

# Assessment of gradient-based samplers in standard cosmological likelihoods

Arrykrishna Mootoovaloo,<sup>1</sup>★† Jaime Ruiz-Zapatero<sup>1b</sup>,<sup>2,3</sup> Carlos García-García<sup>1</sup> and David Alonso<sup>1</sup>

<sup>1</sup>Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

<sup>2</sup>Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

<sup>3</sup>Advanced Research Computing Centre, University College London, 90 High Holborn, London WC1V 6LJ, UK

Accepted 2024 September 7. Received 2024 September 7; in original form 2024 June 20

## ABSTRACT

We assess the usefulness of gradient-based samplers, such as the no-U-turn sampler (NUTS), by comparison with traditional Metropolis–Hastings (MH) algorithms, in tomographic  $3 \times 2$  point analyses. Specifically, we use the Dark Energy Survey (DES) Year 1 data and a simulated dataset for the Large Synoptic Survey Telescope (LSST) survey as representative examples of these studies, containing a significant number of nuisance parameters (20 and 32, respectively) that affect the performance of rejection-based samplers. To do so, we implement a differentiable forward model using JAX-COSMO, and we use it to derive parameter constraints from both data sets using the NUTS algorithm implemented in NUMPYRO, and the Metropolis–Hastings algorithm as implemented in COBAYA. When quantified in terms of the number of effective number of samples taken per likelihood evaluation, we find a relative efficiency gain of  $\mathcal{O}(10)$  in favour of NUTS. However, this efficiency is reduced to a factor  $\sim 2$  when quantified in terms of computational time, since we find the cost of the gradient computation (needed by NUTS) relative to the likelihood to be  $\sim 4.5$  times larger for both experiments. We validate these results making use of analytical multivariate distributions (a multivariate Gaussian and a Rosenbrock distribution) with increasing dimensionality. Based on these results, we conclude that gradient-based samplers such as NUTS can be leveraged to sample high-dimensional parameter spaces in Cosmology, although the efficiency improvement is relatively mild for moderate ( $\mathcal{O}(50)$ ) dimension numbers, typical of tomographic large-scale structure analyses.

**Key words:** methods: data analysis – methods: statistical – cosmological parameters.

## 1 INTRODUCTION

Cosmology has witnessed a transformative integration of machine learning methodologies into its toolbox. This has been partially prompted by the growing complexity of cosmological data sets coupled with the increasingly intricate nature of the theoretical models needed to describe them. In this sense, emulating techniques have become important in modelling complicated functions in cosmology.

In very simple terms, emulation entails finding an approximate function which can model the quantity we are interested in. The idea of emulation is in fact an old concept. For instance, Eisenstein & Hu (1998) derived an analytic expression to describe the linear matter transfer function in the presence of cold dark matter, radiation, and baryons. However, deriving these types of expressions in general, purely in terms of cosmological parameters, requires significant human ingenuity, particularly in the presence of growing model complexity and ever more stringent accuracy requirements. Various types of emulators have been designed, with their own advantages and disadvantages. Techniques such as polynomial regression, neural networks, Gaussian Processes (GPs), and genetic

algorithms have been explored by different groups. For example, Fendt & Wandelt (2007) used polynomial regression to emulate the Cosmic Microwave Background (CMB) power spectra, while Habib et al. (2007) used Gaussian Processes, together with a compression scheme, to emulate the non-linear matter spectrum from simulations. Recently, Aricò, Angulo & Zennaro (2021), Spurio Mancini et al. (2022), and Bonici, Bianchini & Ruiz-Zapatero (2024) developed a neural network framework to emulate different power spectra. Moreover, Bartlett et al. (2023, 2024) used symbolic regression – a technique for finding mathematical expression of the function of interest – to emulate the matter power spectrum. Finally, Mootoovaloo et al. (2022) introduced a Gaussian-Process-based approach to emulate both the linear and non-linear matter power spectra (and Mootoovaloo et al. (2020) explored the combination of emulation and compression in the context of weak lensing). This methodology is further discussed in Section 4.3.

Besides accelerating cosmological calculations, the availability of emulators also enables us to exploit differentiable parameter inference methods. These methods, which include Hamiltonian Monte-Carlo (HMC) samplers, and variational inference schemes, exploit the knowledge of the likelihood derivatives to dramatically improve the acceptance rate of the sample in high-dimensional spaces, or to obtain an approximate form for the marginal posterior (Duane et al. 1987; Blei, Kucukelbir & McAuliffe 2016). Our goal

\* E-mail: [arrykrishna.mootoovaloo@physics.ox.ac.uk](mailto:arrykrishna.mootoovaloo@physics.ox.ac.uk)

† LSST-DA Catalyst Fellow.

in this paper is to explore the extent to which differentiable methods can be used to accelerate the task of parameter inference in standard (e.g. power-spectrum-based) cosmological analyses. In particular we will compare two sampling methods: NUTS, a gradient-based HMC sampler, and the state-of-the-art implementation of the Metropolis–Hastings algorithm in COBAYA (Lewis & Bridle 2002; Hoffman & Gelman 2011). While both methods aim to efficiently explore parameter space, they differ in their approach to incorporating information about the target distribution. NUTS leverages the gradient of the log posterior to dynamically adjust its step size and mass matrix during the warmup phase, allowing it to adapt to the local geometry of the posterior distribution. In contrast, the Markov Chain Monte Carlo (MCMC) method in COBAYA can learn the proposal covariance matrix but may not be optimal. Alternatively, one could also specify a covariance matrix for the proposal distribution, which requires prior specification and might not capture the full complexity of the target distribution. Understanding the trade-offs between gradient-based and non-gradient-based sampling methods is crucial for selecting the most appropriate approach depending on the characteristics of the problem at hand. Although related studies have been carried out recently in the literature, our aim in this work is to make as fair a comparison as possible between both sampling approaches, keeping all other aspects of the analysis (e.g. the hardware platform used, the level of parallelization, the usage of emulators) fixed and equal between both methods. While in this work, we compare gradient versus non-gradient based samplers, Karamanis, Beutler & Peacock (2021) recently developed ZEUS, a non-gradient based sampler and compared it with another non-gradient based sampler, EMCEE.

Our contributions in this work are as follows: (1) We integrate an emulator for the linear matter power spectrum in JAX-COSMO (Campagne et al. 2023) and leverage its existing functionalities for computing power spectra for galaxy clustering and cosmic shear. (2) We take advantage of gradient-based samplers such as NUTS to sample the posterior of the cosmological and nuisance parameters using DES Year 1 data (Abbott et al. 2018) and a future LSST-like survey data. (3) We perform an in-depth assessment of whether differentiability is helpful in this context.

In Section 2, we describe the Gaussian Process framework used here. In Section 3, we elaborate on the different metrics used to assess the performance of the samplers. In Section 4, we use the DES Year 1 data to infer the cosmological and nuisance parameters via emulation and gradient-based sampling techniques, comparing the performance of standard and differentiable sampling methods. Moreover, in Section 5, we investigate how the effective sample size for different samplers varies across as a function of model dimensionality making use of analytical distributions. Finally, in Section 6, we look into the performance gain of differentiable samplers using a future LSST-like data set, before concluding in Section 7. In Appendix A, we also briefly review Hamiltonian Monte Carlo sampling techniques and its variant no-U-turn sampler (NUTS).

## 2 METHOD

In this section, we briefly describe the steps towards building the power spectrum emulator used in this work and the methods used to extract the derivatives of our model.

### 2.1 Gaussian Process emulation

A detailed description of GPs can be found in Rasmussen & Williams (2006), and we only outline the methodology briefly here. A GP is essentially a collection of random variables and a finite subset of the

variables which has a joint distribution which follows a multivariate normal distribution. It is widely applicable not only to regression problems, but also to active learning scenarios.

Suppose we have a training set,  $\{\mathbf{X}, \mathbf{y}\}$ . For simplicity, we will assume noise-free regression, such that the data covariance is  $\Sigma = \sigma^2 \mathbb{I}$ , where  $\sigma$  is a tiny value, of the order of  $10^{-5}$ . We will also assume a zero mean prior on the function we want to emulate,  $f \sim \mathcal{N}(\mathbf{0}, \mathbf{K})$ .  $\mathbf{K}$  is the kernel matrix, for which different functional forms can be assumed. A commonly used kernel is the radial-basis or squared-exponential function (RBF), which is given by:

$$K_{ij} = A \exp \left[ -\frac{1}{2} (\mathbf{x}_i - \mathbf{x}_j)^T \Lambda^{-1} (\mathbf{x}_i - \mathbf{x}_j) \right] \quad (1)$$

where  $A$  is the amplitude of the kernel and  $\Lambda$  is typically a diagonal matrix which consists of the characteristic length-scales for each dimension. Throughout this work, we will use the RBF kernel.

Given the data and the prior of the function, the posterior distribution of the function can be obtained via Bayes' theorem, that is,

$$p(\mathbf{f} | \mathbf{y}, \mathbf{X}) = \frac{p(\mathbf{y} | \mathbf{X}, \mathbf{f}) p(\mathbf{f} | \mathbf{X})}{p(\mathbf{y} | \mathbf{X})}, \quad (2)$$

where the denominator is the marginal likelihood, also a multivariate normal distribution,  $p(\mathbf{y} | \mathbf{X}) = \mathcal{N}(\mathbf{0}, \mathbf{K} + \Sigma)$ . It is a function of the kernel parameters,  $\mathbf{v} = \{A, \Lambda\}$ . The latter are determined by maximizing the marginal likelihood, which is equivalent to minimizing the negative log marginal likelihood, that is,

$$\begin{aligned} \mathbf{v}_{\text{optimum}} &= \arg \min_{\mathbf{v}} -\log p(\mathbf{y} | \mathbf{X}) \\ &= \arg \min_{\mathbf{v}} \frac{1}{2} \mathbf{y}^T (\mathbf{K} + \Sigma)^{-1} \mathbf{y} + \frac{1}{2} \log |\mathbf{K} + \Sigma| + \text{constant}. \end{aligned}$$

In general, we are interested in predicting the function at test points in parameter space. For a given test point,  $\mathbf{x}_*$ , the predictive distribution is another normal distribution with mean and variance given by:

$$\bar{f}_* = \mathbf{k}_*^T (\mathbf{K} + \Sigma)^{-1} \mathbf{y} \quad (3)$$

$$\sigma_*^2 = k_{**} - \mathbf{k}_*^T (\mathbf{K} + \Sigma)^{-1} \mathbf{k}_* \quad (4)$$

where  $\mathbf{k}_* \in \mathbb{R}^{N_\theta}$  is the kernel computed between the test point and all the training points. Note that the mean is very quick to compute since  $\boldsymbol{\alpha} = (\mathbf{K} + \Sigma)^{-1} \mathbf{y}$  is computed only once and is cached. It then requires  $\mathcal{O}(N_\theta)$  operations to compute the mean. On the other hand, computing the predictive variance is expensive since it requires  $\mathcal{O}(N_\theta^2)$  operations assuming the Cholesky factor of  $\mathbf{K} + \Sigma$  is cached.

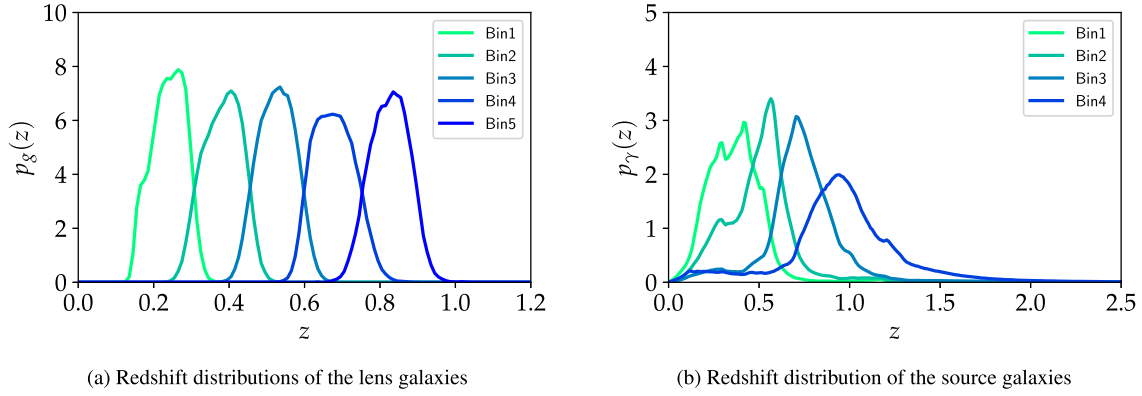
### 2.2 Gradient of the model

If we define  $\boldsymbol{\alpha} = (\mathbf{K} + \Sigma)^{-1} \mathbf{y}$ , the mean function can be written as  $\bar{f}_* = \mathbf{k}_*^T \boldsymbol{\alpha}$ , meaning that, the first derivative of the mean function with respect to the  $i^{\text{th}}$  parameter in the parameter vector  $\boldsymbol{\theta}_*$  (a test point in parameter space) is:

$$\frac{\partial \bar{f}_{*i}}{\partial \theta_{*i}} = \left( \frac{\partial \mathbf{k}_{*i}}{\partial \theta_{*i}} \right)^T \boldsymbol{\alpha}. \quad (5)$$

As shown in Mootooyaloo et al. (2022), analytical expression for the first and second derivatives of the mean function can be derived. In this paper, we leverage automatic differentiation functionalities in JAX to compute the first derivatives.

Emulating an intermediate function, for example, the linear matter power spectrum offers various advantages. It has a lower dimensionality, meaning that few training points suffice to accurately model



**Figure 1.** Panels (a) and (b) show the estimated redshift distributions of the lens and source galaxies. In particular, we have five bins for the lens galaxies and four bins for the source galaxies. These bins are used in the calculation of the power spectra for galaxy clustering and cosmic shear. These distributions were estimated and made publicly available by Abbott et al. (2018).

the function. Moreover, once the non-linear matter power spectrum is computed, the calculation of other power spectra such as the shear power spectra involves a rather straightforward linear algebra (integral of the product of the lensing kernel and the matter power spectrum). The latter can be achieved very quickly with existing libraries.

While the above describes how we can calculate the gradient of the GP with respect to the input parameters, another important quantity is the derivative of the log density of a probability distribution. For example, in a Hamiltonian Monte Carlo sampling scheme, if  $f(\boldsymbol{\Omega})$  is the target distribution and  $\boldsymbol{\Omega} \in \mathbb{R}^p$ , it is desirable to have quick access to  $-\frac{\partial \ln f}{\partial \boldsymbol{\Omega}}$ . A Gaussian likelihood is often adopted in most cosmological data analysis problems. In this spirit, if  $\boldsymbol{\mu}(\boldsymbol{\Omega})$  is the forward theoretical model which we want to compare with the data given a set of parameters  $\boldsymbol{\Omega}$ , and if we define the log-likelihood as  $\mathcal{L} \equiv -\log p(\mathbf{x}|\boldsymbol{\Omega})$ , then

$$\mathcal{L} = \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu}) + \frac{1}{2} \log |\mathbf{C}| + \text{constant} \quad (6)$$

and assuming the covariance is not a function of the parameters, the first derivatives of  $\mathcal{L}$  with respect to  $\boldsymbol{\Omega}$  is:

$$\frac{\partial \mathcal{L}}{\partial \boldsymbol{\Omega}} = -\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\Omega}} \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu}). \quad (7)$$

Note that the derivatives of  $\mathcal{L}$  with respect to  $\boldsymbol{\Omega}$  is a  $p$  dimensional vector. The above equation implies that we can easily compute the derivatives of the negative log-likelihood if we have access to the derivatives of the forward theoretical model with respect to the input parameters. In general, one can get the derivatives of a complex, expensive, and non-linear model by using finite difference method. This is not optimal since for each step in a hybrid Monte Carlo system, one would require  $p$  evaluations, assuming a forward-difference or backward-difference method is adopted. Fortunately, with the emulator, approximate derivatives can be obtained (see equation 5).

Another quantity of interest is the second derivative of  $\mathcal{L}$  at any point in parameter space. This is essentially the Hessian matrix,

$$\mathbf{H}_{\boldsymbol{\Omega}} \equiv \frac{\partial}{\partial \boldsymbol{\Omega}} \left( \frac{\partial \mathcal{L}}{\partial \boldsymbol{\Omega}} \right)^T \quad (8)$$

and differentiating equation (7) with respect to  $\boldsymbol{\Omega}$ ,

$$\mathbf{H}_{\boldsymbol{\Omega}} = -\left[ \frac{\partial}{\partial \boldsymbol{\Omega}} \left( \frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\Omega}} \right)^T \right] \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu}) + \frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\Omega}} \mathbf{C}^{-1} \left( \frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\Omega}} \right)^T. \quad (9)$$

Note that  $\mathbf{H}_{\boldsymbol{\Omega}} \in \mathbb{R}^{p \times p}$ . The parameter vector  $\boldsymbol{\Omega} = \hat{\boldsymbol{\Omega}}$  which yields

$$\left. \frac{\partial \mathcal{L}}{\partial \boldsymbol{\Omega}} \right|_{\boldsymbol{\Omega}=\hat{\boldsymbol{\Omega}}} = 0,$$

corresponds to the maximum likelihood point. For an unbiased estimate of  $\boldsymbol{\Omega} = \hat{\boldsymbol{\Omega}}$ , the Fisher information matrix is an expectation of the Hessian matrix, that is,  $\mathbf{F}_{\boldsymbol{\Omega}} = \langle \mathbf{H}_{\boldsymbol{\Omega}} \rangle$ . The inverse of the Hessian matrix at this point is the covariance matrix of the parameters. See Tegmark, Taylor & Heavens (1997) for a detailed explanation.

### 3 METRICS

This section describes the different metrics we will use to compare the performance of different samplers, both in terms of their ability to sample the posterior, and of the quality of the resulting parameter chains.

#### 3.1 Kullback–Leibler divergence

Suppose we know the exact distribution,  $q(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C})$  and the distribution recovered by a given sampler  $p(\mathbf{x}) \sim \mathcal{N}(\bar{\mathbf{x}}, \hat{\mathbf{C}})$ , and that both of them are multivariate normal. The level of agreement between both distributions can be quantified through the Kullback–Leibler (KL) divergence, which is defined as:

$$D_{\text{KL}}(p||q) = \int_{-\infty}^{+\infty} p(\mathbf{x}) \log \left[ \frac{p(\mathbf{x})}{q(\mathbf{x})} \right] d\mathbf{x}. \quad (10)$$

For two multivariate normal distributions, this can be computed analytically as:

$$D_{\text{KL}} = \frac{1}{2} \left[ \text{tr}(\mathbf{C}^{-1}\hat{\mathbf{C}}) - d + (\boldsymbol{\mu} - \bar{\mathbf{x}})^T \mathbf{C}^{-1}(\boldsymbol{\mu} - \bar{\mathbf{x}}) + \ln \frac{|\mathbf{C}|}{|\hat{\mathbf{C}}|} \right], \quad (11)$$

where  $d$  is the number of variables. The KL divergence can be understood as a distance measure between distributions. For example, in equation (11),  $D_{\text{KL}} \rightarrow 0$  as  $\bar{\mathbf{x}} \rightarrow \boldsymbol{\mu}$  and  $\hat{\mathbf{C}} \rightarrow \mathbf{C}$ . This metric will only be used for the multivariate normal example in Section 5. Computing the KL divergence, involving non-linear models, will require numerical methods for evaluating high-dimensional integrations.

### 3.2 Effective sample size

The effective sample size is an approximate measure of the number of effectively uncorrelated samples in a given Markov chain. It is defined as:

$$n_{\text{eff}} = \frac{m \cdot n}{1 + 2 \sum_{t=1}^{\infty} \rho_t}, \quad (12)$$

where  $m$  is the number of chains,  $n$  is the length of each chain, and  $\rho_t$  is the autocorrelation of a sequence at lag  $t$  (Gelman et al. 2015). Suppose we have the samples,  $\{\theta_i\}_{i=1}^n$ , the autocorrelation,  $\rho_t$  at lag  $t$  can be calculated using

$$\rho_t = \frac{1}{\sigma^2} \int_{\Theta} \theta^{(i)} \theta^{(i+t)} p(\theta) d\theta - \frac{\mu^2}{\sigma^2} \quad (13)$$

for a probability distribution function  $p(\theta)$  with mean  $\mu$  and variance  $\sigma^2$ . Ultimately, we are interested in the computational cost of each effective sample, since we wish to minimize the time taken to obtain a reasonably sampled posterior distribution. For this reason, in this work we will use a scaled effective sample size, which takes into account the total number of likelihood evaluations:

$$N_{\text{eff}} = \frac{n_{\text{eff}}}{N_{\mathcal{L}}}, \quad (14)$$

where  $N_{\mathcal{L}}$  is the total number of calls of the likelihood function. In a typical (non-gradient) MCMC sampler, this is simply the number of times the posterior probability is evaluated, while in HMC, this would correspond to the total number of steps performed in all the leapfrog moves.

### 3.3 Potential scale reduction factor

We also report the potential scale reduction factor,  $R$  (Gelman & Rubin 1992) which measures the ratio of the average of the variance within each chain to the variance of all the samples across the chains. If the chains are in equilibrium, then the value of  $R$  should be close to one.

Let us assume that we have a set of samples,  $\theta_{ij}$ , where  $i \in [1, n]$  and  $j \in [1, m]$ . As in the previous section,  $n$  is the length of the chain and  $m$  is the number of chains. The between-chain variance can be calculated using:

$$B = \frac{n}{m-1} \sum_{j=1}^m (\bar{\theta}_j - \bar{\theta})^2 \quad (15)$$

where  $\bar{\theta}_j = \frac{1}{n} \sum_{i=1}^n \theta_{ij}$  and  $\bar{\theta} = \frac{1}{m} \sum_{j=1}^m \bar{\theta}_j$ . Furthermore, the within-chain variance is defined as:

$$W = \frac{1}{m} \sum_{j=1}^m s_j^2, \quad (16)$$

where  $s_j^2 = \frac{1}{n-1} \sum_{i=1}^n (\theta_{ij} - \bar{\theta}_j)^2$ . The variance estimator of the within- and between- chain variances can be estimated as:

$$v = \frac{n-1}{n} W + \frac{1}{n} B. \quad (17)$$

Then, the potential scale reduction factor is calculated as:

$$R = \sqrt{\frac{v}{W}}. \quad (18)$$

The sampling algorithm is usually started at different initial points in parameter space and in practice,  $R - 1 \sim 0.01$  is deemed to be a reasonable upper limit for convergence to be achieved.

## 4 DES ANALYSIS

This section presents our first comparison between HMC and MCMC using, as an example, the  $3 \times 2$  point analysis of the first-year data set from the Dark Energy Survey (DES-Y1).

### 4.1 Data

We use the data<sup>1</sup> produced in the re-analysis of DES Y1 carried out by García-García et al. (2021). The data vector consists of tomographic angular power spectra, including galaxy clustering autocorrelations of the redMaGIC sample in five redshift bins, cosmic shear auto- and cross-correlations of the Y1 Gold sample in four redshift bins, and all the cross-correlation between clustering and shear samples. See Fig. 1 for the two sets of redshift distributions. The methodology used to construct this data vector was described in detail in García-García et al. (2021), and will not be repeated here. Following García-García et al. (2021), we apply a physical scale cut of  $k_{\text{max}} = 0.15 \text{ Mpc}^{-1}$  for galaxy clustering, and angular scale cut  $\ell_{\text{max}} = 2000$  for cosmic shear. This leads to a total of 405 data points.

Fig. 2 shows the shear–shear component of the data vector, including the measured bandpowers (blue points with error bars) and the best-fitting theoretical prediction (orange curve). Moreover, in Fig. 3, we show the joint correlation matrix for all the data points. The lower block shows the correlation among the data points for galaxy clustering, the middle block along the diagonal for the clustering–shear correlations, and the upper right block for cosmic shear. We assume a Gaussian likelihood of the form

$$p(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\beta}) = \frac{1}{\sqrt{|2\pi\mathbf{C}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right], \quad (19)$$

where  $\mathbf{x}$  is the data vector,  $\boldsymbol{\mu}(\boldsymbol{\theta}, \boldsymbol{\beta})$  is the forward theoretical model as a function of the cosmological parameters,  $\boldsymbol{\theta}$  and the nuisance parameters,  $\boldsymbol{\beta}$ , and  $\mathbf{C}$  is the data covariance matrix.

### 4.2 Theory

In this section, we discuss the forward model used to model the bandpowers. This also entails a set of nuisance parameters to account for the systematics in the model. Throughout this section, we will denote the normalized redshift distributions for the sources as  $p_{\gamma,i}(z)$  and for the lens galaxies as  $p_{g,i}(z)$ . For example, for the source tomographic distribution  $i$ ,

$$p_{\gamma,i}(z) = \frac{n_{\gamma,i}}{\int n_{\gamma,i} dz}, \quad (20)$$

where  $n_{\gamma,i}(z)$  is the unnormalized distribution.

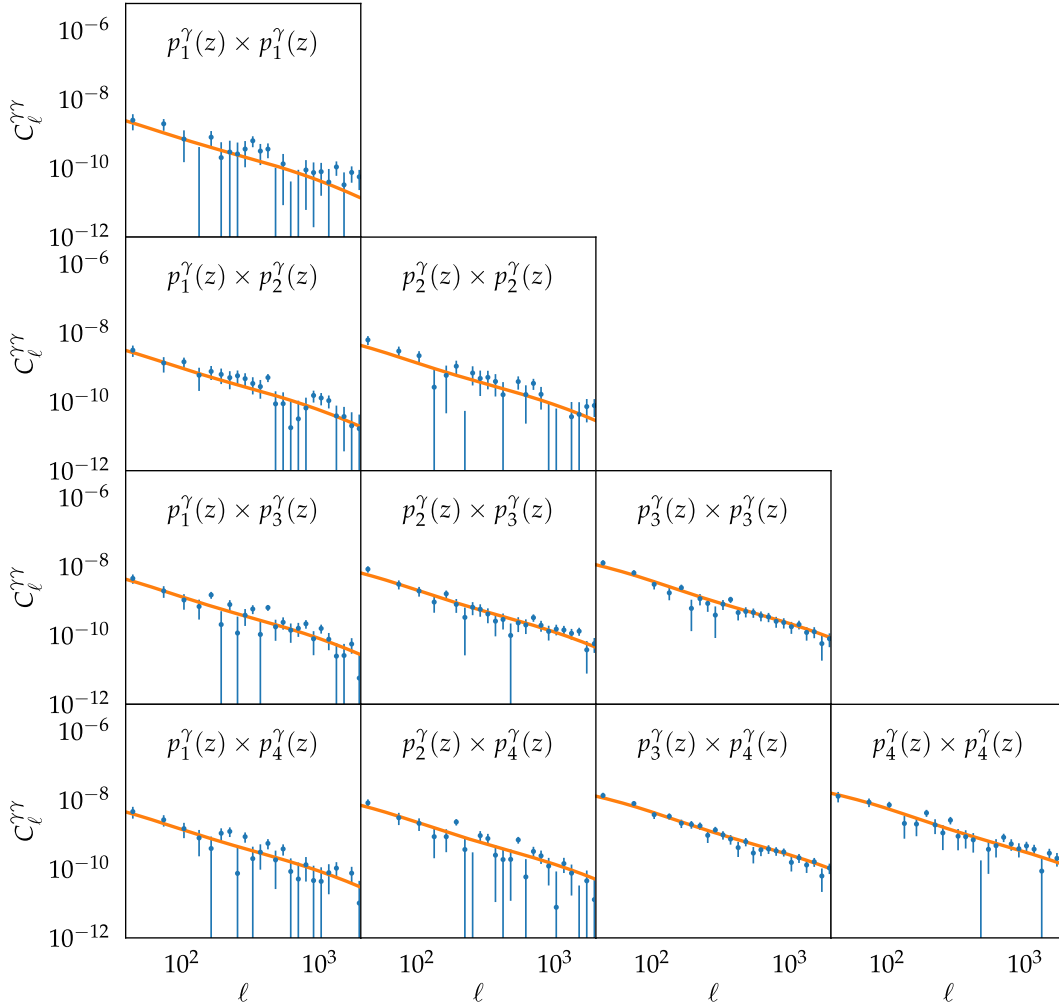
#### 4.2.1 Power spectra

Assuming a simple bias model for clustering, one which does not depend on redshift, the radial weight function for galaxy clustering in terms of the comoving radial distance,  $\chi$  is:

$$q_{g,i}(\chi) = b_i p_{g,i}(z) \frac{dz}{d\chi}, \quad (21)$$

where  $b_i$  is the galaxy bias. Under the Limber approximation (Limber 1953; LoVerde & Afshordi 2008), the power spectrum for galaxy

<sup>1</sup><https://github.com/xC-ell/growth-history/tree/main>



**Figure 2.** The data vector for the cosmic shear bandpowers is shown by the error bars in blue and the assumed set of cosmological parameter is  $\sigma_8 = 0.841$ ,  $\Omega_c = 0.229$ ,  $\Omega_b = 0.043$ ,  $h = 0.717$ ,  $n_s = 0.960$ . This corresponds to the mean of the samples obtained when sampling the posterior distribution with NUTS and the emulator. The orange curves show the cosmic shear power spectra. Since we have four tomographic redshift bins (see Fig. 1), we have a total of 10 auto- and cross-power spectra for cosmic shear.

clustering can be written as:

$$C_{\ell,i}^{gg} = \int d\chi \frac{q_{g,i}^2(\chi)}{\chi^2} P_\delta(k_\ell, z) \quad (22)$$

where  $P_\delta(k, z)$  is the three-dimensional non-linear matter power spectrum and  $k_\ell = (\ell + 1/2)/\chi$ . On the other hand, the lensing efficiency for tomographic bin  $i$  is:

$$q_{\gamma,i} = G_\ell \frac{3H_0^2 \Omega_m}{2c^2} \frac{\chi}{a(\chi)} \int_\chi^{\chi_H} d\chi' p_{\gamma,i} \frac{\chi' - \chi}{\chi}, \quad (23)$$

where  $c$  is the speed of light,  $H_0$  is the Hubble constant,  $a$  is the scale factor and

$$G_\ell = \frac{\sqrt{\ell(\ell+2)(\ell+1)(\ell-1)}}{(\ell+1/2)^2}.$$

The power spectrum due to the correlation of the lens galaxy positions in bin  $i$  with the source galaxy shear in bin  $j$  is given by:

$$C_{\ell,ij}^{\gamma\gamma} = (1 + m_j) \int \frac{q_{g,i}(\chi) q_{\gamma,j}(\chi)}{\chi^2} P_\delta(k_\ell, z) d\chi, \quad (24)$$

where  $m_j$  is the multiplicative shear bias. Moreover, the shear power spectrum is given by:

$$C_{\ell,ij}^{\gamma\gamma} = (1 + m_i)(1 + m_j) \int \frac{q_{\gamma,i}(\chi) q_{\gamma,j}(\chi)}{\chi^2} P_\delta(k_\ell, z) d\chi. \quad (25)$$


Until now, as part of the modelling framework, we have the galaxy bias,  $b_i$  for galaxy clustering and the multiplicative bias,  $m_i$  for the shear bias. Given that we have five tomographic bins for the lens galaxies and four tomographic bins for the source galaxies (see Fig. 1), the total number of galaxy biases and multiplicative biases is nine. In the next section, we elaborate on additional nuisance parameters which are accounted for in the modelling framework.

#### 4.2.2 Intrinsic alignment and shifts

The uncertainty in the redshift distributions for both the lens and source galaxies is modelled as a shift in the redshift, that is,

$$n_i(z) \rightarrow n_i(z - \Delta z_i) \quad (26)$$

hence leading to another set of nine nuisance parameters. The shifts are assumed to follow a Gaussian distribution and we adopt the



**Figure 3.** The correlation matrix for the galaxy clustering and cosmic shear data. The block matrices along the diagonal show the correlation for the different combinations of probes ( $g$  and  $\gamma$ ).

same priors as employed by Abbott et al. (2018). Moreover, we also account for the intrinsic alignment contribution to the shear signal (Joachimi et al. 2011; Abbott et al. 2018; García-García et al. 2021). Intrinsic alignments are described through the non-linear alignment model of Hirata & Seljak (2004) and Bridle & King (2007). In this case, the intrinsic alignment (IA) can be simply included as an additive contribution with a radial kernel which is proportional to the tomographic redshift distributions

$$q_{IA,i}(\chi) = -G_\ell A_{IA}(z) p(z) \frac{dz}{d\chi}, \quad (27)$$

where

$$A_{IA}(z) = 5 \times 10^{-14} \times A_{IA,0} \left( \frac{1+z}{1+z_0} \right)^\eta \frac{\rho_c \Omega_m}{D(z)}, \quad (28)$$

where  $z_0 = 0.62$ ,  $\rho_c = 2.775 \times 10^{11} h^2 M_\odot \text{Mpc}^{-3}$  is the critical density of the Universe,  $\Omega_m$  is the matter density,  $D(z)$  is the growth factor. Two nuisance parameters in this model are  $A_{IA,0}$  and  $\eta$  which are also inferred in the sampling procedure. In short, the total number of parameters in the assumed model is 25, five cosmological parameters and 20 nuisance parameters.

### 4.3 Emulation

The default method for computing the linear matter power spectrum in JAX-COSMO is the fitting formula derived by Eisenstein & Hu (1998). An alternative approach is to emulate the linear matter power spectrum as calculated by a Boltzmann solver such as CLASS (Blas, Lesgourgues & Tram 2011; Lesgourgues 2011), which we adopt in this work. In particular, using a similar approach developed in Mootooyaloo et al. (2022), we decompose the linear matter power spectrum in two parts:

$$P_l(k, z, \theta) = G(z, \theta) P_l^0(k, \theta), \quad (29)$$

where  $P_l^0$  is the linear matter power spectrum evaluated at redshift,  $z = 0$ . The input training points (the cosmological parameters) are generated using Latin Hypercube (LH) Sampling which ensures the points randomly cover the full space. In particular, the emulator is built over the redshift range of  $z \in [0.0, 3.0]$  and the wavenumber range of  $k \in [10^{-4}, 50]$  in units of  $\text{Mpc}^{-1}$ . See Table 1 for the

**Table 1.** Symbols and notations with corresponding meanings.

Symbol	Meaning
$N_\theta$	Number of training points
$N_k$	Number of wavenumbers
$N_z$	Number of redshifts
$x$	The data vector of size $N$
$\theta$	Inputs to the emulator of size $p$
$\Theta$	Input training set of size $N_\theta \times p$
$y_k$	Output of size $N_k$
$y_G$	Output of size $N_z$
$Y_k$	Output training set of size $N_\theta \times N_k$
$Y_G$	Output training set of size $N_\theta \times N_z$

different notations used in this work. Moreover, we use  $N_\theta = 1000$  training points according to the prior range shown in Table 2. In particular, we record the targets ( $G$  and  $P_l^0$ ) over  $N_z = 20$  redshift values, equally spaced in linear scale for the redshift range and  $N_k = 30$  wavenumber values, equally spaced in logarithmic scale for the wavenumber range. This gives us two training sets,  $Y_k \in \mathbb{R}^{N_\theta \times N_k}$  and  $Y_G \in \mathbb{R}^{N_\theta \times N_z}$ . We then build 50 independent models as a function of the cosmological parameters.

As described in Section 2.1, we then use the Gaussian Process formalism to build an emulator which maps the cosmological parameters,  $\Theta \in \mathbb{R}^{N_\theta \times p}$ , where  $N_\theta = 1000$  and  $p = 5$  in our application, to each of the target. The models are trained and we store  $\alpha = (K + \Sigma)^{-1} y$ . During prediction phase, the mean prediction and the first derivative can be calculated using equations (3) and (5), respectively.

One can then use either the emulator or the fitting formula of Eisenstein & Hu (1998) to compute the linear matter power spectrum. Moreover, the existing JAX-COSMO functionalities are untouched, implying that the non-linear matter power spectrum calculation proceeds via the Halofit fitting formula as a Takahashi et al. (2012), which depends directly on the linear matter power spectrum. Furthermore, the shear and galaxy clustering power spectra can be calculated very quickly via numerical integration.

### 4.4 Inference


In this work, we will compare two samplers namely, NUTS and the default MCMC sampler in COBAYA. Throughout this work, we will use the implementation of NUTS in NUMPYRO. In NUMPYRO, the sampling process is divided into two distinct phases: warmup and sampling.

During the warmup phase, the sampler explores the parameter space to adaptively tune its proposal distribution. It does so by generating a series of samples and adjusting its parameters based on these samples. This phase aims to reach a region of high-probability density in the posterior distribution while minimizing the influence of the initial guess. In particular, it focuses on improving the sampler's efficiency by adjusting its step sizes, trajectory lengths, or other parameters to achieve an optimal acceptance rate. The warmup phase is crucial for ensuring that subsequent samples are drawn from the target distribution effectively. In general, if one chooses to adapt the mass matrix during the warmup phase, this can be expensive but once this is learnt, it is fixed and the sampling process can be fast.

Once the warmup phase is complete, the sampler enters the sampling phase. Here, the sampler generates samples from the posterior distribution according to the adapted proposal distribution. These samples are used for inference and analysis, such as estimating posterior means and variances.

**Table 2.** Priors and inferred values of the cosmological parameters. We report the inferred mean and standard deviation for the different experiments we have run. As anticipated, the inferred mean and standard deviation, irrespective of the two samplers employed, are in close agreement with each other given we are using the same method for computing the different power spectra.

Parameters	$\Theta$	Priors	NUTS				COBAYA's MCMC			
			$\mu_{\text{emu}}$	$\sigma_{\text{emu}}$	$\mu_{\text{EH}}$	$\sigma_{\text{EH}}$	$\mu_{\text{emu}}$	$\sigma_{\text{emu}}$	$\mu_{\text{EH}}$	$\sigma_{\text{EH}}$
Amplitude of density fluctuations	$\sigma_8$	$\mathcal{U}[0.60, 1.00]$	0.840	0.064	0.828	0.063	0.840	0.062	0.830	0.061
CDM density	$\Omega_c$	$\mathcal{U}[0.07, 0.50]$	0.229	0.024	0.229	0.026	0.229	0.023	0.228	0.025
Baryon density	$\Omega_b$	$\mathcal{U}[0.03, 0.07]$	0.043	0.007	0.045	0.007	0.043	0.007	0.045	0.007
Hubble parameter	$h$	$\mathcal{U}[0.64, 0.82]$	0.719	0.051	0.711	0.050	0.719	0.049	0.712	0.048
Scalar spectral index	$n_s$	$\mathcal{U}[0.87, 1.07]$	0.957	0.056	0.958	0.056	0.958	0.054	0.959	0.054



**Figure 4.** The left panel shows the accuracy for the linear matter power spectrum, evaluated at  $z = 0$  over a wavenumber range of  $k \in [10^{-4}, 50]$  in units of  $\text{Mpc}^{-1}$  and the right panel shows the accuracy for the quantity,  $G$  evaluated over the redshift range of  $z \in [0.0, 3.0]$ . These quantities can be robustly calculated within an accuracy of 1 per cent.

The warmup phase is essential for the sampler to adapt to the characteristics of the target distribution, while the sampling phase focuses on generating samples for inference. Separating these phases allows Numpyro to achieve efficient sampling while ensuring the quality of the final samples.

The NUTS sampler in Numpyro involves setting a maximum depth parameter, which determines the maximum depth of the binary trees it evaluates during each iteration. The number of leapfrog steps taken is then constrained to be no more than  $2^j - 1$ , where  $j$  is the maximum tree depth (see Fig. 1 in Hoffman & Gelman 2011). The sampler reports both the tree depth and the actual number of leapfrog steps computed, along with the parameters sampled during the process. These parameters provide insights into how the sampler explores the parameter space and allow users to monitor the efficiency of the sampling process. Adjusting the maximum depth parameter can help balance exploration and efficiency, ensuring that the sampler explores the posterior distribution effectively while avoiding excessive computational costs.

For NUTS, we set the maximum number of tree depth to eight and use an initial step size of 0.01. Moreover, we fix the number of warmup steps to 500. We also generate two chains consisting of 15 000 samples each.

COBAYA uses a Metropolis MCMC as described by Lewis (2013) and one could also exploit its fast and slow sampler. The idea is to decorrelate the fast and slow parameters so that sampling becomes very efficient. This can result in large performance gains when there are many fast parameters. Throughout this work, we do not use this feature, as both the emulator and the Eisenstein and Hu (EH) methods are sufficiently fast. Additionally, accurately quantifying the performance gain poses challenges in this context. When using COBAYA's MCMC, the sampler stops when either the number of samples specified is attained or the convergence criterion is met ( $R -$

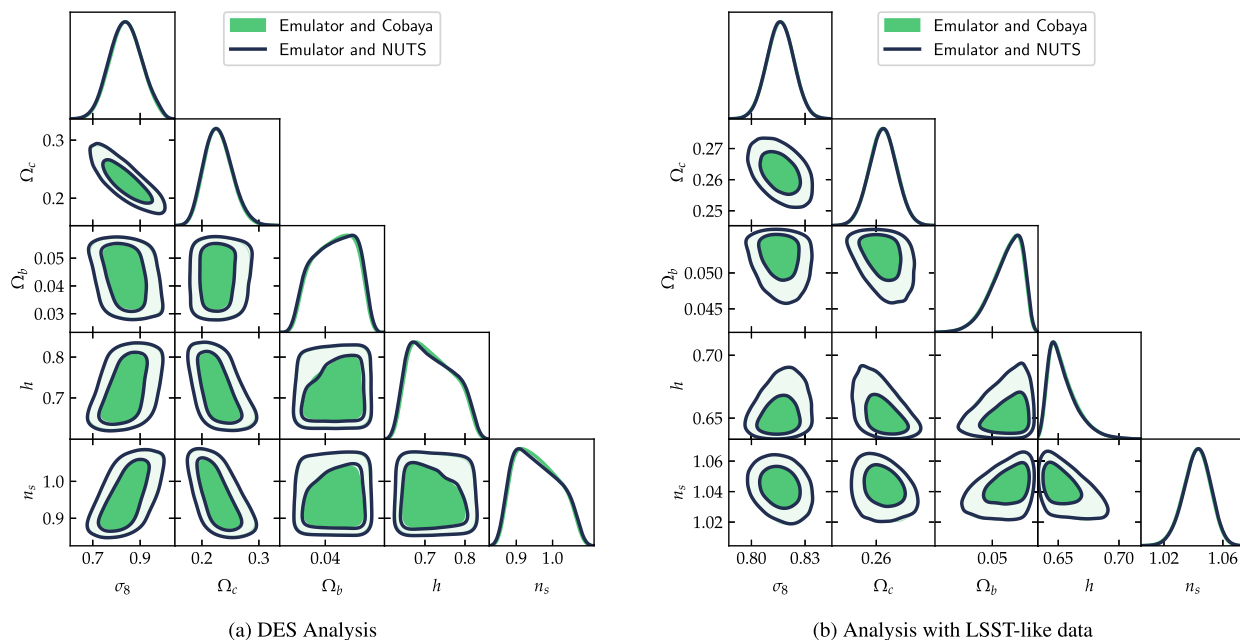
$1 \leq 0.01$ ). In some likelihood analyses, it is also possible to adopt an approximate likelihood by analytically marginalizing over the many nuisance parameters (Hadzhiyska et al. 2023; Ruiz-Zapatero et al. 2023). For COBAYA's MCMC, we run two separate chains and we set the number of MCMC samples to  $5 \times 10^5$ .

In all experiments we have performed in this work, we have provided COBAYA's MCMC with a covariance matrix for the proposal distribution. In doing so, we are essentially providing information about how to explore the parameter space efficiently. This vastly helps the sampler to generate samples that are more likely to be accepted, leading to better sampling performance. In principle, one could also specify a mass matrix in Numpyro for NUTS. The latter dynamically adapts its step size and mass matrix during the warmup phase. It uses the information from the gradient of the log posterior to tune these parameters. This adaptiveness allows NUTS to efficiently explore the parameter space without requiring explicit specification of the proposal distribution covariance.

## 4.5 Results

In this section, we explore the different inference made when using the emulator and the two samplers, NUTS and COBAYA's MCMC.

In Fig. 4, we show the accuracy of the emulator, which was then embedded into the JAX-COSMO pipeline. For the range of redshifts and wavenumber considered and the domain of the cosmological parameters, the quantities  $P_i$  and  $G$  are accurate up to  $\leq 1$  per cent and  $\leq 0.01$  per cent, respectively. Recall that we are using 1000 LH samples to build the emulator. Generating the 1000 training points using CLASS took around 1 h while training the GPs took around 2 h on a desktop computer. The training of GPs is expensive because of the  $\mathcal{O}(N^3)$  cost in computing the Cholesky factor and  $\mathcal{O}(N^2)$  cost in solving for  $\alpha = K_y^{-1}y$ . However, once they are trained and



**Figure 5.** The 1D and 2D marginalized posterior distribution of all the cosmological parameters. The green contours show the distribution when the emulator is used with the COBAYA’s MCMC sampler while the solid black curves correspond to the setup where NUTS is used for sampling the posterior distribution. There is negligible difference in the posterior when comparing the MCMC in COBAYA and NUTS. The left panel shows the posterior obtained when using the DES data while the right panel shows the contours obtained with the simulated LSST-like data.

stored, prediction is very fast and computing the log-likelihood is of the order of milliseconds. Moreover, one can use the fixed  $\alpha$ s and the kernel pre-trained hyperparameters in any GP implementation irrespective of whether we use NUMPY, PYTORCH, TENSORFLOW, and JAX. Moreover, the priors are sufficiently broad that the emulation framework can be used for different probes, for example, weak lensing as in this context.

Fig. 5 compares the marginalized 1D and 2D distributions of the cosmological parameters,  $\Theta$  (see Table 2) using the emulator with COBAYA’s MCMC and NUTS. The inferred cosmological parameters are shown in Table 2 for different setups.

Under this configuration, the potential scale reduction factor is equal to 1.00 for all parameters when either the emulator or EH is used in JAX-COSMO. Sampling the posterior with NUTS takes  $\sim 13$  h for two chains with NUMPYRO using a single GPU. Alternatively, a single run using COBAYA’s MCMC, whether with the emulator or EH in JAX-COSMO, takes approximately 5 h to sample the posterior. Note that the chains generated by both samplers did not contain the same number of samples, and therefore the difference in time above is not reflective of their relative performance. Moreover, in order to quantify the difference between the inferred parameters with either sampler, we use the ‘difference of Gaussians’ statistic:

$$\delta = \frac{|\mu_{\text{NUTS}} - \mu_{\text{COBAYA}}|}{\sqrt{\sigma_{\text{NUTS}}^2 + \sigma_{\text{COBAYA}}^2}}. \quad (30)$$


The maximum difference among the set of parameters considered in this experiment is  $\sim 0.1$ . We also compute the average of the scaled effective sample size,  $N_{\text{eff}}$ , to compare the samplers. We define the efficiency gain as:

$$\gamma = \frac{N_{\text{eff, NUTS}}}{N_{\text{eff, COBAYA}}}. \quad (31)$$

The relative gain in efficiency when using NUTS compared to COBAYA’s MCMC is  $\mathcal{O}(10)$ . When using LIMBERJACK and NUTS in TURING.JL, Ruiz-Zapatero et al. (2024) estimated a gain in  $N_{\text{eff}}$  of  $\sim 1.7$ , compared to the samples obtained using Metropolis–Hastings implemented in COBAYA (García-García et al. 2021). In addition, when using reverse mode automatic differentiation (the default setup in NUMPYRO), the cost of a single gradient calculation to the cost of a single likelihood evaluation is  $\sim 4.5$  (either with the emulator or EH). The gain in efficiency is better compared to the cost of the gradient evaluation. With JULIA, Ruiz-Zapatero et al. (2024) found this ratio to be  $\sim 5.5$  when using forward mode automatic differentiation. Note that the differences in the values above can be attributed to the fact that different samplers will, in general, have different implementations. Taking into account the more expensive gradient evaluation, we find that the overall efficiency gain of NUTS with respect to COBAYA’s MCMC, when measured in terms of computing time on the same platform, is  $\sim 2$ .

In Fig. 6, we show the joint posterior distribution of the  $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$  and  $\Omega_m$  parameters when the posterior is sampled with COBAYA’s MCMC and NUTS. We also use the Planck 2018 samples, which are publicly available, to show the same joint distribution. Note that we are using the baseline  $\Lambda$  cold dark matter (CDM) chains with the baseline likelihoods for Planck. With the public Planck chains,  $S_8 = 0.834 \pm 0.016$ . On the other hand, in our experiments, if we use the emulator with COBAYA’s MCMC and NUTS, then  $S_8 = 0.797 \pm 0.035$  and  $S_8 = 0.797 \pm 0.036$ . With EH,  $S_8 = 0.788 \pm 0.033$  and  $S_8 = 0.788 \pm 0.034$  with COBAYA’s MCMC and NUTS, respectively.

Despite the computational overhead of computing derivatives in NUTS, we observe a gain in the scaled effective sample size when comparing NUTS and COBAYA’s MCMC. However, it raises the question of whether this advantage will persist in higher-dimensional problems. To address this, we investigate three additional cases: a multivariate normal distribution, the Rosenbrock function, and a future LSST-like system with 37 model parameters. This broader



**Figure 6.** The marginalized posterior distribution for the derived parameter,  $S_8 = \sigma_8 \sqrt{\Omega_m / 0.3}$  using the emulator. The black contours show the distribution using the sampler in NUTS while the green contours correspond to the posterior using COBAYA’s MCMC. We also plot the Planck 2018 samples (in orange) of the  $S_8$  and  $\Omega_m$  which are publicly available.

analysis aims to ascertain the scalability and effectiveness of these samplers across various dimensions and problem complexities.

## 5 ANALYTICAL FUNCTIONS

In this section, we will investigate how these metrics scale with other functions, such as the multivariate normal distribution and the Rosenbrock function.

### 5.1 Multivariate normal distribution

The expression for a multivariate normal distribution is:

$$q(\mathbf{x}|\boldsymbol{\mu}, \mathbf{C}) = \frac{1}{\sqrt{|2\pi\mathbf{C}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right]. \quad (32)$$

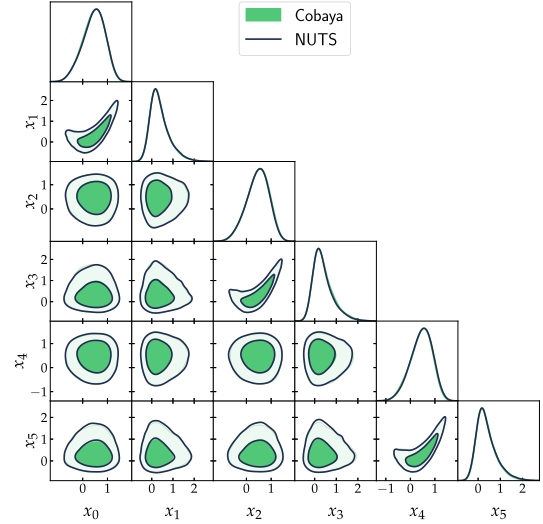
For simplicity, we assume a zero mean and an identity matrix for the covariance of the multivariate normal distribution. The aim is to obtain samples of  $\mathbf{x}$  and to estimate the sample mean,  $\bar{\mathbf{x}}$  and covariance,  $\hat{\mathbf{C}}$  as we increase the dimensionality of the problem. In the limit where we have a large number of unbiased samples of  $\mathbf{x}$ , then  $\bar{\mathbf{x}} \rightarrow \boldsymbol{\mu}$  and  $\hat{\mathbf{C}} \rightarrow \mathbf{C}$ .

### 5.2 Rosenbrock function

The next function we consider is the Rosenbrock function, which is given by:

$$f(\mathbf{x}) = \sum_{i=1}^{N/2} [\zeta(x_{2i-1}^2 - x_{2i})^2 + (x_{2i-1} - 1)^2], \quad (33)$$

where  $\mathbf{x}$  is a vector of size  $N$  and for this particular variant of the Rosenbrock function,  $N$  is even.  $\zeta$  is a factor which controls the overall shape of the final function. If it is set to zero, then, the function is simply analogous to the  $\chi^2$  term in a multivariate normal distribution with mean one and covariance matrix equal to




**Figure 7.** The posterior distribution of the first six parameters in a 100-dimensional Rosenbrock function using COBAYA’s MCMC and NUTS sampler. In particular, this is obtained by setting  $\zeta = 9$  in the function (see Section 5.2 for further details).

the identity matrix. For  $\zeta > 0$ , a quartic term is introduced in the overall function and this leads to banana-like posteriors. In both experiments (multivariate normal distribution and the Rosenbrock function), we set the step size and the maximum tree depth to 0.01 and 6, respectively. We use 500 warmup steps and generate two chains consisting of 5000 samples.

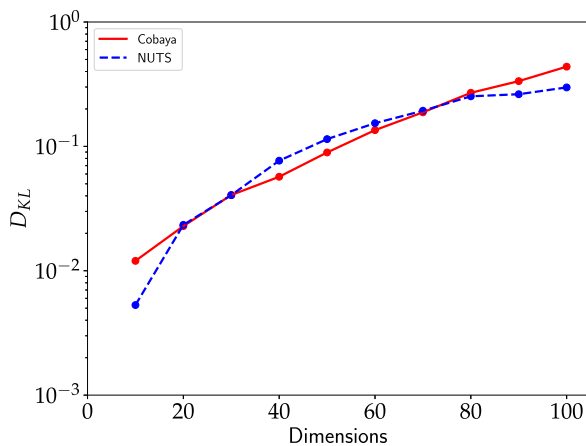
We fix  $\zeta = 9$  in our experiments. Some of the 2D joint posterior distributions follow a banana-like shapes (see Fig. 7), which demonstrates the complexity of these functions, especially in high dimensions.

### 5.3 Results

In Fig. 8, we show how  $N_{\text{eff}}$  changes as we increase the dimensionality of the multivariate normal distribution (in red and blue using COBAYA’s MCMC and NUTS, respectively). When evaluating the effective sample size per likelihood evaluation, COBAYA’s MCMC consistently yields lower values compared to NUTS. This discrepancy arises from their respective sampling strategies. NUTS uses the gradient of the likelihood function to build its proposal distribution leading to a higher acceptance rate compared to COBAYA’s MCMC whose proposal distribution only depends on the last accepted sample. Consequently, while COBAYA’s MCMC requires more likelihood evaluations to achieve a comparable effective sample size to NUTS, the latter tends to provide more informative samples in fewer evaluations. Moreover, Fig. 9 shows how the KL-divergence between the inferred distribution and the exact distribution changes as a function of the dimensionality. The KL-divergence values for both samplers are comparable except at  $d = 10$  and  $d > 80$  where, it is greater for COBAYA’s MCMC compared to NUTS. However, the inferred values of the mean and standard deviation are close to zero and unity, respectively, with either sampler. In order to compute the KL-divergence, we use the same number of samples ( $10^4$ ) for both samplers, which is equal to the number of samples drawn using NUTS. The maximum difference between the expected statistics (mean and standard deviation) and the inferred ones across all dimensions is  $\sim 0.03$  with either sampler.



**Figure 8.** The figure shows the scaled effective sample size as a function of dimensions of the multivariate normal distribution (red for COBAYA’s MCMC and blue for NUTS) and the Rosenbrock function (green for COBAYA’s MCMC and purple for NUTS).  $N_{\text{eff}}$  for NUTS is always higher compared to COBAYA’s MCMC over the dimension considered here. Moreover, as expected, for a tricky function such as the Rosenbrock function,  $N_{\text{eff}}$  is always lower compared to the multivariate normal distribution case.



**Figure 9.** As elaborated in Section 3.1, we also compute the Kullback–Leibler divergence,  $D_{\text{KL}}$ , between the sampled distribution and the expected one. For the dimensions considered here (10 – 100), most the  $D_{\text{KL}}$  values are comparable except at  $d = 10$  and  $d > 80$ , where the  $D_{\text{KL}}$  for COBAYA’s MCMC are greater than that of NUTS. Note that for a fair comparison, we use the same number of MCMC samples in both cases. The inferred mean and the standard deviation in either case are close to 0 and 1, respectively.

In Fig. 8, we show the scaled effective sample size when we sample the Rosenbrock function with NUTS and COBAYA’s MCMC (in purple and green, respectively) for different dimensions. Although the Rosenbrock function is more complex than the multivariate normal distribution,  $N_{\text{eff}}$  for NUTS is superior to that of COBAYA’s MCMC. Nevertheless, as depicted in the Fig. 8, the  $N_{\text{eff}}$  for the Rosenbrock function is anticipated to be lower than that for the multivariate normal distribution when using the same sampler. In the multivariate normal distribution example, we know the analytic expression for the function and hence, we are able to use the analytic expression for the Kullback–Leibler divergence to estimate the distance between the inferred distribution and the expected distribution. For the Rosenbrock function, this is not the case since we are dealing with a non-trivial function. However, as seen in Fig. 7, the

mean and variance are roughly the same for every two dimensions. In this spirit, we calculate the mean of the difference between the mean of the samples and the expected mean, that is,  $\langle \mu_* - \mu \rangle$  for all parameters. We do a similar calculation for the standard deviation, that is,  $\langle \sigma_* - \sigma \rangle$ .  $\mu_*$  and  $\sigma_*$  are the inferred mean and standard deviation while  $\mu$  and  $\sigma$  are the expected mean and standard deviation. For  $\zeta = 9$ ,  $\mu \sim 0.45$  and  $\sigma \sim 0.39$  in the first dimension and  $\mu \sim 0.39$  and  $\sigma \sim 0.48$  in the second dimension. We find that with NUTS and COBAYA’s MCMC, these differences are very close to zero in all dimensions,  $\langle \mu_* - \mu \rangle \lesssim 0.015$  and  $\langle \sigma_* - \sigma \rangle \lesssim 0.015$ .

Furthermore, in both the multivariate normal distribution and the Rosenbrock function examples, the potential scale reduction factor is close to one when either NUTS or COBAYA’s MCMC is used to sample the function. However, with a tricky function such as the Rosenbrock function, we find that the potential scale reduction factor gets worse ( $R \sim 1.0 - 1.4$ ) with COBAYA’s MCMC as the dimensionality increases. Moreover, the acceptance probability when NUTS is used is always  $\gtrsim 0.7$  with either the multivariate normal or the Rosenbrock function. On the other hand, COBAYA’s MCMC has an acceptance probability of  $\sim 0.3$  when sampling the multivariate normal distribution. With the Rosenbrock function, the acceptance probability varies from  $\sim 0.17$  to  $\sim 0.1$  as the dimensionality increases.

Based on the experiments performed in this section, we find that NUTS always produces more effective samples, irrespective of the function employed. With the Rosenbrock function depicting non-Gaussianity – characterized by its non-linear and asymmetric shape – it is expected that samplers will result in a reduction of  $N_{\text{eff}}$ . For  $d \leq 100$ , both samplers are able to recover the correct shape of the posterior distribution. However, NUTS is more likely to scale better to higher dimensions ( $d > 100$ ) as a result of its consistent high  $N_{\text{eff}}$ .

## 6 LSST

Lastly, we look into a  $3 \times 2$  point analysis using simulated data for a future LSST-like survey. The number of nuisance parameters is expected to be higher in order to account for the astrophysical and observational systematic uncertainties in the shear signal.

### 6.1 Data and model

Following the The LSST Dark Energy Science Collaboration (2018) document, we specify ten tomographic spaced by 0.1 in photo- $z$  between  $0.2 \leq z \leq 1.2$  bins for galaxy clustering. For cosmic shear, we assume five tomographic bins (Leonard et al. 2023). We use a physical scale cut of  $k_{\text{max}} = 0.15 \text{ Mpc}^{-1}$  for galaxy clustering, and a less conservative  $\ell_{\text{max}} = 3000$  for cosmic shear. The simulated data consists of the angular power spectra,  $C_\ell$ , with a Gaussian covariance given by the Knox formula. The  $C_\ell$  has been averaged over each of the  $\ell$ -bins. Both auto- and cross-power spectra are included in the analysis for shear–shear and galaxy–shear correlations and only auto-power spectra are included for galaxy–galaxy correlations. The data vector, after applying the scale cuts, consists of 903 elements,  $\mathbf{x} \in \mathbb{R}^{903}$  and a data covariance matrix,  $\mathbf{C} \in \mathbb{R}^{903 \times 903}$ . In this setup, we now have

- (i) ten bias parameters,  $b_i$ ,  $i \in [1, 10]$ ,
- (ii) ten shift parameters,  $\Delta z_j^g$ ,  $j \in [1, 10]$ ,
- (iii) five multiplicative bias parameters,  $m_k$ ,  $k \in [1, 5]$ ,
- (iv) five shift parameters,  $\Delta z_t^y$ ,  $t \in [1, 5]$  and
- (v) two parameters in the intrinsic alignment model ( $A_{IA}, \eta$ )

resulting in a total of 32 nuisance parameters. The total number of parameters in this experiment is 37, the five cosmological parameters (shown in Table 2) and 32 nuisance parameters.

In this experiment, for NUTS, we fix the step size to 0.01 and a maximum tree depth of eight. We also fix the number of warmup steps to 500 and generate two chains of 15 000 samples. For COBAYA’s MCMC, we run two chains and set the number of samples to  $5 \times 10^5$  and the convergence criterion to  $R - 1 \leq 0.01$ .

## 6.2 Results

The Gelman–Rubin convergence test, represented by the potential scale reduction factor, consistently indicates convergence to the target distribution for all parameters when employing either NUTS or COBAYA’s MCMC sampler. This convergence holds true regardless of the forward modelling technique used, whether it be the Eisenstein–Hu method or the emulation method.

Furthermore, when comparing the NUTS and COBAYA’s MCMC samplers, the sampling efficiency,  $\gamma$  as calculated using equation (31), is found to be  $\mathcal{O}(10)$ . The relative cost of the gradient of the likelihood to the likelihood itself is  $\sim 4.5$ . Similar to the analysis of the DES Year 1 data (see Section 4), the overall gain is  $\sim 2$ .

The constraints obtained using NUTS and COBAYA’s MCMC are comparable, with the maximum distance (see equation 30) being  $\sim 0.2$  if we use the emulator and  $\sim 0.1$  if we use EH method. The marginalized 1D and 2D posterior distribution of the cosmological parameters are shown in right panel of Fig. 5.

## 7 CONCLUSION

In this work, we have performed a quantitative assessment of different aspects related to emulation and gradient-based samplers.

In particular, we have integrated an emulator for the linear matter power in JAX-COSMO. The emulator is both accurate (see Fig. 4) and fast, comparable to the speed-up obtained when using Eisenstein & Hu method to calculate the linear matter power spectrum. Note that only 1000 training points have been used to achieve an accuracy of  $\sim 1$  per cent, compared to deep learning frameworks, which are generally more data-hungry, and require many training points to achieve almost the same level of accuracy. Moreover, as shown in Fig. 5, the constraints obtained with different samplers such as NUTS and COBAYA’s MCMC agree with each other. See Table 2 for numerical results. A notable advantage of using emulation technique is that for a particular set of configurations (priors on the cosmological parameters, range of the wavenumber considered and redshift domain), the emulator is trained once, stored and can be coupled to any likelihood code, under the assumption that the scientific problem being investigated does not require configurations beyond those of the emulator.

In the DES analysis, the use of NUTS yields a sampling efficiency gain of approximately 10 for a system with 25 parameters and a non-trivial  $\sigma_8 - \Omega_c$  degeneracy. Alongside assessing the effective sample size, the Gelman–Rubin statistics were calculated to confirm convergence across all chains in each experiment. While the efficiency gain favours NUTS in this scenario, it is not by a significant margin (roughly two when comparing sampling efficiency to the relative cost of gradient calculation). This observation aligns with findings by Ruiz-Zapatero et al. (2024).

Given the DES analysis, we also investigate the advantages of NUTS over standard non-gradient based samplers like COBAYA’s MCMC as a function of dimensionality, and we find that NUTS

is preferable in high dimensions ( $d > 100$ ). The comparison, illustrated with the multivariate normal distribution, reveals several key advantages of NUTS in high-dimensional contexts. First, the small Kullback–Leibler divergence indicates that the chain converges to expected results efficiently. Secondly, the Gelman–Rubin statistics remain close to 1.00, indicating convergence across chains. Thirdly, the scaled effective sample size consistently outperforms COBAYA’s MCMC, demonstrating the effectiveness of NUTS. This analysis underscores the utility of NUTS in efficiently exploring complex parameter spaces, especially in high dimensions where other methods like COBAYA’s MCMC may struggle.

Furthermore, an in-depth examination using the Rosenbrock function confirms the sampling efficiency gain of NUTS. This suggests that NUTS is not only advantageous in terms of convergence and effective sample size but also provides improved exploration of complex, non-trivial functions. Overall, these findings highlight the contexts where NUTS outperforms traditional non-gradient based samplers, making it a valuable tool for Bayesian inference in a wide range of applications.

In exploring a 37-dimensional parameter inference problem with NUTS and COBAYA’s MCMC for a future LSST-like survey data, we have also found that NUTS is more effective than COBAYA’s MCMC by a factor of  $\sim 2$ . NUTS consistently exhibits higher sampling efficiency and provides a greater effective sample size per likelihood call compared to COBAYA’s MCMC. This suggests that, while NUTS is better suited for handling complex parameter spaces with  $\mathcal{O}(40)$  dimensions, the relative improvement factor to be expected on a given platform is mild ( $\sim 2$  instead of orders of magnitude). While the cost of gradient calculation is a drawback of using NUTS, this process can be accelerated by leveraging GPUs. GPUs excel in parallel processing tasks, including gradient computations, offering significant speedup over traditional CPU-based methods. By harnessing GPU acceleration, NUTS becomes more feasible for high-dimensional problems and large-scale cosmological analyses.

In order to fully exploit this possibility, more complete and sophisticated theory prediction frameworks will need to be developed, able to flexibly produce predictions for a wide range of observables of interest to current and future large-scale structure and CMB experiments. Various approaches to this problem have been initiated by the community, making use of tools such as JAX (Campagne et al. 2023; Piras & Spurio Mancini 2023; Piras et al. 2024; Ruiz-Zapatero et al. 2024) and JULIA, and efforts to bring these frameworks to full maturity will significantly improve our ability to obtain both fast and robust parameter constraints from future data.

## ACKNOWLEDGEMENTS

We thank Dr. Zafirah Hosenie for reviewing this manuscript and providing useful feedbacks. We thank Prof. Alan Heavens and Prof. Andrew Jaffe for insightful discussions. AM was supported through the LSST Discovery Alliance (LSST-DA) Catalyst Fellowship project; this publication was thus made possible through the support of Grant 62192 from the John Templeton Foundation to LSST-DA. JR-Z was supported by UK Space Agency grants ST/W001721/1 and ST/X00208X/1. CG-G and DA were supported by the Beecroft Trust. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund.

## DATA AVAILABILITY

The code and part of the data products underlying this article can be found at: <https://github.com/Harry45/DESEMU/>. The full data set and pre-trained GPs can be made available upon request. An unofficial implementation of the emulator implemented by Jean-Eric Campagne – originally developed by Mootoovaloo et al. (2022) – can be found at <https://github.com/jecampagne/Jemu>.

## REFERENCES

- Abbott T. M. C. et al., 2018, *Phys. Rev. D*, 98, 043526  
 Aricò G., Angulo R. E., Zennaro M., 2021, preprint (arXiv:2104.14568)  
 Bartlett D. J. et al., 2023, *A&A*, 686, A209  
 Bartlett D. J., Wandelt B. D., Zennaro M., Ferreira P. G., Desmond H., 2024, *A&A*, 686, A150  
 Bingham E. et al., 2019, *J. Mach. Learn. Res.*, 20, 28:1  
 Blas D., Lesgourgues J., Tram T., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 034  
 Blei D. M., Kucukelbir A., McAuliffe J. D., 2016, preprint (arXiv:1601.00670)  
 Bonici M., Bianchini F., Ruiz-Zapatero J., 2024, *Open J. Astrophys.*, 7, 10  
 Bridle S., King L., 2007, *New J. Phys.*, 9, 444  
 Campagne J.-E. et al., 2023, *Open J. Astrophys.*, 6, 15  
 Duane S., Kennedy A. D., Pendleton B. J., Roweth D., 1987, *Phys. Lett. B*, 195, 216  
 Eisenstein D. J., Hu W., 1998, *ApJ*, 496, 605  
 Fendt W. A., Wandelt B. D., 2007, *ApJ*, 654, 2  
 García-García C., Ruiz-Zapatero J., Alonso D., Bellini E., Ferreira P. G., Mueller E.-M., Nicola A., Ruiz-Lapuente P., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 030  
 Gelman A., Rubin D. B., 1992, *Stat. Sci.*, 7, 457  
 Gelman A., Carlin J., Stern H., Dunson D., Vehtari A., Rubin D., 2015, *Bayesian Data Analysis*, 3rd edn. Chapman and Hall/CRC, New York  
 Habib S., Heitmann K., Higdon D., Nakhleh C., Williams B., 2007, *Phys. Rev. D*, 76, 083503  
 Hadzhiyska B., Wolz K., Azzoni S., Alonso D., García-García C., Ruiz-Zapatero J., Slosar A., 2023, *Open J. Astrophys.*, 6, 23  
 Hajian A., 2007, *Phys. Rev. D*, 75, 083525  
 Hirata C. M., Seljak U., 2004, *Phys. Rev. D*, 70, 063526  
 Hoffman M. D., Gelman A., 2011, preprint (arXiv:1111.4246)  
 Jasche J., Lavaux G., 2019, *A&A*, 625, A64  
 Joachimi B., Mandelbaum R., Abdalla F. B., Bridle S. L., 2011, *A&A*, 527, A26  
 Karamanis M., Beutler F., Peacock J. A., 2021, *MNRAS*, 508, 3589  
 Leonard C. D. et al., 2023, *Open J. Astrophys.*, 6, 8  
 Lesgourgues J., 2011, preprint (arXiv:1104.2932)  
 Lewis A., 2013, *Phys. Rev. D*, 87, 103529  
 Lewis A., Bridle S., 2002, *Phys. Rev. D*, 66, 103511  
 Limber D. N., 1953, *ApJ*, 117, 134  
 LoVerde M., Afshordi N., 2008, *Phys. Rev. D*, 78, 123506  
 Mootoovaloo A., Heavens A. F., Jaffe A. H., Leclercq F., 2020, *MNRAS*, 497, 2213  
 Mootoovaloo A., Jaffe A. H., Heavens A. F., Leclercq F., 2022, *Astron. Comput.*, 38, 100508  
 Neal R., 2011, in Brooks S., Gelman A., Jones G. L., Meng X.-L., eds, *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, New York, p. 113  
 Phan D., Pradhan N., Jankowiak M., 2019, preprint (arXiv:1912.11554)  
 Piras D., Spurio Mancini A., 2023, *Open J. Astrophys.*, 6, 20  
 Piras D., Polanska A., Spurio Mancini A., Price M. A., McEwen J. D., 2024, *Open J. Astrophys.*, 7, 73  
 Rasmussen C. E., Williams C. K. I., 2006, *Gaussian Processes for Machine Learning*. MIT Press, Cambridge  
 Ruiz-Zapatero J., Hadzhiyska B., Alonso D., Ferreira P. G., García-García C., Mootoovaloo A., 2023, *MNRAS*, 522, 5037

- Ruiz-Zapatero J., Alonso D., García-García C., Nicola A., Mootoovaloo A., Sullivan J. M., Bonici M., Ferreira P. G., 2024, *Open J. Astrophys.*, 7, 11  
 Spurio Mancini A., Piras D., Alsing J., Joachimi B., Hobson M. P., 2022, *MNRAS*, 511, 1771  
 Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152  
 Taylor J. F., Ashdown M. A. J., Hobson M. P., 2008, *MNRAS*, 389, 1284  
 Tegmark M., Taylor A. N., Heavens A. F., 1997, *ApJ*, 480, 22  
 The LSST Dark Energy Science Collaboration, 2018, preprint (arXiv:1809.01669)

## APPENDIX A: HMC AND NUTS

## A1 Hamiltonian Monte Carlo

HMC is a sampling technique, developed by Duane et al. (1987) and the method is inspired by molecular dynamics. The motion of molecules follows Newton’s law of motions and it can be formulated as Hamiltonian dynamics. In short, we label the parameters governing the full theoretical model the ‘position variables’, and introduce an equal number of ‘momentum variables’. Starting from an initial state, a final state is proposed, related to the initial state through a Hamiltonian trajectory using the log-posterior as an effective potential, and a kinetic term governed by a mass matrix. A clear benefit of this sampling scheme is that the proposed state has, by construction a high probability of acceptance, due to energy conservation. Moreover, the sampler is now guided by the momentum variables, and is therefore less prone to random transitions. We refer the reader to Neal (2011), who provides an in-depth and pedagogical review on HMC.

In cosmology, HMC has been adopted in various cases, with the main goal being to sample the posterior distribution of some desired quantities, for example, the power spectrum, in an efficient way. Hajian (2007) implemented an HMC to sample cosmological parameters and found it to be more efficient by a factor of 10, compared to existing MCMC method in COSMOMC. Taylor, Ashdown & Hobson (2008) developed an HMC sampling scheme to sample the CMB power spectrum while Jasche & Lavaux (2019) developed an ambitious sampling framework, Bayesian origin reconstruction from galaxies consisting of the HMC sampler, to sample over million of parameters. Most recently, Campagne et al. (2023), Ruiz-Zapatero et al. (2024), and Piras et al. (2024) have developed auto-differentiable frameworks to analyse angular power spectra with the aim of enabling gradient-based samplers.

Here, we describe an implementation of the HMC algorithm tailored towards using the mean and the first derivative. The second derivative can also be used to tune the mass matrix but one could also resort to using a small MCMC chain to estimate it. In this section, we are assuming that one is updating the mass matrix as we sample the parameters of the system using the second derivative. Suppose  $\Omega \in \mathbb{R}^d$  is a  $d$ -dimensional position vector (the parameters of interest) and  $\mathbf{q} \in \mathbb{R}^d$  is a  $d$ -dimensional momentum vector. The full state space has  $2d$  dimensions and the system is described by the Hamiltonian:

$$\mathcal{H}(\Omega, \mathbf{q}) = U(\Omega) + \mathcal{K}(\mathbf{q}). \quad (\text{A1})$$

In Bayesian applications, the potential energy is simply the negative log-posterior, that is,

$$U(\Omega) = - [\log p(\mathbf{x}|\Omega) + \log \pi(\Omega)] \quad (\text{A2})$$

while the kinetic energy term is

$$\mathcal{K}(\mathbf{q}) = \frac{1}{2} \mathbf{q}^T \mathbf{M}^{-1} \mathbf{q} + \frac{1}{2} \log |\mathbf{M}| + \text{constant}. \quad (\text{A3})$$

Note that the determinant of the mass matrix is not ignored in this variant of HMC, since the mass matrix is updated after every leapfrog move. In other words, the mass matrix is in fact a function of position as we sample the parameter space. The partial derivatives of the Hamiltonian control the evolution of the position and momentum variables, that is,

$$\begin{aligned} \frac{d\Omega_i}{dt} &= \frac{\partial \mathcal{H}}{\partial q_i} \\ \frac{dq_i}{dt} &= -\frac{\partial \mathcal{H}}{\partial \Omega_i} \end{aligned} \quad (\text{A4})$$

---

**Algorithm 1** The leapfrog algorithm

---

```

procedure LEAPFROG( $\Omega_0, \mathbf{q}_0, N_L, \epsilon, \mathbf{M}$ )
   $\mathbf{q} = \mathbf{q}(t) - \epsilon \frac{\partial \mathcal{U}}{\partial \Omega}(\Omega_0)/2$ 
  for  $i : 1 \rightarrow N_L$  do
     $\Omega \leftarrow \Omega + \epsilon \mathbf{M}^{-1} \mathbf{q}$ 
     $\mathbf{q} \leftarrow \mathbf{q} - \epsilon \frac{\partial \mathcal{U}}{\partial \Omega}(\Omega)$ 
  end for
   $\mathbf{q} = \mathbf{q} - \epsilon \frac{\partial \mathcal{U}}{\partial \Omega}(\Omega)/2$ 
  return  $(\Omega, -\mathbf{q})$ 
end procedure

```

---

Crucially, as discussed by Neal (2011), three key aspects of this formulation are: (1) the Hamiltonian dynamics are reversible, (2) the Hamiltonian is preserved, and (3) the volume is preserved. In order to solve the system of differential equations, one typically resorts to numerical techniques, for example, Euler's method. A better alternative is the leapfrog algorithm, which is summarized below:

$$\begin{aligned} \mathbf{q}\left(t + \frac{\epsilon}{2}\right) &= \mathbf{q}(t) - \frac{\epsilon}{2} \frac{\partial \mathcal{U}}{\partial \Omega}[\Omega(t)] \\ \Omega(t + \epsilon) &= \Omega(t) + \epsilon \mathbf{M}^{-1} \mathbf{q}\left(t + \frac{\epsilon}{2}\right) \\ \mathbf{q}(t + \epsilon) &= \mathbf{q}\left(t + \frac{\epsilon}{2}\right) - \frac{\epsilon}{2} \frac{\partial \mathcal{U}}{\partial \Omega}[\Omega(t + \epsilon)], \end{aligned}$$

where  $\epsilon$  is the step size parameter. The leapfrog algorithm is shown in Algorithm 1. In particular, one takes a half-step in the momentum at the beginning before doing  $N_L$  full steps in the momentum and position variables, followed by a final half-step in momentum. The new proposed state is accepted according to Algorithm 2. The full HMC algorithm is presented in Algorithm 3.

---

**Algorithm 2** The acceptance criterion

---

```

procedure CRITERION( $\Omega_0, \mathbf{q}_0, \Omega, \mathbf{q}$ )
   $u \sim \mathcal{U}[0, 1]$ 
  if  $u < \exp[\mathcal{H}(\Omega_0, \mathbf{q}_0) - \mathcal{H}(\Omega, \mathbf{q})]$  then
    return  $\Omega$ 
  else
    return  $\Omega_0$ 
  end if
end procedure

```

---

To ensure a good performance with the HMC method, an appropriate mass matrix, step size, and number of leapfrog moves are recommended. We briefly touch upon these concepts below.

### A1.1 Mass matrix

If we can calculate the Hessian matrix, the latter can be used to estimate the mass matrix in the HMC. The inverse of this matrix,  $\mathbf{H}^{-1}$ , gives an estimate of the covariance of the parameters at any point in parameter space. In fact, this covariance is used in the leapfrog algorithm to take a step in the parameters. In Section A1.2, we will look into how we can use this Hessian matrix to set the mass matrix, in order to ensure efficient sampling. Note that we do not adapt the mass matrix at every step of the sampling procedure. Otherwise, this will be very expensive. For example, in the NUMPYRO library, during the warmup phase, one can adapt the mass matrix and then the learned mass matrix is fixed throughout the sampling process (Phan, Pradhan & Jankowiak 2019).

### A1.2 Step size

In the leapfrog algorithm, we also have to specify the step size,  $\epsilon$ . If we assume that the posterior distribution of the parameters are roughly Gaussian, then the derivative of the potential energy is approximately  $\mathbf{H}\tilde{\Omega}$ , where  $\tilde{\Omega}$  is the difference between any point in parameter space and the mean. Hence, we can write a single application of the leapfrog method as:

$$\begin{pmatrix} \tilde{\Omega}(t + \epsilon) \\ \mathbf{q}(t + \epsilon) \end{pmatrix} = \mathbf{Q} \begin{pmatrix} \tilde{\Omega}(t) \\ \mathbf{q}(t) \end{pmatrix} \quad (\text{A5})$$

where

$$\mathbf{Q} = \begin{pmatrix} \mathbb{I} - \frac{\epsilon^2}{2} \mathbf{M}^{-1} \mathbf{H} & \epsilon \mathbf{M}^{-1} \\ \frac{\epsilon^3}{4} \mathbf{H} \mathbf{M}^{-1} \mathbf{H} - \epsilon \mathbf{H} \mathbb{I} & -\frac{\epsilon^2}{2} \mathbf{M}^{-1} \mathbf{H} \end{pmatrix}.$$

The leapfrog move will only be stable if the eigenvalues of the matrix  $\mathbf{Q}$  have squared magnitude of the order of unity. We can write the following characteristic equation:

$$\left| \lambda^2 \mathbb{I} - 2\lambda \left( \mathbb{I} - \frac{\epsilon^2}{2} \mathbf{M}^{-1} \mathbf{H} \right) + \mathbb{I} \right| = 0 \quad (\text{A6})$$

and from the above equation, we can remove the dependence of  $\epsilon$  on  $\mathbf{M}$  by setting it to the Hessian matrix, that is,  $\mathbf{M} = \mathbf{H}$ . Note that one would maximally decorrelate the target distribution if the mass matrix is set to the Hessian matrix. This can be seen from the duality and volume preservation of the  $(\Omega, \mathbf{q})$  phase space. A transformation of the kinetic energy (the momentum variables) leads to an equivalent change in the potential energy (the position variables).

In 1D, if the mass matrix is set to unity, Neal (2011) argues that a step size of  $\epsilon < 2\sigma$ , where  $\sigma$  is the width of the distribution, leads to stable trajectory. In higher dimensions, assuming we have set the mass matrix to the Hessian matrix, solving the characteristic equation (equation A6), leads to

$$\lambda^2 - 2 \left( 1 - \frac{\epsilon^2}{2} \right) \lambda + 1 = 0. \quad (\text{A7})$$

When  $\epsilon < 2$ , this leads to stable trajectories. In practice, we would set  $0 < \epsilon < 2$ . Alternatively, libraries such as NUMPYRO have the option to adapt the step size during the warmup phase via a primal dual averaging (PDA) scheme (Bingham et al. 2019; Phan et al. 2019).

A very small value of  $\epsilon$  would increase the number of steps (and therefore of posterior gradient evaluations), while a large value may lead to numerically unstable trajectories. A similar prescription for choosing the mass matrix and setting the step size, is provided in the appendix of Taylor et al. (2008), where their main objective was to sample the CMB power spectrum.

**Algorithm 3** The Hamiltonian Monte Carlo algorithm

---

```

procedure HMC( $\Omega_0, \epsilon, N_L, N, M$ )
   $i = 0$ 
   $\mathbf{q}_0 \sim \mathcal{N}(\mathbf{0}, M)$ 
   $S = \{\}$ 
  while  $i < N$  do
     $(\Omega, \mathbf{q}) \leftarrow \text{LEAPFROG}(\Omega_0, \mathbf{q}_0, N_L, \epsilon, M)$ 
     $\Omega \leftarrow \text{CRITERION}(\Omega_0, \mathbf{q}_0, \Omega, \mathbf{q})$ 
     $S \leftarrow S \cup \{\Omega\}$ 
     $\mathbf{q}_0 \sim \mathcal{N}(\mathbf{0}, M)$ 
     $\Omega_0 \leftarrow \Omega$ 
     $i++$ 
  end while
  return  $S$ 
end procedure

```

---

*A1.3 Number of leapfrog moves*

Another parameter which we have to set is the number of leapfrog moves. There is no straightforward way to choose this parameter. We would prefer to avoid periodic trajectories that end close to the starting position. We would also like to minimize the number of gradient computations, in order to reach the pre-specified number of samples earlier. If a bad step size is chosen, the Hamiltonian grows with the number of leapfrog moves and the probability of acceptance decreases considerably.

**A2 No-U-turn sampler**

The NUTS has been designed to circumvent the need to tune the step size,  $\epsilon$  and the number of leapfrog moves,  $N_L$ .

In order to adaptively adjust the step size during the sampling process, NUTS uses a technique referred to as PDA. PDA helps address this challenge by adaptively tuning the step size based on the trajectory of the sampler. It does so by maintaining a running average

of the gradients of the log posterior over the course of the sampling process. This average is used to update the step size such that it scales appropriately with the curvature of the log posterior. The primal-dual averaging algorithm ensures that the step size adjusts smoothly and efficiently, allowing NUTS to explore the parameter space effectively without requiring manual tuning from the user. By continuously updating the step size based on the local geometry of the posterior distribution, NUTS with PDA can achieve faster convergence and higher sampling efficiency compared to fixed-step-size methods.

As discussed in the previous section, selecting the number of leapfrog moves can be tricky and the challenge is to find a metric to determine if the trajectory is too short, too long or just right. NUTS dynamically selects the trajectory length during its sampling process based on the “No-U-Turn” criterion. When NUTS starts exploring the parameter space from a given starting point, it extends its trajectory by recursively building a binary tree of states. At each step, NUTS evaluates whether to continue extending the trajectory in a particular direction or to stop. This decision is guided by the no-U-turn criterion, which checks whether the sampler is doubling back on itself. The trajectory length is determined dynamically as NUTS extends the tree. If the sampler is still exploring promising regions of the posterior distribution, it continues to double the trajectory length. However, if the sampler starts to turn back, indicating that it has likely overshoot a relevant part of the distribution, the trajectory is truncated. By dynamically adjusting the trajectory length based on the no-U-turn criterion, NUTS ensures that it explores the parameter space efficiently without needing to specify the number of steps in advance. This adaptive behaviour allows NUTS to effectively explore complex and high-dimensional distributions, leading to faster convergence and higher sampling efficiency compared to fixed-length trajectory methods. For an in-depth explanation of the NUTS algorithm, we refer the reader to Hoffman & Gelman (2011).

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.