



Research



Cite this article: Bartlett D, Pandey S. 2026
Symbolic emulators for cosmology:
accelerating cosmological analyses without
sacrificing precision. *Phil. Trans. R. Soc. A* **384**:
20240585.

<https://doi.org/10.1098/rsta.2024.0585>

Received: 30 May 2025

Accepted: 17 November 2025

One contribution of 15 to a discussion meeting
issue ‘Symbolic regression in the physical
sciences’.

Subject Areas:

cosmology, astrophysics, artificial intelligence

Keywords:

symbolic regression, large-scale structure,
cosmological parameters, machine learning,
power spectrum, weak lensing

Author for correspondence:

Deaglan Bartlett

e-mail: deaglanbartlett@gmail.com

Symbolic emulators for cosmology: accelerating cosmological analyses without sacrificing precision

Deaglan Bartlett^{1,2} and Shivam Pandey³

¹Department of Astrophysics, University of Oxford, Oxford, UK

²Institut d’Astrophysique de Paris, Paris, France

³Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD, USA

DB, 0000-0001-9426-7723

In cosmology, emulators play a crucial role by providing fast and accurate predictions of complex physical models, enabling efficient exploration of high-dimensional parameter spaces that would be computationally prohibitive with direct numerical simulations. Symbolic emulators have emerged as promising alternatives to numerical approaches, delivering comparable accuracy with significantly faster evaluation times. While previous symbolic emulators were limited to relatively narrow prior ranges, we expand these to cover the parameter space relevant for current cosmological analyses. We introduce approximations to hypergeometric functions used for the Λ cold dark matter (Λ CDM) comoving distance and linear growth factor which are accurate to better than 0.001% and 0.05%, respectively, for all redshifts and for $\Omega_m \in [0.1, 0.5]$. We show that integrating symbolic emulators into a Dark Energy Survey Year 1 (DES-Y1)-like 3×2 pt analysis produces cosmological constraints consistent with those obtained using standard numerical methods. Our symbolic emulators offer substantial improvements in speed and memory usage, demonstrating their practical potential for scalable, likelihood-based inference.

This article is part of the discussion meeting issue ‘Symbolic regression in the physical sciences’.

1. Introduction

The aim of any cosmological analysis is to infer parameters which describe our Universe given observed data. This involves computing a prediction for an observable quantity given some set of cosmological parameters and then comparing these to the observed values. One must repeat this procedure typically thousands of times as one searches through cosmological parameters which could reasonably explain the observations.

These predictions can be prohibitively slow, and thus it has become increasingly common to replace exact simulations with surrogate models which are much faster yet are sufficiently accurate to not bias our predictions. These have conventionally taken the form of numerical machine learning emulators, such as AEMULUS [1], BACCO [2–4], COBRA [5], COSMICEMU [6], COSMOPOWER [7,8], DARK EMULATOR [9], EMUPK [10], EUCLIDEMULATOR₁ [11], EUCLIDEMULATOR₂ [12], FOFFITTINGFUNCTION [13], FRANKENEMU [14,15], NGENHALOFIT [16] and PICO [17,18].

Although much faster than solving the full simulation, it is hard to interpret what the emulators are doing ‘under the hood’ and, to continue using these into the future, one must rely on the programming languages and packages which they are written with being maintained and commonly used. As an alternative, one can often approximate the output of these simulations with symbolic models. Such methods were common long before their numerical counterparts, of which the BBKS [19] and Eisenstein & Hu (E&H) [20,21] approximations for the matter power spectrum are particularly popular. Historically, these expressions were obtained through careful, manual, analytic considerations of the outputs of simulations and were able to achieve accuracies of the order of a few to tens of per cent. More recently, the supervised machine learning technique of symbolic regression [22] has been utilized to extract symbolic approximations for cosmological quantities of interest in an automated fashion [23–28]. Not only are these more interpretable than the numerical emulators, but they can be made just as accurate yet orders of magnitude faster [23–25]. Moreover, the output is a symbolic equation using common mathematical expressions; these formulae are easily portable and implemented in any language of choice and do not depend on any particular package. It is possible to analytically differentiate such expressions, which improves efficiency for high-dimensional sampling algorithms. Alternatively, they can be easily imported into programming languages which have automatic differentiation (e.g. JAX) to achieve the same goal.

The goals of this work are threefold. First, we aim to extend the prior range of symbolic emulators to match the broader parameter ranges currently used in cosmological analyses, while maintaining the same level of accuracy. It is crucial that these emulators are accurate to at least the per cent level to be useful for current and upcoming surveys [29]. These are developed in §2. Second, we demonstrate that these extended emulators yield consistent cosmological constraints when embedded within a Dark Energy Survey Year 1 (DES-Y1)-like 3×2 pt analysis. This is the first demonstration that solely symbolic emulators can be used in Stage-III survey applications without causing biases; we leave the study of their ability to be applied to Stage-IV surveys to future work. Finally, we quantify the practical advantages of the symbolic approach by highlighting improvements in computational speed and memory efficiency. This analysis is performed in §3. For simplicity, we focus on Λ cold dark matter (Λ CDM) cosmologies in this work. It is possible to extend this to other cosmological models, such as with massive neutrinos and w_0w_a —with existing tools, as done in previous studies [25]. We conclude in §4.

Throughout, we use ‘log’ to denote the natural logarithm. Whenever the wavenumber, k , appears in a mathematical expression, it is assumed to have units of $h\text{Mpc}^{-1}$.

2. Symbolic emulators for Λ CDM

The aim of this section is to extend some of the SYmbolic Regression ENhanced (SYREN) [23–25] emulators so that they are valid over a wider parameter range than in the original papers. For

Table 1. Prior range used to generate symbolic emulators in this work.

parameter	description	min value	max value
Ω_m	present-day matter density parameter	0.1	0.5
Ω_b	present-day baryon density parameter	0.03	0.07
h	hubble constant ($H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$)	0.5	0.9
n_s	scalar spectral index	0.8	1.2
σ_8	matter fluctuation amplitude	0.6	1.0
z	redshift	0.0	3.0
k	wavenumber ($h \text{ Mpc}^{-1}$)	10^{-4}	10^2

simplicity, we focus on the concordance cosmological model (Λ CDM) which is governed by six key parameters: the present-day baryon and total matter densities (Ω_b and Ω_m , respectively), the Hubble constant ($H_0 = 100 h \text{ km/s/Mpc}$), the reionization optical depth, τ and two parameters describing the amplitude, A_s , and spectral tilt, n_s , of the primordial power spectrum. In this work, we ignore the small effect of τ on the quantities of interest and often use an alternative parametrization where we replace A_s with σ_8 , which is the root mean squared fluctuation of the linearly evolved density field when smoothed with a top-hat filter of radius $8 h^{-1} \text{ Mpc}$ (although we provide an emulator for converting between the two parametrizations). For these parameters, we use the prior range listed in [table 1](#), since this is more appropriate for current cosmological analyses. Given that we are particularly focusing on weak-lensing analyses, we fit our models for redshifts $z \leq 3$, although we demonstrate below that our models have excellent extrapolation behaviour beyond this range. We also produce some new emulators for quantities which are required for a weak lensing analysis and whose evaluation can be slow, namely, the comoving distance and linear growth factor.

In the remainder of this section, we begin by describing the symbolic regression code used (§2a) before outlining the quantities we wish to emulate and our found symbolic approximations (§2b). We assess the performance of our emulators at the level of the nonlinear matter power spectrum in §2c. We note that, owing to factorizing out their dependence, all the emulators developed in this section can be applied to an arbitrary σ_8 or A_s except for the HALOFIT ones, provided the other parameters do not go outside of the ranges listed in [table 1](#).

(a) Symbolic regression

Our symbolic emulators are obtained using the genetic programming-based symbolic regression code OPERON [30]. Genetic programming is an evolutionary algorithm that operates on symbolic expressions represented as trees. In each generation, a population of candidate expressions is evaluated using a fitness function, and less effective expressions are removed. New candidates are generated through crossover—recombining parts of existing expression trees—and mutation, which involves randomly altering tree structures by inserting, deleting or replacing subtrees. Through successive generations, the expression population is expected to improve, yielding increasingly accurate analytic models. In an OPERON, the leaf nodes of expressions (a variable or a parameter) are paired with associated scaling factors. These scaling coefficients are subsequently optimized using the Levenberg–Marquardt algorithm [31,32], as described in [33].

Our Pareto-optimization problem attempts to find an accurate fitting function while remaining relatively simple. During the search, we measure simplicity as the model ‘length’, which is equal to the number of nodes in the tree representation of the function, and accuracy with the root mean

squared error (RMSE). Unless otherwise specified, we seek equations which can be formed from the following symbols: +, −, ×, ÷, √, pow and −, as well as constants and variables. OPERON incorporates ϵ -dominance during the non-dominated sorting stage [34], where two candidate solutions are considered equivalent if their objective values—accuracy and simplicity—differ by less than a specified threshold ϵ . This approach limits the number of near-identical individuals in the population and supports convergence towards a diverse approximation of the Pareto front: the set of solutions for which no improvement in accuracy is possible without increasing complexity.

To balance predictive performance with model compactness, we tested several ϵ values and reported the selected settings in the relevant sections for each emulator. Genetic programming-based symbolic regression is a non-deterministic method, and thus identical set-ups but with different initial populations can lead to very different expressions and loss functions. As such, the variation in our results with different values of ϵ could also be due to different initial populations. Nonetheless, our final chosen settings and initial populations yield sufficiently accurate expressions for our purpose, although the reader could obtain different expressions (of better or worse accuracy) by rerunning our analysis with a different initial random seed, and even using the same training data.

To select our models, we manually inspect all equations returned on the Pareto front after a specified run time. We reject all equations which we deem insufficiently accurate for our purpose or whose losses on the training and validation sets are significantly different; an indication of overfitting. The selection of our preferred model is then somewhat qualitative based on the overall structure of the model (although note that this can be quantified using language models [35]), but we also only select models which have (or can be forced to have) the correct analytic limits in known cases.

(b) Emulators

(i) Radial comoving distance

In Λ CDM, the radial comoving distance to an object at scale factor $a \equiv 1/(1+z)$ is

$$\chi(a) = \frac{c}{H_0} \int_a^1 \frac{da'}{a'^2 \sqrt{\Omega_m a'^{-3} + 1 - \Omega_m}} = \frac{c}{H_0} \left[\frac{2a'}{\Omega_m} {}_2F_1 \left(\frac{2}{3}, 1, \frac{7}{6}; x' \right) \sqrt{\Omega_m a'^{-3} + 1 - \Omega_m} \right]_{a'=a}^1, \quad (2.1)$$

where $x' \equiv a'^3(\Omega_m - 1)/\Omega_m$ and ${}_2F_1$ is the Gaussian hypergeometric function. We can interpret $-x'$ as the ratio of the energy density of dark energy to that of matter at the scale factor a' . Although this has a closed-form solution, this is not always practical for modern computational workflows. In particular, libraries such as JAX.SCIPY (at the time of writing) do not provide built-in support for the hypergeometric function in equation (2.1), and its direct evaluation can be computationally inefficient. This presents a challenge for applications that rely on automatic differentiation, such as those using JAX, where efficient and differentiable expressions are preferred. We therefore wish to find a simpler analytic approximation for $\chi(a)$.

Given equation (2.1), we only need to find an approximation for the hypergeometric function part, since the other terms are trivial to evaluate. To obtain an approximation for the prior range of interest (table 1), we need to have an approximation which is valid for the range $-9.0 \leq x' \leq -0.016$. We therefore generated 200 x points randomly from a uniform distribution in this range for training, and another 200 for validation, and evaluated the hypergeometric function using SCIPY. We search for expressions with a maximum length of 25 and terminate OPERON's search after 2 minutes on 40 cores. We choose $\epsilon = 10^{-5}$ and add the analytic quotient operator ($\text{aq}(x, y) \equiv x/\sqrt{1+y^2}$) to the default basis set outlined in §2a.




Figure 1. Fractional errors on our approximations to the hypergeometric functions required to evaluate the radial comoving distance (left, equation (2.2)) and linear growth factor (right, equation (2.4)) for a Λ CDM cosmology. In both cases, $x \equiv a^3(\Omega_m - 1)/\Omega_m$, for scale factor a and present-day matter density parameter Ω_m . The approximations were obtained for x in the grey shaded region (corresponding to the prior range in table 1), and they thus show very good extrapolation behaviour, particularly for values of x relevant for cosmology (red hatched region).

The longest model which OPERON returns with these settings is at length 15, where our approximation to ${}_2F_1\left(\frac{2}{3}, 1, \frac{7}{6}; x\right)$ is proportional to $(c_0 + (c_1 - x)^{c_2} - c_3x)^{-c_4}$, where $\{c_i\}$ are optimized parameters. Given its relative simplicity and that it is the most accurate model found, we choose this expression. OPERON returns a constant of proportionality; however, we choose not to use this value but instead explicitly enforce that ${}_2F_1\left(\frac{2}{3}, 1, \frac{7}{6}; 0\right) = 1$. This gives

$${}_2F_1\left(\frac{2}{3}, 1, \frac{7}{6}; x\right) \approx \left(\frac{c_0 + c_1^{c_2}}{c_0 + (c_1 - x)^{c_2} - c_3x} \right)^{c_4}, \quad (2.2)$$

for $\{c_i\} = \{0.9207, 0.98617, 1.42499, 0.91875, 0.46516\}$.

We see that this is an exceptionally good approximation in figure 1, where we plot the fractional error on this approximation as a function of $-x$. When considering the range of the training data, the fractional error monotonically increases with $|x|$, from approximately 10^{-7} for small $|x|$ to 10^{-5} at the largest value of x , and thus the error is negligible. In figure 1, we also show the extrapolation behaviour of this function. Given that we only consider positive redshifts in cosmology, we are only concerned with the extrapolation to smaller values of $|x|$. Remarkably, the error continues to decrease as we consider higher redshifts, and thus for all $0.1 < \Omega_m < 0.5$ and for any redshift, we always have errors smaller than 0.001% with this approximation. The error slightly increases if we consider the other (irrelevant) limit for extrapolation, but still remains sub(per cent) until we consider values of x which are 10^4 times larger than the maximum value in our training range.

(ii) Growth factor

We now turn to the linear growth factor. This function gives the time dependence of perturbations at linear order in cosmology and is thus of fundamental importance. The growth factor in a Λ CDM universe is given by [36,37]

$$D(z) = {}_2F_1\left(\frac{1}{3}, 1, \frac{11}{6}; x\right) a, \quad (2.3)$$

where $x \equiv a^3(\Omega_m - 1)/\Omega_m$. We search for an analytic approximation to this, again using 200 training and validation points in the range used for the emulator of $\chi(a)$ and with the same basis functions as in the preceding section. After only 10 s of searching, operon finds a model at length 11 which is proportional to $((b_0 - x)^{b_2} + b_1)^{-1/2}$, for some constants $\{b_i\}$. Although OPERON gives the constant of proportionality, we choose to explicitly enforce that ${}_2F_1\left(\frac{1}{3}, 1, \frac{11}{6}; 0\right) = 1$, which corresponds to ensuring $D(a) \approx a$ at early times. We also know that, for large but negative $|x|$, the hypergeometric function is proportional to $(-x)^{-1/3}$, which sets the value of b_2 to be $2/3$.

After making these adjustments, we re-optimize the parameters by minimizing the RMSE across our training set, yielding the final expression,

$${}_2F_1\left(\frac{1}{3}, 1, \frac{11}{6}; x\right) \approx \frac{\sqrt{b_0^{2/3} + b_1}}{\sqrt{(b_0 - x)^{2/3} + b_1}}, \quad (2.4)$$

where $\{b_i\} = \{0.723, 1.204\}$.

We plot the fractional error on this approximation as a function of $|x|$ in [figure 1](#), where we again see excellent results. For the range of the training data (i.e. for the range of parameters in [table 1](#)), the fractional error on this expression is always below 5×10^{-4} , so is once again negligible. The extrapolation behaviour is the same as before: this approximation improves as one extrapolates to higher redshift, and thus the error is negligible for all values of x relevant to cosmology. For the other limit, the error increases, but stays at approximately per cent level even for much larger values of x than [equation \(2.2\)](#).

(iii) Conversion between A_s and σ_8

Our remaining emulators are for functions and variables which describe the clustering of matter in the Universe. The matter density of the Universe, $\rho(x)$, can be factorized into a spatially constant background density, $\bar{\rho}$, and a density contrast $\delta(x)$, such that $\rho(x) \equiv \bar{\rho}(1 + \delta(x))$. We then define the Fourier transform of $\delta(x)$ to be $\tilde{\delta}(k)$ and assume that the distribution of matter is statistically homogeneous and isotropic. Defining $\langle \dots \rangle$ to be an ensemble average, the matter power spectrum, $P(k)$, is defined to be

$$(2\pi)^3 P(k, \theta) \delta^D(k - k') \equiv \langle \tilde{\delta}(k) \tilde{\delta}^*(k') \rangle, \quad (2.5)$$

where δ^D is the Dirac delta function. If the density contrast is predicted according to linear perturbation theory, then we call this the linear matter power spectrum, $P_{\text{lin}}(k, \theta)$.

An important quantity is the root mean squared fluctuations of these linearly evolved fluctuations within spheres of radius R . This is given by

$$\sigma_R^2(z, \theta) = \int_0^\infty dk \frac{k^2}{2\pi^2} P_{\text{lin}}(k, z, \theta) |W(k, R)|^2, \quad (2.6)$$

where $W(k, R)$ is the Fourier transform of the top-hat window function and is given by

$$W(k, R) = \frac{3}{(kR)^3} (\sin(kR) - kR \cos(kR)). \quad (2.7)$$

It is common in cosmology to replace the parameter A_s by σ_8 , which is simply σ_R for $R = 8 h^{-1}$ Mpc at $z = 0$. However, if we want to change from one parametrization to another, it is convenient to have a simple parametric expression for this conversion rather than having to evaluate and integrate the linear power spectrum.

Since the linear power spectrum is proportional to A_s , we know that $\sqrt{A_s}/\sigma_8$ can only depend on the other cosmological parameters: $\Omega_m, \Omega_b, h, n_s$. We therefore find an approximation for $\log(\sqrt{10^9 A_s}/\sigma_8)$ as a function of these parameters, where we multiply A_s by 10^9 to obtain a target of order 1, and take the logarithm so that we always predict a positive σ_8 for a given A_s . To find this function, we generate 200 training cosmologies using a Latin hypercube and the parameter prior in [table 1](#). For a given A_s , we then find the corresponding σ_8 using CAMB [38] and compute $\log(\sqrt{10^9 A_s}/\sigma_8)$. We repeat this process for a further 200 cosmologies from a different Latin hypercube, which we use as our validation set. We search for expressions with OPERON using the default function set from §2a with $\epsilon = 10^{-6}$ and with a maximum allowed model length of 50.

Table 2. Best-fit parameters for the emulator converting A_s to σ_8 (equation (2.8)).

name	value	name	value	name	value	name	value	name	value
d_0	0.95534	d_1	68.8078	d_2	0.5159	d_3	1.18861	d_4	0.197
d_5	0.53884	d_6	0.01983	d_7	0.76405	d_8	0.29247	d_9	10.8834
d_{10}	0.73004	d_{11}	1.20497	d_{12}	0.75788	d_{13}	2.14175	d_{14}	3.03762
d_{15}	4.71485	d_{16}	5.46729	d_{17}	0.9624				

After 5 min of searching on a single node of 40 cores, we find that the loss for the Pareto front plateaus at a model length of 42, and thus choose this model, which is

$$\log\left(\frac{\sqrt{10^9 A_s}}{\sigma_8}\right) \approx d_0 n_s (d_1 \Omega_m)^{-d_2 h} + (d_3 \sqrt{h} + d_4 n_s) \left(-d_5 n_s + \frac{\Omega_b - d_6}{d_7 \Omega_m - d_8 \Omega_b}\right) + (d_9 \Omega_m)^{-d_{10} h} + \frac{(d_{11} \Omega_m)^{d_{12} h} (d_{13} \Omega_b + (d_{14} h)^{-d_{15} \Omega_m})}{d_{16} \Omega_m + d_{17} \Omega_b}, \quad (2.8)$$

where the parameters $\{d_i\}$ are listed in table 2. This model is very accurate, with an RMSE of 1.3×10^{-3} and 1.6×10^{-3} on the training and validation sets, respectively. Given that we often seek per cent level accurate predictions for the power spectrum, this is far within the required accuracy, since the fractional error on the power spectrum is approximately twice the error on the fit in equation (2.8) (since we have $\sqrt{A_s}$ on the left-hand side of equation (2.8)). This error is comparable to the result of [23] despite the larger prior range considered here.

(iv) Linear matter power spectrum

One of the most costly parts of the analysis that we perform in §3 is evaluating the linear matter power spectrum, and thus an emulator for this is the most effective for speeding up the analysis. This acceleration is not limited to weak lensing analyses; for example, an effective field theory-based analysis of galaxy clustering also requires the linear matter power spectrum, so this could be easily incorporated into such an analysis.

As in previous work [25,39], to obtain an approximation for this, we begin with the approximation from E&H [20] without the baryonic acoustic oscillations (BAOs) ($P_{\text{no wiggle}}(k, \theta)$), and learn the ratio of the true linear power spectrum (evaluated using CAMB) to this,

$$P_{\text{lin}}(k, z, \theta) = D^2(z, \theta) P_{\text{no wiggle}}(k, \theta) F(k, \theta). \quad (2.9)$$

As such, $F(k, \theta)$ fully captures the BAOs, but also provides a correction to the no-wiggle part of the E&H prediction. We choose to learn $\log F$ to enable us to optimize the fractional error and to ensure that our result for F is always positive. Since we are dealing with Λ CDM, F is independent of redshift [40]. We also know that it is independent of A_s since its dependence is fully captured by $P_{\text{no wiggle}}$, and thus only fit this ratio as a function of Ω_m , Ω_b , h , n_s and k .

To perform this fit, we sample 200 cosmologies using a Latin hypercube in the range listed in table 1. We use a further 200 for validation. For each cosmology, we evaluate the matter power spectrum at 200 k values, although these are not evenly spaced in k or $\log k$. We wish to have high resolution near the oscillatory features of the power spectrum (BAOs) and thus use 150 values of k logarithmically spaced between $10^{-2} h \text{ Mpc}^{-1}$ and $1 h \text{ Mpc}^{-1}$. The remaining 50 are used in the rest of the k range listed in table 1 such that the logarithmic spacing is constant for those two regions. Outside the high-resolution region, $\log F$ varies slowly; hence, we are able to use fewer points.




Figure 2. *Left:* Fractional error on symbolic fits to the redshift-zero linear matter power spectrum when compared against CAMB. We plot the 68% error distributions when the cosmological parameters are varied across the range listed in table 1. *Right:* The various contributions to the prediction for $\log F$ at the Planck 2018 cosmology. The high-frequency oscillations of $\cos(f_3)$ are a plotting artefact from using a logarithmic x scale: for large k , f_3 is linear in k .

This is our most challenging emulator to fit, and thus we choose a significantly longer run time for the OPERON of 47 h on 128 cores (after 24 h our results were slightly too inaccurate for our purposes). For numerical stability, we replace division in our function set with aq (again defined to be $aq(x, y) \equiv x/\sqrt{1+y^2}$), and we also include cosine since we know that we shall have to fit oscillatory features. We set $\epsilon = 10^{-4}$ and allow a maximum model length of 200.

We find that the OPERON produces a particularly compact result at a model length of 167, which takes the form

$$\log F(k, \theta) \approx f_1 (f_2 \cos(f_3) + f_4) + f_5, \quad (2.10)$$

where each $\{f_i\}$ are functions of k and θ , except f_5 which only depends on θ . These are given in appendix A. We note that the form of each of $\{f_i\}$ are non-trivial functions of k and the cosmological parameters. As such, although we are able to interpret equation (2.10) in various limits below, this is not generally true for arbitrary k from just looking at the equations. Nonetheless, this fit is accurate and, once the results are plotted, we can begin to see how each of the terms in equation (2.10) corrects the E&H formula.

OPERON returns a constant for f_5 ; however, we choose not to use this since we must obey a physical constraint. Since the E&H transfer function is unity on large scales, we require that $\log F \rightarrow 0$ as $k \rightarrow 0$. We therefore analytically compute $\tilde{f}_i \equiv \lim_{k \rightarrow 0} f_i$ and, to obey the physical constraint, enforce that

$$f_5 = -\tilde{f}_1 (\tilde{f}_2 \cos(\tilde{f}_3) + \tilde{f}_4). \quad (2.11)$$

It is not guaranteed *a priori* nor forced in the search that the limits of these expressions as $k \rightarrow 0$ are well behaved; however, for our chosen model, each f_i is finite in this limit. We find that $\tilde{f}_2 = 1$ whereas the other \tilde{f}_i depend on the cosmological parameters.

We compute the error on this fit compared to CAMB as a function of k and averaged across the prior range of table 1, and plot the results in figure 2. We compare the results to the errors on the E&H formula which includes the wiggles [21] as implemented in JAX-COSMO [41]. For both cases, we only use the transfer function with the true A_s rather than integrating and changing the normalization so that σ_8 matches. One immediately sees that the approximation obtained here is much more accurate than that of E&H; when averaged across all cosmologies and wavenumbers of table 1, we obtain an RMSE of 0.6% error, compared to 3% for E&H. Moreover, from figure 2 we see that the error on our fit is relatively constant with wavenumber, but is a strong function of k for E&H; our approximation's error band stays within 1% for all k , whereas the E&H result grows to up to 10% error on the smallest scales.

In figure 2, we also illustrate the various components contributing to the symbolic fit, using the Planck 2018 cosmological parameters. Each term in the model has a distinct interpretative role. The function f_1 accounts for the large- k inaccuracies of the E&H fitting formula, which arise owing to it

neglecting baryon pressure effects [20] on small scales. The terms f_2 and f_4 act as approximate step functions that effectively localize the oscillatory features to a finite range in k , suppressing them outside this window. The oscillatory component $\cos(f_3)$ behaves approximately as $\cos(\omega k + \phi)$, where ω and ϕ are asymptotically constant in the small- and large- k limits, with smooth corrections in the intermediate regime. Finally, f_5 enters the expression as a constant offset to ensure the correct limit as $k \rightarrow 0$.

(v) Halofit variables

Up until now, we have only considered the linear evolution of matter in the Universe. Although a good approximation on large scales, for smaller scales, one cannot neglect the nonlinear effects of gravity, and thus one must correct the matter power spectrum. This has the effect of distorting the BAOs and introducing a non-trivial time dependence, such that $P(k)$ is larger on small scales and the ratio of the nonlinear to linear power spectrum on these scales grows with time.

The most accurate way to capture this behaviour is through running N -body simulations; however, these are far too computationally expensive to be used within an inference pipeline, and thus surrogate models are required to estimate the effects of nonlinearity on $P(k)$. Many numerical emulators based on these have been created [2,4,11,12], and the SYREN-HALOFIT [24] and SYREN-NEW [25] symbolic emulators have been shown to demonstrate similar accuracy to these methods but at a fraction of the computational cost per evaluation. Unfortunately, these emulators were trained on N -body simulations with prior ranges smaller than that of table 1, and thus in this work, we shall not consider these emulators further.

Instead, we wish to emulate a semi-analytic prescription of nonlinear physics, namely, HALOFIT [42–44]. In this work, we use the commonly used parameters of [44], which are optimized to match the results of 16 N -body simulations around the Wilkinson Microwave Anisotropy Probe cosmologies [45], although we note an updated set of parameters is presented in [24].

HALOFIT gives the nonlinear matter power spectrum as a function of the linear spectrum, the cosmological parameters, as well as three derived variables computed from the linear spectrum. The first of these is the nonlinear scale, k_σ , which is defined to be the wavenumber at which

$$\sigma_G^2(k_\sigma^{-1}) \equiv 1, \quad \sigma_G^2(R) \equiv \int \Delta_{\text{lin}}^2(k) \exp(-k^2 R^2) \, d \log k, \quad (2.12)$$

for

$$\Delta_{\text{lin}}^2(k) \equiv P_{\text{lin}}(k) \frac{k^3}{2\pi^2}, \quad (2.13)$$

and $P_{\text{lin}}(k)$ is the linear matter power spectrum at the given redshift. Once we have k_σ , we also define

$$n_{\text{eff}} + 3 \equiv - \left. \frac{d \log \sigma_G^2(R)}{d \log R} \right|_{\sigma_G=1}, \quad C \equiv - \left. \frac{d^2 \log \sigma_G^2(R)}{d \log R^2} \right|_{\sigma_G=1}, \quad (2.14)$$

where the derivatives can be computed by fitting splines to $\sigma_G^2(R)$.

Obtaining these variables can be slow to evaluate since one needs to perform integrals of the linear power spectrum, run a root-finding algorithm, fit splines and then take derivatives of these. As such, the SYREN-HALOFIT emulators [24] find simple analytic expressions for k_σ , n_{eff} and C as a function of cosmological parameters and redshift. In this section, we extend these emulators to the prior of table 1. Given the wider prior range than SYREN-HALOFIT, we require more training samples than in that work; we use 1000 for each variable rather than 200, with parameters sampled uniformly on a Latin hypercube given the prior range of table 1. We compute k_σ by running a root-finding algorithm to find the value of R which obeys equation (2.12), where we compute the linear power spectrum with CAMB. Once this is obtained, we find $\sigma_G^2(R)$ with a fifth-order smoothing spline and take derivatives of this to compute n_{eff} and C . To approximate these variables, we




Figure 4. Fractional error on symbolic fits to the HALOFIT approximation to the nonlinear matter power spectrum when compared against CAMB’s implementation of HALOFIT. We plot the 68% and 95% error distributions when the cosmological parameters are varied across the range listed in table 1, and in each panel, we choose randomly sampled redshifts in the ranges given by the titles. Our approximation utilizes the symbolic approximation for the linear matter power spectrum (equation (2.10)), the growth factor (equation (2.4)) and the HALOFIT variables (equations (A 1)–(A 3)).

three equal redshift ranges. We see that the errors on our predictions are relatively independent of z , with all three panels of figure 4 appearing very similar. We observe that the 68% error band is always within 1% for all k and z , and that the 95% band is often at or within this level. As such, we see that our emulators are able to produce per cent level accurate predictions for this quantity, even when combined in this way.

3. Mock cosmological analysis

Until now, we have focused on the performance of the individual symbolic emulators at the level of the quantities which they are emulating. Although we have seen good performance, it is essential to assess how well they perform in a cosmological inference; even if the accuracy levels appeared high in previous sections, if they lead to biased cosmological inferences, then they would be unusable. In this section, we perform this test, where we perform a mock analysis of a DES-Y1 lensing and clustering cosmological analysis [46,47]. We describe the set-up of our test in §3a, and the results are presented in §3b.

(a) Analysis set-up

To evaluate the performance of our emulator in a realistic setting, we embed it in a synthetic version of the DES-Y1 3×2 pt analysis. This framework combines three two-point correlation functions: cosmic shear (the auto-correlation of galaxy shapes), galaxy–galaxy lensing (the cross-correlation between galaxy positions and shapes) and galaxy clustering (the auto-correlation of galaxy positions). Together, these observables jointly constrain the underlying matter distribution and the galaxy–matter connection. We adopt the DES-Y1 set-up as a representative case study owing to its well-documented likelihood pipeline, realistic survey characteristics and relevance to current and future cosmological experiments.

If we consider two probes (either clustering or shear), labelled by (i, j) , the angular power spectrum C_ℓ^{ij} is given under the Limber approximation as [48]

$$C_\ell^{ij} \approx \left(\ell + \frac{1}{2}\right)^{\mu_i + \mu_j} \int_{a_{ij}}^1 \frac{da}{c^2 \chi^2} \frac{d\chi}{da} W_i(\chi) W_j(\chi) P\left(k = \frac{\ell + 1/2}{\chi}, z\right), \quad (3.1)$$

where $\mu_i = 0$ for galaxy clustering and $\mu_i = -2$ for weak lensing, and the integration range (specified by a_{ij}) is determined by the maximum redshift considered. The number density of tracers in

a given tomographic redshift bin is given by $W_i(z)$ which, for galaxy number counts, is assumed to be

$$W_i^g(z) = n_i(z)b(z)H(z), \quad (3.2)$$

where $n_i(z)$ is the distribution of redshifts for the sample and $b(z)$ is the bias function. For weak lensing, the form of $W_i(z)$ is

$$W_i^k(z) = \frac{3H_0^2\Omega_m}{2c^2} \left(\frac{(\ell+2)!}{(\ell-2)!} \right)^{1/2} (1+z)\chi(z) \int_z^\infty p_i(z') \frac{\chi(z') - \chi(z)}{\chi(z')} dz' + W_{\text{IA}}(z), \quad (3.3)$$

where $p_i(z) = n_i(z)/\bar{n}_i$ for $\bar{n}_i = \int_0^{z_{\text{max}}} n_i(z) dz$. We model the intrinsic alignment term, $K_{\text{IA}}(z)$, as [49]

$$W_{\text{IA}}(z) = \left(\frac{(\ell+2)!}{(\ell-2)!} \right)^{1/2} p_i(z)b(z)H(z) \frac{\Omega_m}{D(z)} A_{\text{IA}}(z), \quad (3.4)$$

where

$$A_{\text{IA}}(z) = A_{\text{IA}} \left(\frac{1+z}{1.62} \right)^{\eta_{\text{IA}}}, \quad (3.5)$$

gives the intrinsic alignment amplitude, and is parametrized by the free parameters A_{IA} and η_{IA} .

Our approach in this section broadly follows the analyses of [41,46,47]. We use the $n_i(z)$ relevant for the four source and five lens samples of DES-Y1. We assume that these may be slightly miscalibrated, and thus add in one free parameter for each sample, such that

$$n_i(z) = n_i^{\text{PZ}} (z - \Delta z^i), \quad (3.6)$$

where n_i^{PZ} is the initial estimate of $n_i(z)$. We consider a linear galaxy bias $b_i(z) = b_i$, which is also a free parameter, and a multiplicative bias for the weak lensing measurements, $\{m_i\}$. We list the priors for each of these parameters used in the inference, as well as for the cosmological parameters, in table 3. This corresponds to 25 parameters to infer (5 cosmological and 20 ‘nuisance’ parameters), representing a moderately high-dimensional inference problem.

For our ‘true’ inference which we compare against, we use the DES-Y1 likelihood implemented in COSMOSIS¹ [50]. For both our mock data and inference, we choose to generate data according to the Takahashi *et al.* [44] implementation of HALOFIT, which is computed using CAMB. We use the MULTINEST [51–53] sampler with 500 live points and allowing a maximum number of 50 000 iterations. We terminate the sampling when our evidence estimate reaches an error of 0.1, and we do not use the constant efficiency mode. The target efficiency for the importance nested sampling is 0.3. Otherwise, we use the default settings of COSMOSIS. We use a single node with 128 CPU cores to perform our inference.

We compare the inference results from COSMOSIS with an implementation of this likelihood written in JAX, for which we use many of the functionalities of JAX-COSMO [41]. To enable a fair comparison, we modify JAX-COSMO’s photometric redshift implementation from a kernel density estimator to a cubic spline, which is what is used by COSMOSIS. We found that the difference between these two interpolation methods was much greater (up to a few per cent) than the difference between the spectra produced by the symbolic emulators and that from using CAMB.

Utilizing gradients to improve efficiency, we choose to sample with a No-U-Turn Sampler [54], as implemented in NUMPYRO [55,56]. We run 16 parallel chains, each with 1000 warm-up steps and 10 000 samples, using a single Nvidia RTX 2080 GPU. For the SYREN runs, we are able to use a single node of 8 CPU cores with a combined memory of 30.3 GB. This is insufficient memory when we do not use the symbolic emulators, so we use two nodes for those cases (but with a single GPU still).

To compare our run time against that using a numerical emulator, we also run our inference using COSMOPower-JAX [8]. This emulator predicts a different approximation from HALOFIT, and

¹<https://cosmosis.readthedocs.io/en/latest/>.

Table 3. Priors used for the mock cosmological analysis. For ‘uniform’ priors, the arguments give the minimum and maximum allowed values. For ‘normal’ priors, the prior is a Gaussian distribution with a mean and standard deviation given by the first and second arguments, respectively.

parameter	prior	parameter	prior
cosmology		lens photo-z shift	
Ω_m	uniform (0.1, 0.5)	$\Delta z'_1$	normal (0, 0.007)
$10^{10} A_s$	uniform (5, 50)	$\Delta z'_2$	normal (0, 0.007)
n_s	uniform (0.8, 1.2)	$\Delta z'_3$	normal (0, 0.006)
Ω_b	uniform (0.03, 0.07)	$\Delta z'_4$	normal (0, 0.010)
h	uniform (0.5, 0.9)	$\Delta z'_5$	normal (0, 0.010)
lens galaxy bias		shear calibration	
$b_i (1 \leq i \leq 5)$	uniform (0.8, 3.0)	$m_i (1 \leq i \leq 4)$	normal (0, 0.023)
source photo-z shift		intrinsic alignment	
Δz^s_1	normal (0, 0.016)	A_{IA}	uniform (−5, 5)
Δz^s_2	normal (0, 0.013)	η_{IA}	uniform (−5, 5)
Δz^s_3	normal (0, 0.011)		
Δz^s_4	normal (0, 0.022)		

thus, we do not compare its results at the level of posteriors to that obtained with `COSMOSIS`, but we use it as a useful timing comparison.

(b) Results

The primary goal of an emulator (whether numerical or symbolic) is for increased computational speed. In our inference, there are two primary ways in which our symbolic emulators help in this regard. First, a single likelihood evaluation is much faster since evaluating $P(k)$ is much slower when using `CAMB` than our emulator. Second, given the differentiable nature of our emulators, we can gain sampling efficiency through the use of gradients. As such, we find that our inference takes far less time than using `COSMOSIS`. Using the settings outlined in the preceding section, the ‘exact’ inference required approximately 72 h of run time on 128 CPU cores, whereas our inference only took 1.2 h on a single GPU: approximately 60 times faster.

It is unsurprising that using our emulator is much faster than the `COSMOSIS` inference. However, an interesting result can be seen by comparing our emulator to `JAX-COSMO` and `COSMOPOWER`. For `JAX-COSMO`, we adapt the analysis of [41] for a Λ CDM cosmology, which notably involves applying `HALOFIT` to the E&H approximation. For the `COSMOPOWER` result, we replace E&H + `HALOFIT` with the `COSMOPOWER` emulator. Upon doing this, we find that `COSMOPOWER` requires 3.5 h of run-time, whereas `JAX-COSMO` requires 9.3 h. We therefore find that using all our emulators results in a speed-up of at least a factor of 3. We find that this increase in speed is dominated by the approximations for the comoving distance and growth factor, since we find more modest increases in speed of around 20% if we only replaced `COSMOPOWER` by our power spectrum emulator. Since these are our most accurate emulators (they are better than 0.001% and 0.05% accurate, respectively), even if one wished to use a neural network for estimating $P(k)$, using our approximations for these quantities would still result in a significant increase in speed since one no longer must numerically solve an integral or differential equation. We note that neither the neural network approach nor our symbolic emulators has been optimized for speed. As such, either or both of these

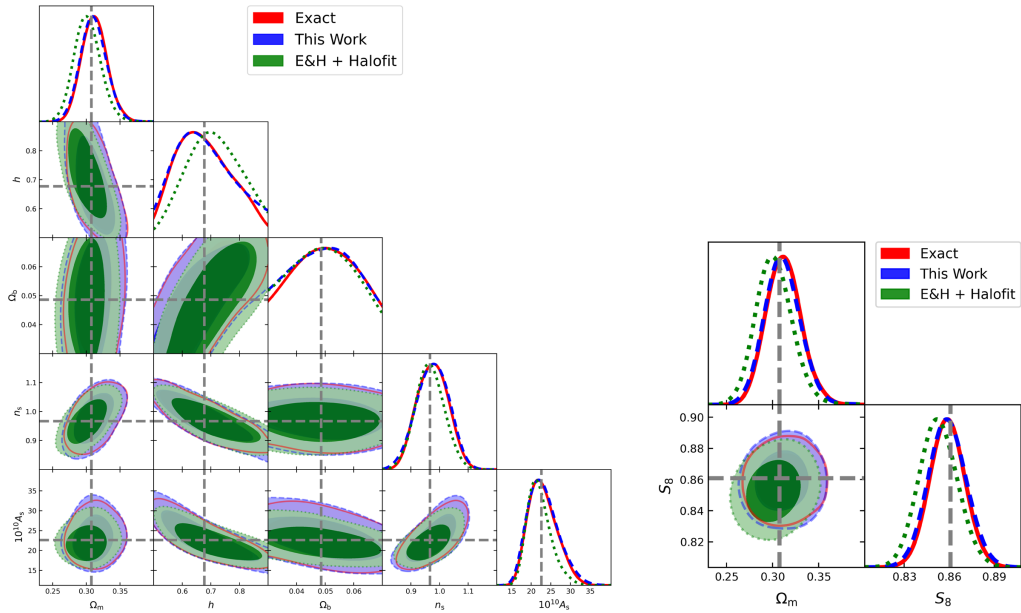


Figure 5. One- and two-dimensional posterior distributions of the cosmological parameters for our mock cosmological DES-Y1-like analysis. The red contours are obtained using an ‘exact’ model (CAMB) and are compared to those obtained using the symbolic approximations produced in this work (blue contours), and to the combination of E&H and HALOFIT, as implemented in JAX-COSMO (green contours). The grey dashed lines indicate the true parameters. In the left-hand panel, we consider the sampled cosmological parameters (which all have uniform priors), whereas in the right-hand panel, we also consider a derived parameter, S_8 . The high level of consistency between our contours and the exact model indicates that the emulators are sufficiently accurate for such an analysis. This is not true for the E&H+HALOFIT model.

emulation strategies could be made faster. Nonetheless, for currently available symbolic models and neural networks, we find the former to be faster.

One of the motivations for producing symbolic emulators for various parts of the pipeline, rather than just the linear power spectrum, is to reduce memory usage. In particular, obtaining the HALOFIT variables involves computing integrals, and calculating the growth factor requires solving a differentiable equation. To demonstrate the advantage of these emulators, we compare the memory requirements for computing the mean of the likelihood at the true parameters using a JAX-COSMO approach, compared to relying solely on our symbolic emulators. We find that the memory needed to evaluate the mean is comparable; however, evaluating the gradient of the JAX-COSMO pipeline with respect to the input parameters requires 50% more memory than for our symbolic emulators. This additional saving in memory could be exploited in a computational analysis by running more Markov chain Monte Carlo simulations in parallel per GPU.

Of course, accelerating an inference is only worth doing if the results do not change after doing so. To assess this, in figure 5, we plot the one- and two-dimensional posterior distributions for inferred cosmological parameters, plotted using GETDIST [57]. We see that the distributions inferred using COSMOSIS and our symbolic emulators are very similar, in particular, for the parameters that are well-constrained in these analyses: Ω_m , n_s and A_s . There is a small difference between the posteriors of h and Ω_b ; however, these are relatively unconstrained by such an analysis, and thus, this is not problematic. These results are in contrast to the combination of E&H and HALOFIT, with significant shifts in the peaks of the posterior of Ω_m and n_s , and an overly confident inference of A_s .

It is common to not only consider the cosmological parameters as listed in [table 1](#), but also consider the derived parameter S_8 , defined to be

$$S_8 \equiv \sigma_8 \sqrt{\frac{\Omega_m}{0.3}}, \quad (3.7)$$

since S_8 and Ω_m represent two approximately independent parameters which are constrained in such an analysis. Hence, in [figure 5](#), we also show the posteriors in the Ω_m - S_8 plane, as is conventional. Again, we see that the differences from the truth for the emulators developed in this work are negligible, and thus, the symbolic emulators are sufficiently accurate to be used to infer these parameters in such an analysis. Again, this is not true for E&H+HALOFIT, with an underestimation of both Ω_m and S_8 .

During our inference, as well as cosmological parameters, we also sampled parameters describing intrinsic alignments, bias, shear calibration errors and errors in the photometric redshift distributions. To determine whether there are any biases in these parameters, in [figure 6](#), we plot the two-dimensional marginalized posterior distribution of all parameters sampled in the Markov chain Monte Carlo simulation. As with the cosmological parameters, we find that the two sets of contours agree well, further demonstrating that the inference with symbolic emulators is accurate. Once again, we see that this is not the case if one uses the HALOFIT formalism but with the linear matter power spectrum from E&H. Hence, the corrections to E&H obtained in this work are necessary to achieve accurate posteriors.

Therefore, we conclude that the symbolic emulators developed in this work are sufficiently accurate to lead to accurate cosmological inference, but this is not true if one uses E&H for the linear part of the power spectrum.

4. Conclusion

Symbolic emulators offer a powerful alternative to traditional numerical surrogates for cosmological inference. By learning fast, interpretable approximations to expensive forward models, they enable efficient and transparent exploration of cosmological parameter space. In this work, we have extended the validity of symbolic emulators to broader parameter ranges relevant for current analyses, demonstrating that these models retain the accuracy needed for unbiased cosmological inference. We also introduced a new emulator for the hypergeometric function needed to compute the comoving distance in a Λ CDM cosmology, which was accurate to better than one part in 10^5 for all cosmological applications. The extrapolation behaviour of this emulator is particularly notable, as it becomes more accurate as one extrapolates beyond the range of the training data (in the direction relevant for cosmology).

We embedded these extended symbolic emulators within a DES Y1-like 3×2 pt likelihood pipeline and found that the resulting parameter constraints are consistent with those obtained using CAMB. Crucially, the symbolic models enable over an order-of-magnitude improvement in evaluation speed, making them especially attractive for applications where computational resources are a limiting factor. We find that our pipeline with all symbolic emulators is both faster and more memory efficient than if using a combination of numerical emulators for the power spectrum and solving for the growth factor numerically.

In this work, we focused on a Stage-III cosmological analysis as a first step towards deploying symbolic emulators in Stage-IV surveys. The errors on the emulators produced in this work are sub-dominant compared to the errors on the quantities they are emulating (HALOFIT is insufficiently accurate for Stage-IV surveys). Other symbolic emulators in the SYREN family overcome this problem by fitting more accurate power spectra [24,25] but with narrower prior ranges than considered here. We leave the extension of their predictions to wider priors and testing their use within Stage-IV likelihoods to future work.

As cosmological datasets continue to grow in precision and volume, the ability to perform fast and accurate inference will only become more essential. The results of this work show that symbolic emulators are not only a viable replacement for numerical ones, but may in fact be preferable in scenarios demanding interpretability, efficiency and portability. We expect symbolic techniques to play an increasingly central role in the future of cosmological data analysis.

Data accessibility. For the purposes of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. The data underlying this article will be shared on reasonable request to the corresponding authors. We provide implementations of our emulators at https://github.com/DeaglanBartlett/symbolic_pofk.

Declaration of AI use. Minor assistance with coding and plotting routines, as well as some suggested phrasings for the manuscript text based on authors' notes and draft sentences.

Authors' contributions. D.B.: conceptualization, data curation, investigation, methodology, visualization, writing—original draft, writing—review and editing; S.P.: data curation, methodology, software, writing—review and editing.

Both authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. The authors are supported by the Simons Collaboration on 'Learning the Universe'. D.B. is supported by Schmidt Sciences through The Eric and Wendy Schmidt AI in Science Fellowship. This work has made use of the Infinity Cluster hosted by Institut d'Astrophysique de Paris; we thank Stephane Rouberol for running this cluster smoothly for us.

Acknowledgements. We thank Harry Desmond, Pedro Ferreira, Lukas Kammerer, Gabriel Kronberger, Constantinos Skordis and Benjamin Wandelt for useful comments and suggestions. We thank Jonathan Patterson for smoothly running the Glamdring Cluster hosted by the University of Oxford, where some of the data processing was performed. For the purposes of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. The data underlying this article will be shared on reasonable request to the corresponding authors. We provide implementations of our emulators at https://github.com/DeaglanBartlett/symbolic_pofk.

Appendix A. Symbolic approximations for the linear matter power spectrum

In this section, we give the symbolic expressions for the various terms given in equation (2.10). After some manipulations to merge superfluous constants and rounding some which are close to 0 or 1, we obtain the following expressions (see figure 6):

$$f_1 = e_0 \left[k + e_1 \left(1 - e_2 \Omega_m^{-e_3} k \right) \right. \\ \left. \left(-\Omega_m + \frac{e_4 \Omega_b + e_5 k}{\sqrt{e_6 + (\Omega_b + e_7 k)^2}} - \frac{e_8 (e_9 \Omega_m)^{-e_{10} \Omega_b} k}{\sqrt{(e_{11} + h^2) (e_{12} k^2 + \Omega_m^2)}} + \frac{e_{13} (e_{14} \Omega_m)^{-e_{15} \Omega_b + e_{16} k} k}{\sqrt{e_{17} + (e_{18} \Omega_m)^{e_{19} h} (e_{20} \Omega_m - k)^2}} \right) \right. \\ \left. \left(e_{21} + \left(e_{22} + \frac{(e_{23} - e_{24} \Omega_b - k)^2}{e_{25} + (k - e_{26} \Omega_b)^2} \right)^{-1} \right)^{-\frac{1}{2}} \left(e_{27} + (e_{28} \Omega_b + e_{29} - k)^2 \right)^{-\frac{1}{2}} \right], \quad (\text{A } 1)$$

$$f_2 = \left(1 + \left(\frac{(e_{30} \Omega_m)^{-e_{31} \Omega_b} \cos(e_{32} k) - \cos(e_{33} k)}{\sqrt{1 + \cos^2(e_{34} k)}} + e_{35} \frac{(e_{36} \Omega_m)^{e_{37} h} (e_{38} k)^{-e_{39} \Omega_b}}{\sqrt{e_{40} + h^2}} k^2 \right)^2 \right)^{-\frac{1}{2}}, \quad (\text{A } 2)$$

$$f_3 = \left(\frac{e_{41} (e_{42} \Omega_m)^{e_{43} h}}{\sqrt{e_{44} + h^2}} - e_{45} (e_{46} \Omega_m)^{e_{47} \Omega_b h - e_{48} h^2} \right) k + \frac{e_{49}}{\sqrt{e_{50} + (\Omega_m e_{51} - k)^2}}, \quad (\text{A } 3)$$

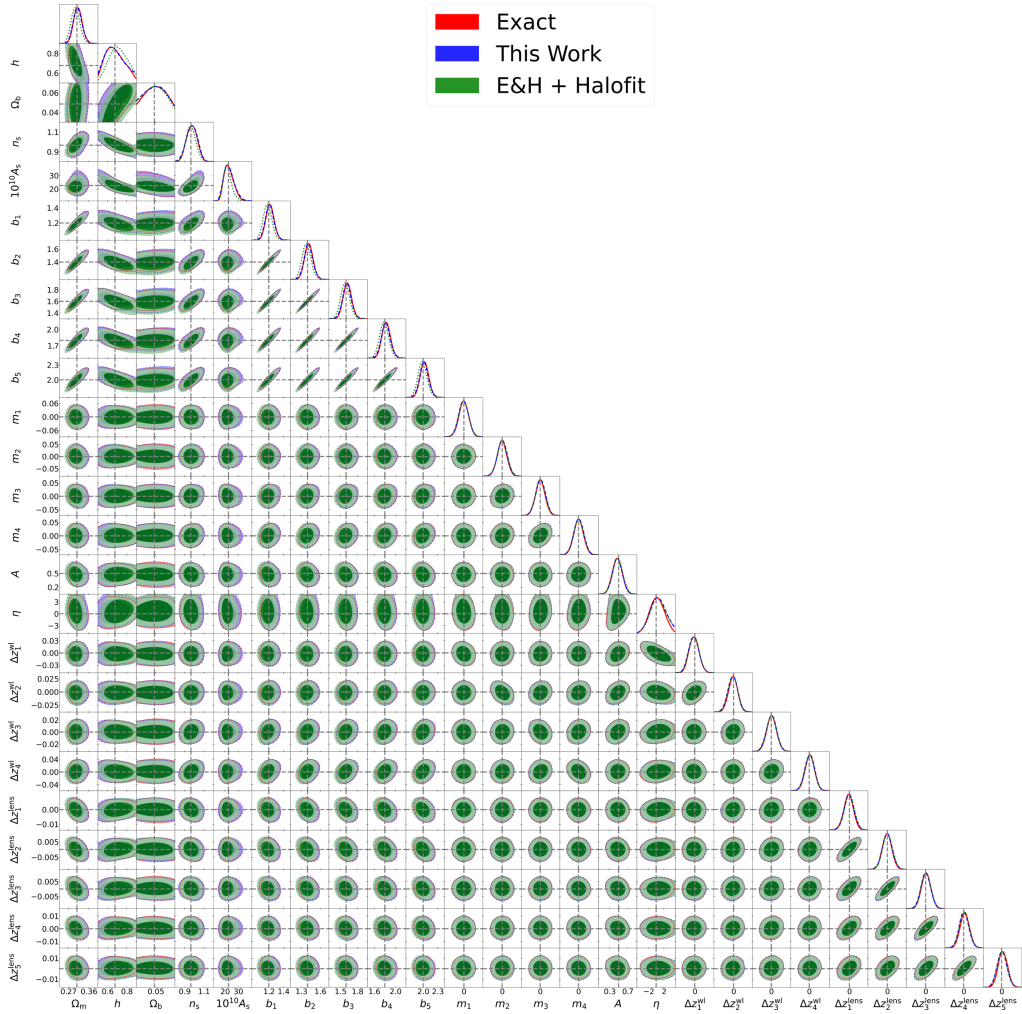


Figure 6. Marginalized posterior distributions of all parameters considered in our mock DES-Y1 analysis. The ‘exact’ model is computed using CAMB, and we see that its contours (red) are very similar to those obtained using our symbolic emulators (blue), demonstrating that the emulators are sufficiently accurate for such an analysis. This is in contrast to applying HALOFIT to the E&H approximation (green contours), for which significant discrepancies are seen.

$$f_4 = \left(\frac{e_{52} (e_{53} \Omega_m)^{e_{54} h} k}{\sqrt{e_{55} + (e_{56} \Omega_b - k)^2} (\Omega_b + e_{57} k)} + \frac{e_{58} k + e_{59}}{\sqrt{e_{60} + (e_{61} \Omega_m)^{-e_{62} h} k^2}} - e_{63} - \frac{e_{64}}{\sqrt{e_{65} + k^2}} \right) \left(1 + e_{66} \left(-\Omega_m + \frac{e_{67} \Omega_b + e_{68} \frac{\Omega_m}{\Omega_b} (e_{69} \Omega_m)^{-e_{70} h} + e_{71} k}{\sqrt{e_{72} + (k - e_{73})^2}} \right)^2 \right)^{-\frac{1}{2}}, \quad (\text{A } 4)$$

where the parameters $\{e_i\}$ are listed in [table 4](#).

Table 4. Best-fit parameters for the linear matter power spectrum emulator (equation (2.10)).

name	value	name	value	name	value	name	value
e_0	0.013924	e_1	5.1771	e_2	20.636	e_3	0.49092
e_4	3.4224	e_5	0.35621	e_6	0.016739	e_7	0.18401
e_8	0.58832	e_9	5.108	e_{10}	10.783	e_{11}	0.0043879
e_{12}	0.11547	e_{13}	1.3869	e_{14}	2.9301	e_{15}	32.014
e_{16}	0.002192	e_{17}	0.002926	e_{18}	316.2	e_{19}	1.2158
e_{20}	0.095822	e_{21}	0.074921	e_{22}	0.0067841	e_{23}	0.0093912
e_{24}	0.022678	e_{25}	9.3762×10^{-4}	e_{26}	0.64834	e_{27}	0.006882
e_{28}	0.6518	e_{29}	0.11855	e_{30}	2.6863	e_{31}	42.261
e_{32}	121.6	e_{33}	123.21	e_{34}	14.51	e_{35}	1378.1
e_{36}	22.375	e_{37}	1.0369	e_{38}	75.439	e_{39}	3.6282
e_{40}	0.15279	e_{41}	41.884	e_{42}	1052.3	e_{43}	0.22687
e_{44}	0.022003	e_{45}	236.72	e_{46}	0.16852	e_{47}	0.55815
e_{48}	0.040376	e_{49}	0.4352	e_{50}	0.011574	e_{51}	0.043423
e_{52}	0.001834	e_{53}	205.7	e_{54}	0.26165	e_{55}	0.0073544
e_{56}	0.99135	e_{57}	0.18444	e_{58}	0.089257	e_{59}	0.0074827
e_{60}	7.7431×10^{-5}	e_{61}	1.1788	e_{62}	0.44205	e_{63}	0.098723
e_{64}	0.01075	e_{65}	0.007492	e_{66}	0.82882	e_{67}	1.1882
e_{68}	383.17	e_{69}	397.51	e_{70}	4.4661	e_{71}	0.81134
e_{72}	0.0052959	e_{73}	0.15939				

Appendix B. Symbolic approximations for HALOFIT variables

In this section, we give our chosen analytic approximations to the HALOFIT variables defined in equations (2.12) and (2.14). For the nonlinear scale, k_σ , we choose

$$\begin{aligned}
 \log k_\sigma = & -\psi_0 \sigma_8 - (\psi_1 h)^{\psi_{2z}(\psi_3 \Omega_m)^{\psi_4 \sigma_8 - \psi_5 \sqrt{z}}} + \frac{\psi_6}{h} (\psi_7 z)^{\psi_8 n_s + \psi_9 \sigma_8 - \psi_{10} z} \\
 & + \left[-\psi_{11} z \left(-\psi_{12} \Omega_b + (\psi_{13} \Omega_m + \log(\psi_{14} n_s))^{-\psi_{15} z (\psi_{16} z)^{-\psi_{17} h}} \right) \right. \\
 & - \left(\psi_{18} \sigma_8 (\psi_{19} \Omega_m)^{-\psi_{20} z} - (\psi_{21} n_s)^{(-\psi_{22} n_s - \psi_{23} z) (\psi_{24} h)^{\psi_{25} z}} \right) \\
 & \left. \times \left(-\psi_{26} h + \psi_{27} n_s + (-\psi_{28} \Omega_b + \psi_{29} \Omega_m)^{-\psi_{30} \Omega_b - \psi_{31} \sigma_8} \right) \right] \\
 & \times \left[\psi_{32} \Omega_m + \psi_{33} \sigma_8 - \psi_{34} + (\psi_{35} n_s)^{\psi_{36} n_s + \psi_{37} z} + (-\psi_{38} \Omega_b + \psi_{39} \Omega_m)^{\psi_{40} \sigma_8} + (\psi_{41} n_s)^{\psi_{42} \sigma_8 (\psi_{43} h)^{-\psi_{44} z}} \right]^{-1},
 \end{aligned} \tag{B 1}$$

where the parameters $\{\psi_i\}$ are listed in table 5.

Table 5. Best-fit parameters for the emulator of the HALOFIT variable k_{σ} equation (A 1).

name	value	name	value	name	value	name	value	name	value
ψ_0	0.3458	ψ_1	0.01477	ψ_2	0.0825	ψ_3	4.642	ψ_4	0.4738
ψ_5	0.3847	ψ_6	2.005	ψ_7	0.02206	ψ_8	0.2958	ψ_9	0.4962
ψ_{10}	0.03355	ψ_{11}	0.6467	ψ_{12}	1.139	ψ_{13}	8.498	ψ_{14}	4.57
ψ_{15}	0.6448	ψ_{16}	0.1022	ψ_{17}	0.3782	ψ_{18}	1.239	ψ_{19}	18.27
ψ_{20}	0.1043	ψ_{21}	0.3435	ψ_{22}	0.2178	ψ_{23}	0.1644	ψ_{24}	0.2413
ψ_{25}	0.03482	ψ_{26}	0.5142	ψ_{27}	0.4983	ψ_{28}	1.1	ψ_{29}	0.8871
ψ_{30}	0.7256	ψ_{31}	0.1405	ψ_{32}	0.614	ψ_{33}	1.027	ψ_{34}	3.036
ψ_{35}	0.9132	ψ_{36}	0.4545	ψ_{37}	0.2444	ψ_{38}	113.5	ψ_{39}	97.35
ψ_{40}	0.1534	ψ_{41}	0.9639	ψ_{42}	1.309	ψ_{43}	1.616	ψ_{44}	0.3708

Table 6. Best-fit parameters for the emulator of the HALOFIT variable n_{eff} equation (A 2).

name	value	name	value	name	value	name	value	name	value
χ_0	3.3208	χ_1	6.3738	χ_2	0.2304	χ_3	0.1642	χ_4	0.064
χ_5	0.1461	χ_6	0.2171	χ_7	0.8835	χ_8	0.7457	χ_9	0.0537
χ_{10}	0.268	χ_{11}	6.4778	χ_{12}	2.3502	χ_{13}	1.3872	χ_{14}	0.6122
χ_{15}	0.8784	χ_{16}	0.6466	χ_{17}	512.827	χ_{18}	0.0894		

Table 7. Best-fit parameters for the emulator of the HALOFIT variable C equation (A 3).

name	value	name	value	name	value	name	value
φ_0	4.917	φ_1	0.04262	φ_2	1461	φ_3	2181
φ_4	11.15	φ_5	0.4784	φ_6	0.09069	φ_7	0.0343
φ_8	0.04317	φ_9	0.0372	φ_{10}	0.09107	φ_{11}	0.151
φ_{12}	0.04674	φ_{13}	0.04854	φ_{14}	0.05496	φ_{15}	0.03631
φ_{16}	0.1805	φ_{17}	0.1707	φ_{18}	0.2315	φ_{19}	0.4075
φ_{20}	0.593	φ_{21}	1.84	φ_{22}	1.028	φ_{23}	0.02645
φ_{24}	0.06507	φ_{25}	0.06477	φ_{26}	0.192	φ_{27}	0.003867
φ_{28}	5.56×10^{-4}	φ_{29}	8.51×10^{-4}	φ_{30}	1.77×10^{-4}	φ_{31}	0.03328
φ_{32}	0.04181	φ_{33}	0.06002				

Our approximation to the effective slope, n_{eff} , is

$$\begin{aligned}
 n_{\text{eff}} \approx & (\chi_0 \sqrt{n_s} - \chi_1) (\chi_2 n_s + \chi_3 \sigma_8 + \chi_4 z - \chi_5)^{\chi_6 \sigma_8} \\
 & \times \left(\chi_7 \sqrt{\chi_8 \Omega_m - \Omega_b} - \frac{\chi_9}{h} (\chi_{10} \sqrt{n_s} + (\chi_{11} \Omega_m)^{-\chi_{12} z - \chi_{13}})^{-\chi_{14} \Omega_m} \right) \\
 & \times (\chi_{15} h + \chi_{16} \sigma_8 + (\chi_{17} z)^{-\chi_{18} z}),
 \end{aligned} \tag{B 2}$$

where the optimized values of $\{\chi_i\}$ are listed in table 6.

Finally, our approximation to the curvature variable, C , is

$$\begin{aligned}
 C \approx & (\varphi_0 z)^{-\varphi_{1z}} \\
 & \times (\varphi_2 \Omega_m - \varphi_3 \Omega_b + \varphi_4 z)^{\varphi_5 + (\varphi_6 \Omega_m)^{-\varphi_7 h + \varphi_8 n_s - \varphi_9 \sigma_8} - (\varphi_{10} z + \varphi_{11} n_s)^{-\varphi_{12} \Omega_m - \varphi_{13} h - \varphi_{14} \sigma_8 - \varphi_{15} z}} \\
 & \times \left(-\varphi_{16} h - \varphi_{17} \sigma_8 + (\varphi_{18} \sigma_8)^{(\varphi_{19} \Omega_m - \varphi_{20} \Omega_b + \varphi_{21})^{-\varphi_{22} h}} (\varphi_{23} \Omega_m + \varphi_{24} h + \varphi_{25} n_s + \varphi_{26} \sigma_8 - \varphi_{27} z) \right. \\
 & \left. - (-\varphi_{28} \Omega_m + \varphi_{29} n_s + \varphi_{30} z)^{\varphi_{31} h + \varphi_{32} n_s + \varphi_{33} \sigma_8} \right), \tag{B 3}
 \end{aligned}$$

with the optimized coefficients $\{\varphi_i\}$ listed in table 7.

References

- Zhai Z, Tinker JL, Becker MR, DeRose J, Mao YY, McClintock T, McLaughlin S, Rozo E, Wechsler RH. 2019 The aemulus project. III. Emulation of the galaxy correlation function. *Astrophys. J.* **874**, 95. (doi:10.3847/1538-4357/ab0d7b)
- Angulo RE, Zennaro M, Contreras S, Aricò G, Pellejero-Ibañez M, Stücker J. 2021 The BACCO simulation project: exploiting the full power of large-scale structure for cosmology. *Mon. Not. R. Astron. Soc.* **507**, 5869–5881. (doi:10.1093/mnras/stab2018)
- Aricò G, Angulo R, Zennaro M. Accelerating large-scale-structure data analyses by emulating boltzmann solvers and lagrangian perturbation theory. *Open Res. Europe* **1**, 152. (doi:10.12688/openreseurope.14310.2)
- Zennaro M, Angulo RE, Pellejero-Ibañez M, Stücker J, Contreras S, Aricò G. 2023 The BACCO simulation project: biased tracers in real space. *Mon. Not. R. Astron. Soc.* **524**, 2407–2419. (doi:10.1093/mnras/stad2008)
- Bakx T, Chisari NE, Vlah Z. 2024 COBRA: optimal factorization of cosmological observables. (doi:10.48550/arXiv.2407.04660)
- Lawrence E *et al.* 2017 The mira-titan universe. II. Matter power spectrum emulation. *Astrophys. J.* **847**, 50. (doi:10.3847/1538-4357/aa86a9)
- Spurio Mancini A, Piras D, Alsing J, Joachimi B, Hobson MP. 2022 CosmoPower: emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys. *Mon. Not. R. Astron. Soc.* **511**, 1771–1788. (doi:10.1093/mnras/stac064)
- Piras D, Spurio Mancini A. 2023 CosmoPower-JAX: high-dimensional Bayesian inference with differentiable cosmological emulators. *Open J. Astrophys.* **6**, 20. (doi:10.21105/astro.2305.06347)
- Nishimichi T *et al.* 2019 Dark quest. I. Fast and accurate emulation of halo clustering statistics and its application to galaxy clustering. *Astrophys. J.* **884**, 29. (doi:10.3847/1538-4357/ab3719)
- Mootooyaloo A, Jaffe AH, Heavens AF, Leclercq F. 2022 Kernel-based emulator for the 3D matter power spectrum from CLASS. *Astron. Comput.* **38**, 100508. (doi:10.1016/j.ascom.2021.100508)
- Knabenhans M *et al.* 2019 *Euclid* preparation: II. The euclidemulator – a tool to compute the cosmology dependence of the nonlinear matter power spectrum. *Mon. Not. R. Astron. Soc.* **484**, 5509–5529. (doi:10.1093/mnras/stz197)
- Knabenhans M *et al.* 2021 *Euclid* preparation: IX. EuclidEmulator2 – power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations. *Mon. Not. R. Astron. Soc.* **505**, 2840–2869. (doi:10.1093/mnras/stab1366)
- Winther HA, Casas S, Baldi M, Koyama K, Li B, Lombriser L, Zhao GB. 2019 Emulators for the nonlinear matter power spectrum beyond Λ CDM100. *Phys. Rev. D* **100**, 123540. (doi:10.1103/PhysRevD.100.123540)
- Heitmann K, Higdon D, White M, Habib S, Williams BJ, Lawrence E, Wagner C. 2009 The coyote universe. II. Cosmological models and precision emulation of the nonlinear matter power spectrum. *Astrophys. J.* **705**, 156–174. (doi:10.1088/0004-637X/705/1/156)
- Heitmann K, Lawrence E, Kwan J, Habib S, Higdon D. 2014 The coyote universe extended: precision emulation of the matter power spectrum. *Astrophys. J.* **780**, 111. (doi:10.1088/0004-637X/780/1/111)

16. Smith RE, Angulo RE. 2019 Precision modelling of the matter power spectrum in a planck-like universe. *Mon. Not. R. Astron. Soc.* **486**, 1448–1479. (doi:10.1093/mnras/stz890)
17. Fendt WA, Wandelt BD. 2007 Computing high accuracy power spectra with pico. *arXiv-Prints* arXiv:0712.0194. (doi:10.48550/arXiv.0712.0194)
18. Fendt WA, Wandelt BD. 2007 Pico: parameters for the impatient cosmologist. *Astrophys. J.* **654**, 2–11. (doi:10.1086/508342)
19. Bardeen JM, Bond JR, Kaiser N, Szalay AS. 1986 The statistics of peaks of Gaussian random fields. *Astrophys. J.* **304**, 15. (doi:10.1086/164143)
20. Eisenstein DJ, Hu W. 1998 Baryonic features in the matter transfer function. *Astrophys. J.* **496**, 605–614. (doi:10.1086/305424)
21. Eisenstein DJ, Hu W. 1999 Power spectra for cold dark matter and its variants. *Astrophys. J.* **511**, 5–15. (doi:10.1086/306640)
22. Kronberger G, Burlacu B, Kommenda M, Winkler SM, Affenzeller M. 2024 *Symbolic regression*. London, UK: Chapman & Hall/CRC Press.
23. Bartlett DJ, Kammerer L, Kronberger G, Desmond H, Ferreira PG, Wandelt BD, Burlacu B, Alonso D, Zennaro M. 2024 A precise symbolic emulator of the linear matter power spectrum. *A&A* **686**, A209. (doi:10.1051/0004-6361/202348811)
24. Bartlett DJ, Wandelt BD, Zennaro M, Ferreira PG, Desmond H. 2024 SYREN-HALOFIT: a fast, interpretable, high-precision formula for the Λ CDM nonlinear matter power spectrum. *A&A* **686**, A150. (doi:10.1051/0004-6361/202449854)
25. Sui C, Bartlett DJ, Pandey S, Desmond H, Ferreira PG, Wandelt BD. 2025 SYREN-NEW: precise formulae for the linear and nonlinear matter power spectra with massive neutrinos and dynamical dark energy. *A&A* **698**, A1. (doi:10.1051/0004-6361/202452854)
26. Orjuela-Quintana JB, Nesseris S, Sapone D. 2024 Machine learning unveils the linear matter power spectrum of modified gravity. *Phys. Rev. D* **109**, 063511. (doi:10.1103/PhysRevD.109.063511)
27. Orjuela-Quintana JB, Sapone D, Nesseris S. 2024 Analytical emulator for the linear matter power spectrum from physics-informed machine learning. *arXiv E-Prints* arXiv:2407.16640.
28. Kammerer L, Bartlett DJ, Kronberger G, Desmond H, Ferreira PG. 2025 syren-baryon: analytic emulators for the impact of baryons on the matter power spectrum. *A&A* **701**, A284. (doi:10.1051/0004-6361/202555887)
29. Taylor PL, Kitching TD, McEwen JD. 2018 Preparing for the cosmic shear data flood: Optimal data extraction and simulation requirements for stage IV dark energy experiments. *Phys. Rev. D* **98**, 043532. (doi:10.1103/PhysRevD.98.043532)
30. Burlacu B, Kronberger G, Kommenda M. 2020 Operon C++: an efficient genetic programming framework for symbolic regression. In *Proc. 2020 Genetic and Evolutionary Computation Conf. Companion GECCO '20*, pp. 1562–1570. New York, NY, USA: Association for Computing Machinery. (doi:10.1145/3377929.3398099)
31. Levenberg K. 1944 A method for the solution of certain non-linear problems in least squares. *Quart. Appl. Math.* **2**, 164–168. (doi:10.1090/qam/10666)
32. Marquardt DW. 1963 An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.* **11**, 431–441. (doi:10.1137/0111030)
33. Kommenda M, Burlacu B, Kronberger G, Affenzeller M. 2020 Parameter identification for symbolic regression using nonlinear least squares. *Genet. Program. Evolvable Mach.* **21**, 471–501. (doi:10.1007/s10710-019-09371-3)
34. Laumanns M, Thiele L, Deb K, Zitzler E. 2002 Combining convergence and diversity in evolutionary multiobjective optimization. *Evol. Comput.* **10**, 263–282. (doi:10.1162/106365602760234108)
35. Bartlett DJ, Desmond H, Ferreira PG. 2023 Priors for symbolic regression. In *The Genetic and Evolutionary Computation Conf. 2023*, Lisbon Portugal. New York, NY: ACM. (doi:10.1145/3583133.3596327)
36. Heath DJ. 1977 The growth of density perturbations in zero pressure Friedmann-Lemaitre universes. *Mon. Not. R. Astron. Soc.* **179**, 351–358. (doi:10.1093/mnras/179.3.351)
37. Peebles PJE. 1980 *The large-scale structure of the universe*. Princeton, NJ: Princeton University Press.

38. Lewis A, Challinor A, Lasenby A. 2000 Efficient computation of cosmic microwave background anisotropies in closed friedmann-robertson-walker models. *Astrophys. J.* **538**, 473–476. (doi:10.1086/309179)
39. Bartlett DJ, Desmond H, Ferreira PG. Exhaustive symbolic regression. *IEEE Trans. Evol. Computat.* **28**, 950–964. (doi:10.1109/TEVC.2023.3280250)
40. Dodelson S. 2003 *Modern cosmology*. Amsterdam, The Netherlands: Academic Press.
41. Campagne JE *et al.* 2023 JAX-COSMO: An end-to-end differentiable and GPU accelerated cosmology library. *Open J. Astrophys.* **6**, 15. (doi:10.21105/astro.2302.05163)
42. Smith RE, Peacock JA, Jenkins A, White SDM, Frenk CS, Pearce FR, Thomas PA, Efstathiou G, Couchman HMP. 2003 Stable clustering, the halo model and non-linear cosmological power spectra. *Mon. Not. R. Astron. Soc.* **341**, 1311–1332. (doi:10.1046/j.1365-8711.2003.06503.x)
43. Bird S, Viel M, Haehnelt MG. 2012 Massive neutrinos and the non-linear matter power spectrum. *Mon. Not. R. Astron. Soc.* **420**, 2551–2561. (doi:10.1111/j.1365-2966.2011.20222.x)
44. Takahashi R, Sato M, Nishimichi T, Taruya A, Oguri M. 2012 Revising the halofit model for the nonlinear matter power spectrum. *Astrophys. J.* **761**, 152. (doi:10.1088/0004-637X/761/2/152)
45. Spergel DN *et al.* 2003 First-year wilkinson microwave anisotropy probe (WMAP) observations: determination of cosmological parameters. *Astrophys. J. Suppl. Ser.* **148**, 175–194. (doi:10.1086/377226)
46. Abbott TMC *et al.* 2018 Dark energy survey year 1 results: cosmological constraints from galaxy clustering and weak lensing. *Phys. Rev. D* **98**, 043526. (doi:10.1103/PhysRevD.98.043526)
47. Krause E *et al.* 2017 Dark energy survey year 1 results: multi-probe methodology and simulated likelihood analyses. *arXiv E-Prints* arXiv:1706.09359. (doi:10.48550/arXiv.1706.09359)
48. LoVerde M, Afshordi N. 2008 Extended limber approximation. *Phys. Rev. D* **78**, 123506. (doi:10.1103/PhysRevD.78.123506)
49. Joachimi B, Mandelbaum R, Abdalla FB, Bridle SL. 2011 Constraints on intrinsic alignment contamination of weak lensing surveys using the MegaZ-LRG sample. *A&A* **527**, A26. (doi:10.1051/0004-6361/201015621)
50. Zuntz J, Paterno M, Jennings E, Rudd D, Manzotti A, Dodelson S, Bridle S, Sehrish S, Kowalkowski J. 2015 CosmoSIS: modular cosmological parameter estimation. *Astron. Comput.* **12**, 45–59. (doi:10.1016/j.ascom.2015.05.005)
51. Feroz F, Hobson MP. 2008 Multimodal nested sampling: an efficient and robust alternative to markov chain monte carlo methods for astronomical data analyses. *Mon. Not. R. Astron. Soc.* **384**, 449–463. (doi:10.1111/j.1365-2966.2007.12353.x)
52. Feroz F, Hobson MP, Bridges M. 2009 MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Mon. Not. R. Astron. Soc.* **398**, 1601–1614. (doi:10.1111/j.1365-2966.2009.14548.x)
53. Feroz F, Hobson MP, Cameron E, Pettitt AN. Importance nested sampling and the multinest algorithm. *Open J. Astrophys.* **2**, 10. (doi:10.21105/astro.1306.2144)
54. Hoffman MD, Gelman A. 2011 No-U-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *arXiv e-Prints* arXiv:1111.4246. (doi:10.48550/arXiv.1111.4246)
55. Bingham E *et al.* 2018 Pyro: deep universal probabilistic programming. *arXiv E-Prints* arXiv:1810.09538. (doi:10.48550/arXiv.1810.09538)
56. Phan D, Pradhan N, Jankowiak M. 2019 Composable effects for flexible and accelerated probabilistic programming in numpyro. *arXiv e-Prints* arXiv:1912.11554. (doi:10.48550/arXiv.1912.11554)
57. Lewis A. 2019 GetDist: a python package for analysing monte carlo samples. *arXiv e-Prints* arXiv:1910.13970. (doi:10.48550/arXiv.1910.13970)