosmoPower` run we sample over H_0 and obtain θ_s as a derived parameter using our DER emulator in post-processing of the chains, whereas the reference `class` chains were computed with θ_s as an input parameter and H_0 was saved as a derived parameter.

[14]The 'mnu' results on this figure are for one massive and two massless neutrino states.

[15]The testing data are made of the exact computations which are not used for training the emulator neural networks. We use 20 per cent of our data set for testing.

[16]https://github.com/ACTCollaboration/pyactlike

[17]CLASS2p8_ACTPol_lite_DR4_leakfix_yp2_baseLCDM_taup_hip_R0p01 downloaded from LAMBDA.

As a second test, we perform a maximum-likelihood analysis using the full ACT DR4 likelihood, including foregrounds. To do so we use a public python implementation `bpllike`.[18] The `CosmoPower` results (blue contours) are shown in Fig. 15 and compared with the reference high-precision `camb` chains that we ran using the original full Fortran ACT DR4 likelihood[19] (red contours). Again, the overlap between the contours is nearly perfect. The foreground parameter contours (for the thermal SZ and kinematic SZ amplitudes and CIB SED power-law index and amplitude) match exactly. Cosmological parameter contours agree to better than 0.1σ, except for θ_s. This is simply because of different definitions in `class` and `camb`. Note that our `cosmopower` chains reach $R-1=0.05$ in ≈20 min on a laptop with 10 threads.

5.2 Accelerated *Planck* lensing + DES + BAO analysis

To further demonstrate the efficiency and wide range of our emulators, we reproduce cosmological parameters extracted from the *Planck* CMB lensing power spectrum + DES + BAO analysis. We compare with the publicly available *Planck* chains.[20] The chains of interest are labeled `DES_lenspriors_lensing_BAO`.

We use the Python-native re-implementation of the *Planck* reconstructed lensing power spectrum likelihood, assuming their fiducial (conservative) multipole range. This likelihood is available in `cobaya` as `planck_lensing_2018` (Planck Collaboration 2020b). We apply the priors used in the *Planck* lensing analysis: Gaussian priors on $n_s=0.96\pm0.02$ and $\Omega_b h^2=0.0222\pm0.0005$, a flat prior $40<H_0/(\mathrm{km\,s^{-1}\,Mpc^{-1}})<100$, and a fixed $\tau=0.055$ (see Planck Collaboration 2020b, for details).

We use the DES-Y1 cosmic shear + galaxy auto- and cross-correlation ('3 × 2-point') likelihood, as implemented in `des_y1` in `cobaya`, which is as described by Troxel et al. (2018), DES Collaboration (2018), and Krause et al. (2017).

Finally, we use the BAO likelihoods corresponding to BOSS DR12 (Alam et al. 2017), the SDSS Main Galaxy Sample (Ross et al. 2015), and the 6dF survey (Beutler et al. 2011). These are implemented in `cobaya` as `sdss_dr12_consensus_bao`, `sdss_dr7_mgs`, and `sixdf_2011_bao`.

For the `CosmoPower` runs, this likelihood combination involves the PKNL, H, DA, and PP emulators. Because this calculation requires the computation of the non-linear $P(k)$ at many redshifts, it is slightly more time-consuming. Indeed, we need to evaluate the matter power spectrum on a 2D grid in (k, z), hence requiring us to call the emulator as many times as there are points in the z dimension. We set the $P(k, z)$ grid to cover redshifts between 0 and 4 with 15 linearly spaced points. With this choice, `cobaya` performs ≈16 evaluations per second on an ARM64 MacBook Pro. We reach $R-1\simeq0.15$ after ≈30 min on a laptop, with four chains taking a total of 56 000 accepted steps with acceptance rate of 0.2.

The results are shown in Fig. 16. The overlap between the `CosmoPower` and reference `camb` contours is perfect. As derived parameters, we show $S_8\equiv\sigma_8(\Omega_m/0.3)^{0.5}$ and $S_8^{0.25}\equiv\sigma_8(\Omega_m/0.3)^{0.25}$, which are most constrained by galaxy weak lensing and CMB lensing data, respectively. Our `CosmoPower` runs recover the reference constraints on S_8 and $S_8^{0.25}$ exactly. This validates the matter power spectrum emulators for current galaxy WL surveys.

5.3 RSD full shape analysis

To demonstrate that we can also compute $f\sigma_8$ and use it in a likelihood analysis, using our S8 emulator (see Section 2), we perform a likelihood analysis and use `sdss_dr12_consensus_fullshape` with *Planck* lensing. We compare `class` and `CosmoPower` chains in Fig. 17.

5.4 Extensions

To validate our emulators in extended cosmology we run maximum-likelihood analyses in wCDM, ΛCDM + $N_{\rm eff}$, and ΛCDM + Σm_ν. For the data and likelihood, we choose the official `clik` *Planck* 2018 `plikHMTTTEEE+lowl + lowE` (Planck Collaboration 2020a) implementation, which we call through `cobaya`.

We show the wCDM, ΛCDM + $N_{\rm eff}$, and ΛCDM + Σm_ν constraints in Figs 18–20, respectively. In each extension, we overplot the reference contours from `camb` chains (empty red contours) that we obtain online (see footnote 20). Our `CosmoPower` contours are in blue. In all cases the agreement with the reference *Planck* chains is excellent. To reach $R-1\approx0.1$, `CosmoPower` needs approximately 25 min, while `class` or `camb` chains typically take $\mathcal{O}$(1 day) to converge.

6 CONCLUSIONS

We have extended the work of Spurio Mancini et al. (2022) by creating emulators for CMB temperature and polarization angular anisotropy power spectra, CMB lensing convergence power spectra, linear and non-linear matter power spectra, and the redshift evolution of $H(z)$, $D_A(z)$ and $\sigma_8(z)$, in ΛCDM and extensions, namely, wCDM, ΛCDM + $N_{\rm eff}$ and ΛCDM + Σm_ν with one or three massive states (see Section 2). All of these quantities are computed using high-precision `class` settings (see Section 4). They are sufficiently precise for Stage-IV CMB analyses (see Section 3) and current galaxy weak lensing and clustering surveys. We have tested all of our emulators in maximum-likelihood analyses, involving high-ℓ CMB, galaxy weak lensing and clustering, BAO and RSD data (see Section 5).

The main outcome of this work is to open the door to fast parameter inference, via the widely used MCMC method, with accelerated Boltzmann computations thanks to our `CosmoPower` emulators. For instance, as we demonstrate here, most of Stage-III parameter inference is now feasible on a laptop.

In a forthcoming paper, we plan release wrappers of our emulators so that they can be readily used within `cobaya` (Torrado & Lewis 2021) and `cosmosis` (Zuntz et al. 2015).

We created an online repository to store our emulators and will continue updating it. For CMB polarization, we have made emulators for E modes, which could be used to accelerate *LiteBIRD* (LiteBIRD Collaboration 2022) analyses dedicated on reionization history constraints (Zaldarriaga et al. 2008; Allys et al. 2022). We defer B-modes power spectra emulators to future work.

In addition to the applications considered here, these emulators can also be used to save time at the initial step of libraries for computations of cosmological LSS observables like `class_sz` (Bolliet et al. 2018; Bolliet et al. 2022), `ccl` (Chisari et al. 2019), and any other code which relies on the quantities emulated here. In fact, `class_sz` now contains a fully-fledged wrapper for our emulators (see Bolliet et al. 2023, the online code repository[21] and

[18] https://github.com/ACTCollaboration/bplike
[19] ACT DR4 Fortran likelihood, named actpolfull there.
[20] https://wiki.cosmos.esa.int/planck-legacy-archive
[21] https://github.com/CLASS-SZ/class_sz

tutorial notebooks[22]). The `class_sz` wrapper has recently been employed in a new likelihood code package for cluster number counts: `cosmocnc`[23] (Zubeldia & Bolliet 2024).

Recent works (Cabass et al. 2022; Philcox & Ivanov 2022) have used the effective field theory of LSS (Baumann et al. 2012; Carrasco, Hertzberg & Senatore 2012) in order to derive constraints from spectroscopic surveys, based on the one-loop galaxy power spectrum and bispectrum. These calculations are time-consuming (MCMC convergence takes $\mathcal{O}$(1day) using `class_pt` (Chudaykin et al. 2020). The bottleneck is the calculation of higher-order correlators. The work presented here does not directly help to accelerate such analyses, but none the less the method can directly be applied in this context as well.

We note that `class` has been recently updated[24] so it can compute high-ℓ CMB lensing faster and more accurately than what is possible with the version of the code used here. Therefore, it should be possible to generate `class` based emulators with even more demanding high-accuracy settings. Other possible improvement include more refined recombination treatment, e.g. using `CosmoRec` (Chluba & Thomas 2011; Chluba & Thomas 2013) or `HyRec` (Ali-Haïmoud & Hirata 2011; Lee & Ali-Haïmoud 2020). As noted and discussed in (Giardiello et al. 2024; Jense et al. 2024), this will be useful for upcoming and future Stage-IV CMB analyses. Finally, the uncertainty in CMB temperature (which fixes the radiation density), even though it is small, will eventually become important to take into account as accuracy requirements increase. We leave this point of investigation to future work.

ACKNOWLEDGEMENTS

We are very grateful to Antony Lewis and Julien Lesgourgues for many valuable discussions which improved the work presented here, as well as Ian Harrison, Jens Chluba, Yacine Ali-Haimoud and an anonymous referee for comments on the manuscript. We thank Nick Carriero and Robert Blackwell for key computational advice.

We acknowledge the use of computational resources at the Flatiron Institute. The Flatiron Institute is supported by the Simons Foundation. ASM acknowledges support from the MSSL STFC Consolidated Grant ST/W001136/1. This work was partially enabled by funding from the University College London (UCL) Cosmoparticle Initiative and by collaborative visits funded by Princeton University and the Cosmology and Astroparticle Student and Postdoc Exchange Network (CASPEN). ASM thanks Princeton University and the Flatiron Institute for their hospitality during his visit, during which this project was started. JCH acknowledges support from NSF grant AST-2108536, NASA grant 21-ATP21-0129, DOE grant DE-SC00233966, the Sloan Foundation, and the Simons Foundation. BB acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 851274). EC and HJ acknowledge support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 849169). JD acknowledges support from NSF grant AST-2108126.

[22] https://github.com/CLASS-SZ/notebooks
[23] https://github.com/inigozubeldia/cosmocnc
[24] See v3.2 comment here: http://class-code.net

DATA AVAILABILITY

The emulators are available online at CosmoPower Organisation, including tutorial notebooks. If you use the emulators, please cite our work as well as Spurio Mancini et al. (2022).

REFERENCES

Abadi M. et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. https://www.tensorflow.org/
Abazajian K. et al., 2019, preprint (arXiv:1907.04473)
Aihara H. et al., 2018, PASJ, 70, S4
Aiola S. et al., 2020, J. Cosmol. Astropart. Phys., 12, 047
Akeson R. et al., 2019, preprint (arXiv:1902.05569)
Alam S. et al., 2017, MNRAS, 470, 2617
Albers J., Fidler C., Lesgourgues J., Schöneberg N., Torrado J., 2019, J. Cosmol. Astropart. Phys., 09, 028
Ali-Haïmoud Y., Hirata C. M., 2011, Phys. Rev. D, 83, 043513
Allys E. et al., 2022
Alsing J. et al., 2020, ApJS, 249, 5
Amon A., Efstathiou G., 2022, MNRAS, 516, 5355
Aricò G., Angulo R. E., Zennaro M., 2021, preprint (arXiv:2104.14568)
Auld T., Bridges M., Hobson M. P., Gull S. F., 2007, MNRAS, 376, L11
Balkenhol L. et al., 2022, preprint (arXiv:2212.05642)
Balkenhol L., Trendafilova C., Benabed K., Galli S., 2024
Baumann D., Nicolis A., Senatore L., Zaldarriaga M., 2012, J. Cosmol. Astropart. Phys., 07, 051
Benson B. A. et al., 2014, in Holland W. S., Zmuidzinas J., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII. p. 91531P, preprint (arXiv:1407.2973)
Beutler F. et al., 2011, MNRAS, 416, 3017
Blas D., Lesgourgues J., Tram T., 2011, J. Cosmol. Astropart. Phys., 2011, 034
Bolliet B., Comis B., Komatsu E., Macías-Pérez J. F., 2018, MNRAS, 477, 4957
Bolliet B., Hill J. C., Ferraro S., Kusiak A., Krolewski A., 2022, preprint (arXiv:2208.07847)
Bolliet B. et al., 2023, in mm Universe 2023: Observing the mm Universe at mm wavelengths. preprint (arXiv:2310.18482)
Bonici M., Bianchini F., Ruiz-Zapatero J., 2023
CMBS4, 2022, ApJ, 926, 54
Cabass G., Ivanov M. M., Philcox O. H. E., Simonović M., Zaldarriaga M., 2022, Phys. Rev. Lett., 129, 021301
Carrasco J. J. M., Hertzberg M. P., Senatore L., 2012, JHEP, 09, 082
Chisari N. E. et al., 2019, ApJS, 242, 2
Chluba J., Thomas R. M., 2011, MNRAS, 412, 748
Chluba J., Thomas R. M., 2013, Astrophysics Source Code Library, record ascl:1304.017
Chluba J., Fung J., Switzer E. R., 2012, MNRAS, 423, 3227
Choi S. K. et al., 2020, J. Cosmol. Astropart. Phys., 12, 045
Chudaykin A., Ivanov M. M., Philcox O. H. E., Simonović M., 2020, Phys. Rev. D, 102, 063533
Consiglio R., de Salas P. F., Mangano G., Miele G., Pastor S., Pisanti O., 2018, Comput. Phys. Commun., 233, 237
DES Collaboration, 2018, Phys. Rev. D, 98
DESI Collaboration et al., 2016, preprint (arXiv:1611.00036)
Dark Energy Survey Collaboration, 2016, MNRAS, 460, 1270
Dawson K. S. et al., 2013, AJ, 145, 10
Doré O. et al., 2014, preprint (arXiv:1412.4872)
Fendt W. A., Wandelt B. D., 2007, ApJ, 654, 2
Fixsen D. J., Cheng E. S., Gales J. M., Mather J. C., Shafer R. A., Wright E. L., 1996, ApJ, 473, 576
Giardiello S. et al., 2024, preprint (arXiv:2403.05242)
Günther S., Lesgourgues J., Samaras G., Schöneberg N., Stadtmann F., Fidler C., Torrado J., 2022, J. Cosmol. Astropart. Phys., 11, 035
Henderson S. W. et al., 2016, J. Low Temp. Phys., 184, 772

Hill J. C. et al., 2022, Phys. Rev. D, 105, 123536
Ivezić v. et al., 2019, ApJ, 873, 111
Jense H. et al., 2024, preprint (arXiv:2405.07903)
Krause E. et al., 2017, preprint (arXiv:1706.09359)
Kuijken K. et al., 2015, MNRAS, 454, 3500
Laureijs R. et al., 2011, preprint (arXiv:1110.3193)
Lee N., Ali-Haïmoud Y., 2020, Phys. Rev. D, 102, 083517
Lesgourgues J., 2011a, preprint (arXiv:1104.2932)
Lesgourgues J., 2011b, preprint (arXiv:1104.2934)
Lewis A., 2019, preprint (arXiv:1910.13970)
Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
LiteBIRD Collaboration, 2022, preprint (arXiv:2202.02773)
McCarthy F., Hill J. C., Madhavacheril M. S., 2022, Phys. Rev. D, 105, 023517
Mead A. J., Peacock J. A., Heymans C., Joudaki S., Heavens A. F., 2015, MNRAS, 454, 1958
Mead A., Heymans C., Lombriser L., Peacock J., Steele O., Winther H., 2016, MNRAS, 459, 1468
Mead A. J., Brieden S., Tröster T., Heymans C., 2021, MNRAS, 502, 1401
Mootoovaloo A., Jaffe A., Heavens A., Leclercq F., 2022, Astron. Comput., 38, 100508
Philcox O. H. E., Ivanov M. M., 2022, Phys. Rev. D, 105, 043517
Planck Collaboration, 2020a, A&A, 641, A6
Planck Collaboration, 2020b, Astron. Astrophys., 641, A8
Ross A. J., Samushia L., Howlett C., Percival W. J., Burden A., Manera M., 2015, MNRAS, 449, 835
Scott D., Moss A., 2009, MNRAS, 397, 445
Seager S., Sasselov D. D., Scott D., 1999, ApJ, 523, L1
Sehgal N. et al., 2019, in Bulletin of the American Astronomical Society. p. 6, preprint (arXiv:1906.10134)
Simons Observatory, 2019, J. Cosmol. Astropart. Phys., 2019, 056
Smith R. E. et al., 2003, MNRAS, 341, 1311
Spergel D. et al., 2015, preprint (arXiv:1503.03757)
Spurio Mancini A., Piras D., Alsing J., Joachimi B., Hobson M. P., 2022, MNRAS, 511, 1771
Torrado J., Lewis A., 2021, J. Cosmol. Astropart. Phys., 05, 057
Troxel M. A. et al., 2018, Phys. Rev. D, 98, 043528
Zaldarriaga M. et al., 2008
Zubeldia I. n., Bolliet B., 2024
Zuntz J. et al., 2015, Astron. Comput., 12, 45

This paper has been typeset from a TEX/LATEX file prepared by the author.