

Playing dice with the universe: Bayesian statistical analyses of cosmological models and new observables
Cañas Herrera, G.

Citation

Cañas Herrera, G. (2022, October 19). Playing dice with the universe: Bayesian statistical analyses of cosmological models and new observables. Casimir PhD Series. Retrieved from https://hdl.handle.net/1887/3483658

Version: Publisher's Version

License: License agreement concerning inclusion of doctoral thesis in the

Institutional Repository of the University of Leiden

Downloaded from: https://hdl.handle.net/1887/3483658

Note: To cite this publication please use the final published version (if applicable).

Part I Constraints on inflationary models

Chapter 2

Bayesian reconstruction of the inflaton's speed of sound using CMB data

This chapter is based on:

Bayesian reconstruction of the inflaton's speed of sound using CMB data Guadalupe Cañas-Herrera, Jesús Torrado Cacho and Ana Achúcarro (2021, June), Physical Review D 103, 12, 123531, arXiv:2012.04640.

2.1 Introduction

The standard cosmological model (Λ CDM) is currently favored by the available data. It assumes that primordial fluctuations are Gaussian and defined by an almost scale-invariant primordial power spectrum (see Chapter 1). These assumptions do not point to any particular origin, although the simplest inflationary model, canonical slow-roll single-field inflation, naturally predicts them. By contrast, other models of inflation predict deviations from the near scale-invariant spectrum in the form of features. If ever detected, they would open a new window of research in the field of primordial dynamics. See i.e. (Chluba et al., 2015; X. Chen, 2010; Slosar et al., 2019; Palma, 2015).

The study of features of primordial origin can be done within an Effective Field Theory approach. Within this scenario, features can be produced by the time dependence of primordial functions such as the slow-roll parameters or the speed of sound of the effective inflaton (the adiabatic mode). In particular, small, soft and transient reductions in the inflaton's speed of sound produce such correlated localized oscillatory features in the n-point correlation functions. In the 3-point function (or bispectrum), these localized oscillations present a distinct difference in phase between the squeezed and equilateral configurations (Achúcarro, Gong, Palma, & Patil, 2013).

The Planck Collaboration (Planck Collaboration, Aghanim, Akrami, Ashdown, et al., 2020b) searched for deviations of the canonical scenario in its last release of data. Nevertheless, they did not find strong evidence in the context of features in

2.1. INTRODUCTION

the primordial power spectrum (Planck Collaboration, Akrami, et al., 2020). They included, for the first time, a joint search of correlated simple features in the primordial power spectrum and in the bispectrum, also without significant results. However, the Planck Collaboration has not studied in detail different feature templates such as the above mentioned ones due to small and transient reductions of the inflaton's speed of sound. This motivates us to continue our previous study (Achúcarro, Atal, Ortiz, & Torrado, 2014; Achúcarro, Atal, Hu, Ortiz, & Torrado, 2014; Hu & Torrado, 2015; Torrado, Hu, & Achúcarro, 2017), in preparation for a future release of the Planck bispectrum likelihood or for future investigation in light of incoming Large Scale Structure surveys.

Most of the time, the study of features in both observables (primordial power spectrum and higher correlation functions) are model-dependent, both regarding their physical origin and the ansatz used. In our latest paper in this series (Torrado et al., 2017), we already pointed out the need for testing more flexible feature templates to mitigate the dependence on the ansatz. Within this approach, we can test whether multiple and consecutive reductions of the inflaton's speed of sound can take place consecutively, a possibility already pointed out in the previous work (Torrado et al., 2017). Furthermore, reconstructing the inflaton's speed of sound allows us to test more complex feature templates with variable amplitude and oscillation frequency, which implies more possibilities to fit well-motivated deviations from Λ CDM beyond those that only used a pre-defined ansatz for the features.

Reconstructions at the level of the primordial power spectrum have already been attempted (Planck Collaboration et al., 2016; Hazra, Shafieloo, & Souradeep, 2013; Hunt & Sarkar, 2014; Ravenni, Verde, & Cuesta, 2016; Durakovic, Hunt, Mukherjee, Sarkar, & Souradeep, 2018; Ballardini, 2019; Durakovic, Hunt, Patil, & Sarkar, 2019; W. J. Handley, Lasenby, Peiris, & Hobson, 2019). However, there is not enough constraining power in Planck's power spectrum alone to decide on a particular model for the features. Model-informed reconstructions have the advantage of increasing the constraining power by adding the information contained in higher-order correlation functions; but this is only possible if the constraints of the theoretical model are properly imposed on the reconstructed spectrum, so that it will always lead to a consistent prediction. The task of imposing these physical constraints along the reconstruction is non trivial (Appleby, Gong, Hazra, Shafieloo, & Sypsas, 2016). It is advisable, instead, to reconstruct the primordial dynamics directly. In our case, we reconstruct the inflaton's speed of sound: the timing, intensity and rate of its reduction. Since we are reconstructing the underlying function leading to the correlated features, it is not only guaranteed that we will obtain a consistent bispectrum feature prediction using power spectrum data alone, but we will also be able to use both data sets simultaneously to get a more stringent reconstruction once a bispectrum likelihood has been released.

In this chapter, we develop a new analysis pipeline that uses Gaussian Processes (GPs), a hyper-parametric regression technique, to model the inflaton's speed of sound profile. The analytic nature of GPs makes easy to impose the constraints of the theoretical model, which involve derivatives of the reconstructed function. For a given number of nodes in the GP, we construct a prior on the hyper-parameters of

the GP model (the position of the nodes and the correlation length), that maximizes entropy with respect to the bare physical constraints. In this way, we verify that nodes are not placed wherever they would lead to an unphysical reconstruction, and that the density with which the hyper-parameters are explored reproduces the measure of the physical prior (W. Handley & Millea, 2019; Gariazzo & Mena, 2019; Planck Collaboration, Aghanim, Akrami, Ashdown, et al., 2020b).

We test our new pipeline against Planck 2018 temperature, polarization and lensing CMB angular power spectrum data, obtaining corresponding posteriors of the parameters of interest and several *maxima a posteriori*. Our results do not only reproduce our previous findings (Torrado et al., 2017), but allow for combinations of multiple consecutive reductions as well as more complex shapes.

This chapter is organized as follows. In section 2.2 we review the theoretical framework for inflationary correlated features in the primordial power spectrum due to transient reductions in the speed of sound. In section 2.3, we explain the methodology used to generate features in the primordial power spectrum: the parametrization for the reduction in the speed of sound (2.3.1), the chosen priors for the the different parameters (2.3.2) and the computational procedure (2.3.3). In section 5.5, we present the results corresponding to the fitting of features using the CMB angular power spectrum. Finally, we discuss the results, draw our conclusion and show prospective work for the future in section 2.5.

2.2 Theoretical model

We follow the Effective Field Theory (EFT) of inflationary perturbations (Cheung et al., 2008) to characterize the fluctuations of comoving curvature perturbations around an inflating cosmological background. It starts with an effective action for the Goldstone boson of cosmic time diffeomorphisms $\pi(t, \mathbf{x})$. This Goldstone boson is related to the comoving curvature perturbation $\mathcal{R}(t, \mathbf{x})$ through the relation $\mathcal{R} \approx -H(t)\pi(t, \mathbf{x})$, with the Hubble parameter $H(t) \equiv \dot{a}/a$, with a being the scale factor (where the dot denotes derivatives with respect to cosmic time t). The effective single field action for π up to second order is given by

$$S_2 = \int d^4x a^3 M_P^2 \epsilon_1 H^2 \left[-\frac{\dot{\pi}^2}{c_s^2} + \frac{(\partial_i \pi)^2}{a^2} \right], \tag{2.1}$$

where $M_P = 1/\sqrt{8\pi G}$ is the reduced Planck Mass in natural units $c = \hbar = 1$, $\epsilon_1 \equiv -\dot{H}/H^2$ is the first slow-roll parameter and $c_s \equiv c_s(t)$ is the time-dependent speed of sound.

The effective single field action up to third order, neglecting higher order slow-roll corrections ($\sim \mathcal{O}(\epsilon_1^2)$) and assuming $\dot{\pi}^3$ to be small and approximately constant reads

2.2. THEORETICAL MODEL

(Achúcarro, Gong, et al., 2012),

$$S_{3} = \int \left[d^{4}x a^{3} M_{P}^{2} \epsilon_{1} H^{2} - 2H s c_{s}^{-2} \pi \dot{\pi}^{2} - (1 - c_{s}^{-2}) \dot{\pi} \left(\frac{\dot{\pi}^{2}}{c_{s}^{2}} - \frac{(\partial_{i} \pi)^{2}}{a^{2}} \right) \right], \tag{2.2}$$

where $s \equiv s(t)$ parameterizes the change in the speed of sound $c_s(t)$ defined as

$$s \equiv \frac{\dot{c}_s(t)}{(c_s(t)H)}. (2.3)$$

The physical details of the theory are encoded in the speed of sound c_s and in its corresponding rate of change given by s. The speed of sound c_s accounts for the effects of integrating out the heavy fields within the effective action. To get an insight of what this variable $c_s(t)$ means, we look at the particular case of an effective theory for the comoving curvature perturbation \mathcal{R} , when a strong turn in the inflationary trajectory in multifield space is supported by a heavy field \mathcal{F} with "effective mass" M_{eff} . In this case, the curvature perturbation \mathcal{R} is kinetically coupled to the heavy field \mathcal{F} . This effective action is similar to the EFT of inflation equation (2.1), with the speed of sound c_s of the adiabatic perturbation \mathcal{R} given by (Achúcarro, Gong, Hardeman, Palma, & Patil, 2011a, 2011b)

$$c_s^{-2} = 1 + \frac{4\Omega^2}{k^2/a^2 + M_{\text{eff}}},$$
 (2.4)

where Ω is the the angular velocity when there is a turn in the inflationary trajectory, inducing a momentary reduction on the speed of sound c_s (Achúcarro, Gong, et al., 2012). The effect of this variable speed of sound c_s can be seen in the primordial power spectrum $\mathcal{P}_{\mathcal{R}}$, in the bispectrum $\mathcal{B}_{\mathcal{R}}$ and in higher-order correlation functions. In particular, transient variations of c_s produce localized oscillatory and correlated features in both $\mathcal{P}_{\mathcal{R}}$ and $\mathcal{B}_{\mathcal{R}}$ (Achúcarro, Atal, et al., 2012). Generally, $c_s(t)$ encodes the effect of derivative interactions.

Under the assumption of small, mild and transient reductions of the speed of sound c_s , the modifications in the primordial power spectrum of curvature perturbations with respect to the almost-invariant power law spectrum, $\Delta \mathcal{P}_{\mathcal{R}}/\mathcal{P}_{\mathcal{R}0}$, were already calculated (Achúcarro et al., 2013). The quadratic action of EFT of inflation, equation (2.1), is divided into a free part (resembling single field inflation, with $c_s = 1$) and a small perturbation:

$$S_{2} = \int d^{4}x \, a^{3} M_{P}^{2} \epsilon H^{2} \left(\dot{\pi}^{2} - \frac{(\partial_{i} \pi)^{2}}{a^{2}} \right) - \int d^{4}x \, a^{3} M_{P}^{2} \epsilon H^{2} \left(\left(1 - c_{s}^{-2} \right) \dot{\pi}^{2} \right), \tag{2.5}$$

Transitioning from cosmic time t to the conformal time τ , so that $d\tau = dt/a(t)$, using the *in-in* formalism (Weinberg, 2005) and the following definition of the variable u,

$$u(\tau) \equiv (1 - c_s^{-2}(\tau)),$$
 (2.6)

the change in the primordial power spectrum $\Delta \mathcal{P}_{\mathcal{R}}$ is given by the Fourier transform of the reduction in the speed of sound c_s :

$$\frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R},0}} = k \int_{-\infty}^{0} d\tau u(\tau) \sin(2k\tau), \qquad (2.7)$$

where $\mathcal{P}_{\mathcal{R}0}$ is the power law expressed in equation (1.39).

2.3 Methodology

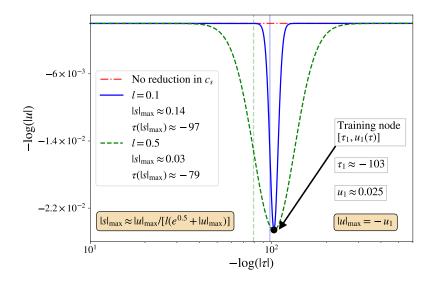


Figure 2.1: Example of reductions of the inflaton's speed of sound $u(\tau)$ in logarithmic space GP with a single node at (τ_1, u_1) . The reductions peak at $\tau_1 \approx 103$ with a maximum reduction value $-u_1 = |u|_{\text{max}} \approx 0.025$. The width of each of the reductions (given by the correlation length l, which, for a single GP node, plays the role of the standard deviation) is different, being the green-dashed parametrization (l = 0.5) milder than the solid blue one (l = 0.1). The rate of change of $u(\tau)$, see equation (2.3), can be approximated as $|s|_{\text{max}} = |u|_{\text{max}}/[l(e^{0.5} + |u|_{\text{max}})]$ using this parametrization. The vertical lines indicate the values of τ at which $|s|_{\text{max}}$ is reached in each case.

The reduction of the speed of sound and its rate of change are encoded in $u(\tau)$ and $s(\tau)$ respectively. We aim to use current cosmological data (the temperature, polarization and lensing power spectrum of Planck 2018) to estimate them given the theoretical framework presented in section 2.2. To do that, we use *Bayesian inference*, as explained Chapter 1.

2.3.1 Reconstruction model for the reduction in the speed of sound

We use Gaussian Processes (GPs) (Rasmussen & Williams, 2005) as an interpolator for reconstructing the speed of sound of the inflaton. The *mean* curve of a GP,

2.3. METHODOLOGY

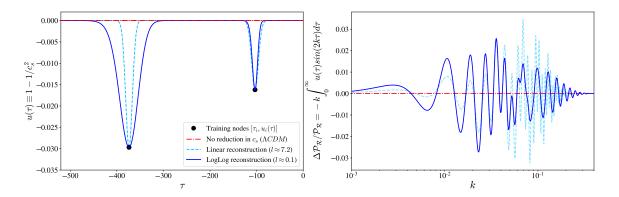


Figure 2.2: **Left panel**: Reconstruction of $u(\tau)$ using a GP on $\log|u(\log|\tau|)|$ (continuous line) and on $u(\tau)$ (dashed line), both using as training nodes $(\tau, u) \in \{(-100, -0.016), (-375, -0.03)\}$, and correlation length for each model such that the width of the mode corresponding to the first node is similar (notice the difference in width of the mode of the second one). As the reconstruction is done in conformal time τ , the x-axis is always negative and $\tau = 0$ indicates the end of inflation. **Right panel**: Corresponding feature in the primordial power spectrum. Notice how, despite the similar position of the training nodes, their features look quite different: similar in the leftmost oscillations (corresponding to the node at $\tau = -100$), but very different after that, due to the broader width of the first mode.

which we use to represent the speed of sound evolution, is smooth by construction (given our choice of kernel, see below) and naturally returns to a baseline value away from the *nodes* of the interpolator, which is useful for representing the transient character of the speed of sound reductions. The length scale over which the return to the baseline happens is called *correlation length*, and it is the same for all individual nodes in a particular realisation. The particular properties of the correlation between nodes is given by the *kernel* of the GP, which we choose to be a *squared exponential* kernel. This means that when the interpolator is defined by a single node placed at (x_1, y_1) , the interpolating curve looks like a (non-normalised) Gaussian peaking at the node's position, with standard deviation equal to the correlation length l, i.e. $y(x) = y_1 \exp \left[-1/2 (x - x_1)^2/l^2\right]$.

We aim to reconstruct $u(\tau) \equiv (1 - c_s^{-2}(\tau))$. Since $u(\tau)$ is a negative quantity, it makes sense to reconstruct the logarithm of -u = |u|, to guarantee that the GP interpolator, once exponentiated, conserves sign. On the other hand, there is a choice to be made about the scale of the conformal time axis: whether to reconstruct $\log |u(\tau)|$ or $\log |u(\log |\tau|)|$. We show results for the latter case in this section for illustration purposes.

As explained above, in the case where a single node is placed at $(\tau_1, \log(-u_1))$ the reconstruction of $\log(-u(\tau))$ corresponds to a single transient reduction given by a log-normal function of conformal time, whose maximum occurs exactly at the node. The parameters that we would try to infer from the data would be the position of the node $(\tau_1, \log(-u_1))$ and the correlation length l representing the standard deviation of the log-normal. The rate of change of the reduction, of interest in our

theoretical framework, would peak approximately at $|s|_{\text{max}} \approx 0.5 l |u_1|/(\exp(0.5) + |u_1|)$. See Figure 2.1 for an example.

To reconstruct $\log|u(\log|\tau|)|$ as a generalization of the previous case using GPs, we choose a number i of training nodes¹ $(\log|\tau_i|, \log|u_i|)$ where $\log|u_i| = \log|u(\log|\tau_i|)|$ (see Figure 2.2), to which we fit a GP with kernel function

$$\kappa(\log|\tau_i|, \log|\tau_{i+1}|; l) = c^2 \exp\left\{-\frac{1}{2} \left(\frac{\log|\tau_i| - \log|\tau_{i+1}|}{l}\right)^2\right\},$$
(2.8)

where c is the output scale, and l the correlation length. The output scale c plays no role in this approach, and can be fitted using maximum likelihood and then ignored. The correlation length will be sampled together with the position of the nodes. To compute the GPs, we use the Python package sklearn (Scikit-learn 0.19.1 documentation: Gaussian Processes, 2018). The mean of the GP is used as an interpolator for $\log |u(\log |\tau|)|$, and reads, in terms of the training nodes, as the matrix product:

$$\log|u(\log|\tau|)| = \kappa(\log|\tau|, \log|\tau_i|; l) \times \left[\kappa(\log|\tau_i|, \log|\tau_j|; l)\right]^{-1} \log|u_j|, \qquad (2.9)$$

where the first kernel function κ is a vector of evaluations at the requested $\log |\tau|$ combined with each of the training $\log |\tau_i|$, the second one is the matrix of evaluations of κ for each pair of training nodes (i,j), and the final term is the vector of training $\log |u_j|$. Once $u(\tau)$ is generated, we calculate $s(\tau)$ from equation (2.3) numerically, which we can rewrite more conveniently as

$$s(\tau) = \frac{1}{2} \frac{u}{1 - u} \frac{d \log u}{d \log |\tau|}$$

$$= \frac{1}{2} \frac{u}{u - 1} \frac{1}{l^2} \left[(\log |\tau| - \log |\tau_i|) \kappa(\log |\tau|, \log |\tau_i|; l) \right] \times$$

$$\left[\kappa(\log |\tau_i|, \log |\tau_i|; l) \right]^{-1} \log |u_i|, \qquad (2.10)$$

where we have taken the derivative after substituting u by the mean of the GP defined in equation (2.9). Notice that this reproduces the matrix product in equation (2.9), just changing the first vector. Finally, we compute the power spectrum feature of equation (2.7) from a fine sampling of the GP using the FFTLog algorithm (Hamilton, 2000; fftlog - python wrapper FFTLog, 2016). The density and limits of the $\log |\tau|$ sampling for the FFTLog are chosen adaptively to minimise computational costs and guarantee the accurate computation of the transform.

The most consequential difference of the choice between linear and logarithmic τ in the GP will show up whenever we have nodes separated by a distance much larger than the correlation length, appearing as isolated (log)Gaussians: in the linear case, their width in τ will be similar, whereas for the logarithmic one, the width will scale logarithmically (see Figure 2.2). Looking at the first equality in equation (2.10), and seeing how s depends on the logarithmic derivative on τ , it is easy to see that the

 $^{^{1}}$ Notice that our use of GPs as interpolators does not involve machine-learning, but we are borrowing the term training from its literature.

2.3. METHODOLOGY

linear parameterization is going to struggle to place two or more nodes away from each other, since $s(\tau)$ will peak at highly different values in each of them, making it hard not to violate the perturbativity bounds on s (see subsection 2.3.2).

Thus for the primary results in this chapter, we model $\log |u(\log |\tau|)|$. To mitigate excessive sensitivity to the prior of our results, we also perform a reconstruction in $u(\tau)$. Notice that by modelling u and not $\log |u|$ we need to deal with cases in which $u(\tau)$ goes positive, by assigning it null prior density. However, those are generally disfavored by the data (require large l compared to the distance between nodes), and the large difference between the $\log |u(\log |\tau|)|$ and the $u(\tau)$ reconstruction is useful for assessing prior sensitivity.

2.3.2 Parameters and Priors

The action described in equations (2.1) and (2.2) is perturbative in terms of $(1 - 1/c_s^2)$. It implies that the reduction in the speed of sound, c_s , cannot be too big $(|u| \ll 1)$ and the rate of change in the reduction cannot be too fast $(|s| \ll 1)$. Also, the contributions of the slow-roll corrections ϵ_1, ϵ_2 have to be smaller than those of the variable speed of sound c_s . We need to impose these conditions for all values of τ , but it is enough to restrict to the point where $u(\tau)$ and $s(\tau)$ take their maximum value $(|u|_{\text{max}}, |s|_{\text{max}})$. Note that imposing the perturbative limit on $|s|_{\text{max}}$ satisfies the consistency conditions in (Céspedes, Atal, & Palma, 2012; Adshead & Hu, 2014; Cannone, Bartolo, & Matarrese, 2014). In short:

$$\max(\epsilon_1, \epsilon_2) \ll \max(|u|_{\max}, |s|_{\max}) \ll 1. \tag{2.11}$$

In (Torrado et al., 2017) we argued that this condition could be naturally imposed by a prior Beta(5,5) on $\max(\log_{10}|u|_{\text{max}},\log_{10}|s|_{\text{max}})$ between the extremes in equation (2.11) (see left part of Figure 2.3), the logarithm coming from the difference in order of magnitude between both bounds.

Contrary to (Torrado et al., 2017), in this work $|u|_{\max}$ and $|s|_{\max}$ are not sampled directly. Instead, the parameter space for the feature consists of the position of the nodes $\{(\tau_i, u_i = u(\tau_i))\}$ and the correlation length l. Thus, the total number of feature parameters is 2N+1 for a number N of nodes. Imposing the Beta prior described above is not as simple as sampling the GP parameters from some prior, computing $|u|_{\max}$ and $|s|_{\max}$ along the reconstruction, and multiplying by the Beta density. That procedure will likely introduce undesired information in the shape of under- or overdensities in the probability induced on $(|u|_{\max}, |s|_{\max})$, which would finally diverge significantly from a Beta. The correct way to proceed so that the induced probability on the physical parameters is the desired one is by constructing a distribution on the parameters of the nodes that $maximises\ entropy$ with respect to the desired one, which can be computed, according to (W. Handley & Millea, 2019), as

$$\pi_{\text{pert, maxent}}(\tau_i, u_i, l) = \frac{\pi_0(\tau_i, u_i, l) \pi_{\text{pert}}(|u|_{\text{max}}, |s|_{\text{max}})}{P(|u|_{\text{max}}, |s|_{\text{max}} | \pi_0)},$$
(2.12)

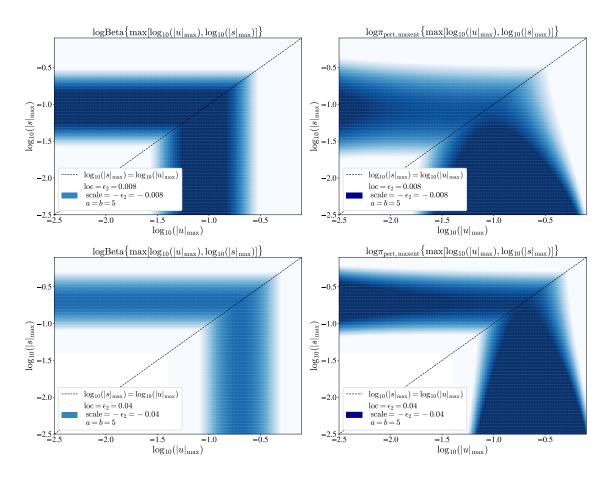


Figure 2.3: A description of the Bayesian priors adopted in this study. **Left panels**: Desired Beta density distribution for the parameters $log_{10}(|u|_{\text{max}})$ and $log_{10}(|s|_{\text{max}})$, with $\epsilon_2 = 0.008$ and $\epsilon_2 = 0.04$, respectively. **Right panels**: probability density distribution induced by π_0 on $log_{10}(|s|_{\text{max}})$ and $log_{10}(|u|_{\text{max}})$ (defined as P in the denominator of equation (2.12)), for a Gaussian-Processes reconstruction in $log|u(log|\tau|)|$ using 2 training nodes. The initial prior $\pi_0(\tau_i, u_i, l)$ induces overdensities at low $log_{10}(|s|_{\text{max}})$ due to this quantity being doubly-correlated to early-time position of the nodes and large correlation lengths; equation (2.12) corrects for this effect.

where $\pi_0(\tau_i, u_i, l)$ is some initial prior on the node parameters and $\pi_{\text{pert}}(|u|_{\text{max}}, |s|_{\text{max}})$ the Beta prior described above. The term in the denominator is the probability density induced by π_o on $(|u|_{\text{max}}, |s|_{\text{max}})$, which we compute from a Monte Carlo sample from π_0 using PolyChord (W. J. Handley et al., 2015a, 2015b). The Monte Carlo sample is fed to GetDist (Lewis, 2019) to construct a density estimator. Both the Beta prior and the maxent prior can be seen in Figure 2.3.

For π_0 , we choose log-uniform prior distributions on the correlation length l and on the training node location $(\tau_i, u_i)^2$. The bounds for the time-positions τ_i are chosen so that the feature falls in the CMB window function (features running from scales

²Note that we are sampling the training nodes and the correlation length in a logarithmic scale as we expect them to vary several orders of magnitude.

 $k \approx 10^{-3}$ to $k \approx 3 \times 10^{-1}$, though a larger region has been scanned as a consistency check (see subsection 2.4.1). The bounds for the amplitude of the reductions at the nodes, u_i , are chosen to generously fulfill the EFT condition in equation (2.11). Summarizing:

$$\pi_{0}(\log_{10}(|\tau_{1}|), \dots, \log_{10}(|\tau_{n}|), \log_{10}(|u_{1}|), \dots, \log_{10}(|u_{n}|), \log_{10}l) = \mathcal{U}(1.8 < \log_{10}(|\tau_{n}|) < 3.3) \times \prod_{i=n}^{2} \mathcal{U}(\log_{10}(|\tau_{i}|) < \log_{10}(|\tau_{i-1}|) < 3.3) \times \prod_{i=1}^{n} \mathcal{U}(-4 < \log_{10}(|u_{i}|) < 0) \times \pi(l)$$
(2.13)

where \mathcal{U} means a uniform distribution, and the prior on the time positions of the nodes includes sorting so that $\tau_i < \tau_{i-1}$ (*i* runs from 1...*n*). The prior on *l* is chosen so that it produces reasonable values of $|s|_{\text{max}}$. For each of the two reconstructions studied here, $\log |u(\log |\tau|)|$ and $u(\tau)$, the boundaries can be chosen as³

$$\pi_{\log|u(\log|\tau|)|}(l) = \mathcal{U}(-2 < \log_{10} l < 2) \quad \text{and}
\pi_{u(\tau)}(l) = \mathcal{U}(-2 < \log_{10} l < 3.3).$$
(2.14)

As an improvement on (Torrado et al., 2017), in this work we do not fix the value of the slow-roll parameters ϵ_1 and ϵ_2 as bounds of the perturbativity condition equation (2.11). Instead we let the bounds of the Beta distribution run dynamically, marginalizing over the slow-roll parameters. We use as priors uniform distributions $\mathcal{U}(0.0001 < \epsilon_1 < 0.05)$ and $\mathcal{U}(-0.06 < \epsilon_2 < 0.06)$, which encompass the Λ CDM posterior found for them in Planck 2018 (Planck Collaboration et al., 2016).

A diagram showing the different assumptions that enter the final prior is shown in Figure 2.4.

2.3.3 Data sets and sampler

To constrain the reduction of the speed of sound, we use the Planck 2018 polarized CMB and lensing data. In particular we use the product of the low multipole likelihoods lowT and lowE, the *unbinned* high- ℓ likelihood plik_TTTEEE and the lensing likelihood. We use the unbinned likelihoods because of the fast frequency of oscillations in the features, as was already pointed out in (Torrado et al., 2017).

We compute the changes to the CMB power spectra C_{ℓ} using the Boltzmann code CAMB (Lewis, 2013), modified accordingly to account for the increased sampling in k needed by the oscillatory features $\Delta \mathcal{P}_{\mathcal{R}}/\mathcal{P}_{\mathcal{R}0}$ in the primordial power spectrum. We sample over the parameter space described in section 2.3.2, i.e. the positions of the

³Notice that, while smaller values of l will result in sharper reductions with too high, forbidden $|s|_{\text{max}}$ values, larger values of l would result in small $|s|_{\text{max}}$ values which are actually allowed as long as $|u|_{\text{max}}$ fulfills equation (2.11). In any case, we are imposing these upper l boundaries for the main runs, since reductions with very small $|s|_{\text{max}}$ tend not to be easily distinguishable from changes in the background cosmological model (see subsection 2.4.1).

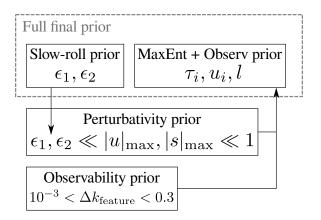


Figure 2.4: Diagram showing the structure of the prior. The conformal time τ_i of the nodes must fulfil that the feature happens within the observable CMB window. The physically-motivated perturbativity condition of equation (2.11) is imposed using maximum entropy (see text) on the position of the nodes and the correlation length. Since the value of the slow-roll parameters influences the prior on the position of the nodes via the lower bound of the perturbativity condition, the full prior (dashed grey box) is non-separable.

training nodes $\{(\tau_i, u_i = u(\tau_i))\}$, the correlation length l of the GP, and the kinetic slow-roll parameters ϵ_1 and ϵ_2 . We also allow for the possibility of tensor modes, as changes in the Sachs-Wolfe plateau caused by them could possibly be correlated with features at very large scales. We track as derived parameters the scalar tilt n_s , the tensor-to-scalar ratio r and the EFT parameters $(|u|_{\text{max}}, |s|_{\text{max}})$. We fix the rest of cosmological parameters of Λ CDM to the best fit of Planck 2018 with the present likelihoods, as well as the *nuisance* parameters of the likelihoods. Fixing the Λ CDM parameters is justified by previous sensitivity analyses in (Torrado et al., 2017), that we repeat here for the background Λ CDM parameters by exploring a broader range of τ and l than the one indicated above (see subsection 2.4.1, where we have also assessed the impact of fixing the nuisance parameters of the Planck likelihoods).

We obtain the posterior distribution of the parameters using the sampler PolyChord (W. J. Handley et al., 2015a, 2015b). We use this nested sampler since, from previous searches, we expect the posterior distributions of u_{max} and τ_i to be multi-modal. The handling of the priors, likelihoods, Boltzmann code and sampler is managed by the Bayesian framework Cobaya (Torrado & Lewis, 2021, 2019). The analysis of the posterior distributions is carried out using GetDist (Lewis, 2019).

We sample the posterior of two different parameterizations of the GP sound speed reconstruction: $\log|u|(\log|\tau|)$ and $u(\tau)$, in the following called simply logarithmic and linear parameterizations, respectively. We know the logarithmic parameterization is more stable numerically, as it consistently makes the reconstruction of $u(\tau)$ negative. However, we still use the runs in the linear parameterization for the purposes of assessing prior sensitivity. For the first sampling processes (up to three GPs nodes), we run Cobaya in parallel launching 8 MPI processes, each allowed to thread across 3 CPU cores. In the case of 4 nodes, we run Cobaya with 32 MPI processes, each

2.4. RESULTS

allowed to thread across only one single CPU core. The nested sampler PolyChord has been run with 1000 live points (which is far above the requirements for the current number of dimensions in the parameter space) and a stopping criterion of 0.01. The computation time varies depending on the number of training nodes in the GPs: from a few days with only 1 node, up to several weeks with 4 nodes.

All the maxima a posteriori (MAP) presented in the next section have been obtained running Py-BOBYQA (Cartis, Fiala, Marteau, & Roberts, 2018; Cartis, Roberts, & Sheridan-Methven, 2018) (a Python implementation of the BOBYQA algorithm (Powell, 2009), available via Cobaya), initialized on the relevant local maxima of the PolyChord samples.

2.4 Results

2.4.1 Consistency checks

Before presenting our results, we shortly discuss whether the assumptions made in previous sections were justified. In particular, we have tested whether we find clear posterior modes outside the (τ_i, l) prior region described section 2.3.2 (the *CMB window* prior), and whether in posterior modes either in our initial prior or in the broader region, the assumption of no-correlation with background cosmological parameters is fulfilled.

To do that, we produced a 1-node posterior sample in the logarithmic parameterization in the enlarged prior region $0 < \log_{10}(|\tau_i|) < 4.3$ and $2 < \log_{10} l < 10$, and let the background Λ CDM parameters vary. No significant modes were found outside the original, reduced prior region. We found mild modes in the region $1.8 < \log_{10}(|\tau_i|) < 3.3$ and $2 < \log_{10} l < 3.3$, which presented some degeneracy between Ω_m , n_s and the reconstruction parameters ($\rho \approx 0.17$), due to the fact that these features can be confused with the shape of the first and second acoustic peaks (already observed in (Torrado et al., 2017)). This justifies restricting ourselves to the prior described in 2.3.2, since any mode found outside of it would not be distinguishable from background cosmology.

The check for degeneracies between the c_s reconstruction parameters and the slow-roll parameters is of particular importance, since the latter determine the perturbative prior limits on the former (see equation 2.11 and Figure 2.3). We have found no significant degeneracies, neither in the tests described above nor in the final runs. We have reproduced the Planck Λ CDM posterior on the slow-roll parameters in all cases (see Figure 2.5).

Most of the results below have been run both in the linear and logarithmic parameterizations for the Gaussian Process (GP) reconstruction of the speed of sound profile $u(\tau)$. The results agree with each other, in particular, for the maxima a posteriori found at late conformal time (i.e. towards the end of inflation, with $-\tau_i$ of a few hundreds, where both parameterizations look similar). However, the logarithmic parameterization differs from the linear one when training nodes are thrown at early values of conformal time (i.e. $-\tau_i$ over 800), see Figure 2.2). This is due to modes

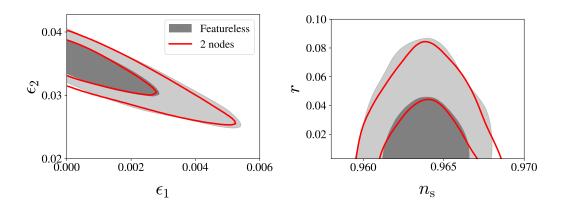


Figure 2.5: Posterior distributions of the primordial parameters: the kinetic slow-roll parameters ϵ_1 and ϵ_2 , and the derived spectral index n_s and tensor-to-scalar ratio r, for a 2-node reconstruction in the logarithmic parameterization (red line). The grey contours correspond to the featureless Λ CDM scenario. The correspondence between both posteriors is due to the absence of degeneracies between the c_s reconstruction parameters and the slow-roll parameters. Similar results are found for 1, 3 and 4 nodes.

of constant width in logarithmic scale getting broader the further we go along the axis of conformal time. In the 1-node case, the linear parameterization reproduces the results in (Torrado et al., 2017) (which uses a Gaussian ansatz in $u(\tau)$), whereas the logarithmic parameterization produces different 1-node posterior modes (see Sec. 2.6.1 in appendix A).

It is worth remarking that the use of the logarithmic parametrization does not compromise the flexibility of our reconstruction of $u(\tau)$. Even though the logarithmic parameterization reconstructs naturally profiles of $u(\tau)$ with broad reductions at earlier conformal times and narrower reductions at later times (which is preferred so that $|s|_{\text{max}}$ is not violated), narrow reductions at early conformal times can always be achieved by adding further nodes that would force the profile to return to zero. If the data and EFT conditions did allow for a narrow reduction at earlier conformal times, we would have seen it during the analysis of the posterior distributions when more than one training node was used.

Moreover, we have assessed the effect of fixing the nuisance parameters of the Planck likelihoods by running a minimizer around both the baseline, featureless Λ CDM model and the MAP features presented below, now letting the nuisance parameters vary. We find that this choice has almost no impact in our analysis, changing the $\Delta\chi^2$ values by less an unit.

Finally, we have assessed the impact of using separately each of the high- ℓ unbinned TT and EE Planck 2018 data sets, in order to check which subset dominates the posterior around each of the fits. To do that, we have combined each of these subsets with low- ℓ temperature and polarisation, and lensing data, fixing LCDM and nuisance parameters, and assuming a single dip. Using high- ℓ TT data alone, we recover the main single-dip maxima a posteriori described in the next section and appendix 2.6

2.4. RESULTS

(i.e. $\tau_1 \approx -100, -200, -400, -1000$). When using high- ℓ EE data alone, we clearly recover the dip at $\tau_1 \approx -100$, whereas just barely the peaks at $\tau_1 \approx -200$ and $\tau_1 \approx -400$, and none of the earlier-time dips found in combination with TT data. This shows that our posterior is temperature-dominated in most of the parameter range. Concordance at late times is encouraging, but it needs to be explored further with future less noisy polarised data from space- and ground-based surveys.

2.4.2 Reconstruction of the inflaton's speed of sound profile $u(\tau)$

In this section, we present the results of the GP reconstruction in the logarithmic parameterization using the Planck data as described above, and imposing the Maximum-Entropy prior described in 2.3.2 for the derived quantities $|u|_{\text{max}}$ and $|s|_{\text{max}}$.

When presenting our results, we use an effective $\Delta \chi^2$ where we have subtracted the χ^2 of the MAP of the featureless Λ CDM (obtained by using a minimizer) for the same likelihood combination (see subsection 2.3.3). Note that this effective $\Delta \chi^2$ is not meant for model selection purposes and it is used for illustration only. As an example, a triangle plot of the posterior distribution for the 2-nodes case can be seen in Figure 2.6.

We have reconstructed the inflaton's speed of sound profile $u(\tau)$ using up to four training nodes. We have stopped there after checking that the Akaike Information Criterion (AIC) (Akaike, 1974) has a minimum for three training nodes and stabilizes after that. The profile $u(\tau)$ shows different patterns depending on how many training nodes are used in the GP reconstruction. We have decided to classify all possible profiles $u(\tau)$ based on whether they show differentiated and non-overlapping reductions (that we denominate dips) or they present some kind of substructure:

- One single dip: usually present at either late values of conformal time (-100, -200, -400), or at earlier times (-800, \sim -1000). Early-time dips produce features in the CMB power spectra localised in ℓ 's up to the first acoustic peak, whereas features from late-time dips affect the power spectra along the full ℓ range. Similar profiles were already found in previous studies (see subsection 2.4.1). Details on this posterior modes can be found in appendix 2.6.1.
- Combination of non-overlapping reductions (2, 3 and 4 dips): appearing when more than one training node is used, they consist of consecutive, isolated reductions in the speed of sound.⁴ Details can be found in appendices 2.6.2, 2.6.3 and 2.6.4. These combinations can be classified as (for details see appendix 2.6):
 - All dips at late conformal times: when at least two training nodes are considered, there is a preference for two of the possible dips remaining at

⁴Notice that the number of training nodes is not always equal to the number of dips: reconstructions with m dips found with m GP nodes usually re-appear as posterior modes in the m+1 GP nodes case, where one of the nodes is placed at $u_i \approx 0$.

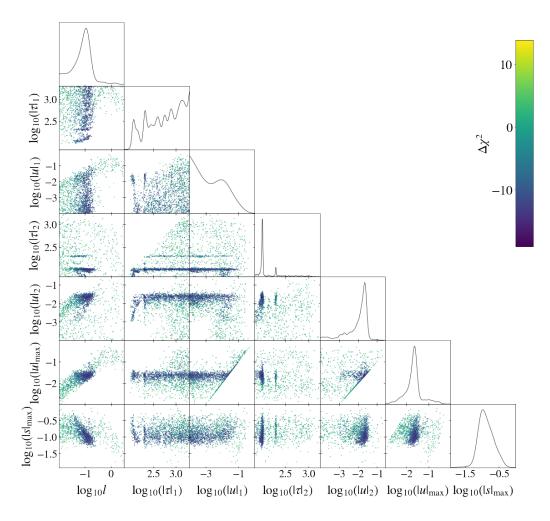


Figure 2.6: Posterior distribution for the reconstruction of the speed of sound's profile $u(\tau)$ using 2 training nodes and the logarithmic parametrization shown in Figure 2.2. We use $\Delta \chi^2 = \chi^2_{\rm model} - \chi^2_{\rm baseline}$ as the variable for the scatter plot's colour scale, the reference $\chi^2_{\rm baseline}$ corresponding to the MAP of the baseline Λ CDM model to the same datasets. We show the parameters of the training nodes (τ_i, u_i) and the correlation length l (described in 2.3.1; priors in 2.3.2). We also show the posteriors of the EFT parameters $(|u|_{\rm max}, |s|_{\rm max})$ (described in section 2.2, and not sampled directly, but derived from the nodes parameters). It can be seen how longer correlation lengths (broader reductions) lead to lower values of $|s|_{\rm max}$, and vice-versa. The posterior distributions for different numbers of nodes display similar patterns. For all cases, the posterior distributions are clearly multi-modal.

late-time values of τ_i , combining either -100 and -200, or -400 and -200. Their effect in the CMB power spectrum overlap each other along a large range of ℓ 's.

- Combination of early- and late-time dips: these appear typically as a combination of features at both low ℓ 's (from the early-time dips) and high ℓ 's

2.5. CONCLUSIONS

(from the late time ones), e.g. from the presence of dips both at -800 and -100.

• Dips with substructure: We have found some maxima a posteriori where the reconstructed $u(\tau)$ does not show clearly separated reductions, but a more complex profile with some degree of substructure. These substructures are presented either at early and late τ_i , trying to fit some of the characteristic features of the CMB angular power spectrum (i.e: $\ell \approx 20 - 40$ feature). The fits are presented in Appendix A, in subsection 2.6.5.

As noticed in previous works (Torrado et al., 2017), we do not have a highly predictive posterior of the maximum of the rate of change of the sound speed, $|s|_{\text{max}}$, whose value is mostly constrained by the prior information. By contrast, the positions of the nodes (the oscillation frequency of the features in the power spectrum) are tightly constrained within each of the multiple posterior modes, specially for nodes at late conformal time.

Using the sampling results of the profile $u(\tau)$ with four training nodes in the GPs, we have reconstructed the allowed confidence contours for $u(\tau)$ given Planck 2018 data. The result can be seen in Figure 5.3. As expected from the 1-D marginalized posterior distributions, the confidence contours are narrower around $\tau_i = [-100, 200]$. These modes were found in every single reconstruction of the inflaton's speed of sound independently of the number of training nodes (and were also observed in previous studies (Torrado et al., 2017)), and usually show the highest individual dip $\Delta \chi^2$ with respect to Λ CDM (since they produce features at a long range of ℓ for which Planck has low error bars). On the other hand, the confidence contours are broader for earlier conformal times $\tau_i < -400$. This is the range of τ where we have found the modes at $\tau_i = [-800, -1000]$ and some degree of substructure. In this range, the posterior distributions are not very predictive (see again Figure 2.6, where the posterior peaks are small for $\tau_i < -800$), since they produce low-multipole features hidden by cosmic variance.

2.5 Conclusions

We have searched for features in the primordial power spectrum as given by the last release of Planck 2018 data. Following an EFT of inflation approach, we have focused our search on features coming from reductions of the sound speed of the inflaton, assuming these reductions to be small, mild and transient. These feature templates were not tested by the Planck Collaboration.

We have improved over previous studies (which used a single-reduction Gaussian ansatz) by developing a reconstruction technique for the speed of sound's profile based on Gaussian Processes. We have also marginalised over the slow roll parameters to allow for a dynamical prior. In this new pipeline, the parameters of the reconstruction (the position of the training nodes and the correlation length) are fitted to the Planck 2018 data. The physical constraints of the model are imposed on the reconstruction parameters by means of a Maximum-Entropy prior defined on the EFT quantities

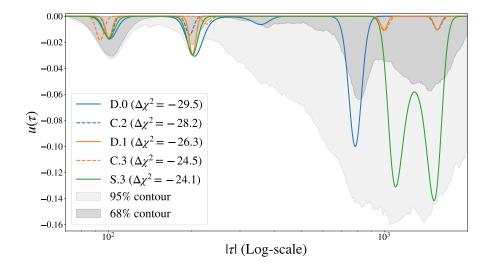


Figure 2.7: Reconstruction of the inflaton's speed of sound profile $u(\tau)$ based a 4-nodes GP, where the confidence contours (68% and 95%) are shown. We are able to constrain the shape the inflaton's speed of sound more stringently at late times (up to ~ -200), whereas the confidence intervals get larger at earlier times (i.e. starting from -800). This difference in the constraining power between early and late conformal times is mostly due to early-time reductions being associated to low-multipole features where cosmic variance is largest. The best maxima a posteriori are also plotted on top of the confidence contours: two 3-dipped cases (labelled C.2, dashed blue, and C.3, dashed orange), two 4-dipped cases (labelled D.0, solid blue, and D.1, solid orange) and a 3-dipped case in which one of the dips possesses some substructure (labelled S.3, solid green). For additional fits and a more detailed presentation of them see appendix 2.6.

 $(|u|_{\text{max}}, |s|_{\text{max}})$, which define the consistency bounds of the model. We have also tracked as derived parameters n_s and r.

This template-free reconstruction of $u(\tau)$ has allowed us to make an exhaustive search of more flexible features' templates, constrained only by EFT conditions. The analysis of the result of Bayesian parameter inference on the Planck 2018 data has demonstrated that there are many possible different and complex $u(\tau)$ profiles which are consistent with Planck's CMB power spectra. As expected, none of these fits is preferred with respect to Λ CDM (their $\Delta \chi^2$'s are not significant), although show some interesting results in terms of new feature templates. First, we have argued that there is a strong preference for two consecutive reductions of the speed of sound to coexist at late times around $\tau_i \approx -200$ and $\tau_i \approx -100$. Also, combinations of modes at late conformal time $\tau_i \approx -100$ and early conformal time $\tau_i \approx -800$ are also possible. Second, we have found certain profiles which show some degree of sub-structure at early and late conformal times. Finally, we have been able to obtain reconstruction

confidence contours for the $u(\tau)$ profile given the results obtained with four training nodes.

In the future, we plan to exploit this robust and novel pipeline in the search of features using new sets of data (in particular, Large Scale Structure surveys or the CMB bispectrum). Furthermore, the improvement of current data (for example, the polarization of the CMB) will also help to reduce the noise and, therefore, the uncertainty we have at large scales. If the noise is reduced, we could discern how realistic the reductions at earlier conformal times are. Moreover, we also consider introducing new features coming from a variable first slow-roll parameter (Durakovic et al., 2019) to perform a joint search of both patterns: features induced by a variable $c_s(\tau)$ and $\epsilon(\tau)$.

2.6 Appendix: detailed results up to 4 nodes of $u(\tau)$ in log-log parameterization

In this appendix, we explain in detail the several maxima a posteriori found during the sampling runs when the profile of the inflaton's speed of sound $u(\tau)$ was reconstructed using Gaussian Processes up to 4 training nodes.

2.6.1 One dip (denoted by "A")

These profiles of $u(\tau)$ show only one single reduction of the inflaton's speed of sound or one dip. These reductions can be found using just one training node in the Gaussian Process. The modes at late conformal time τ_i (see the first row of Figure 2.8) present a well defined oscillation frequency τ_i (at -100, -200 and -400). The value of $|u|_{\text{max}}$ is around 0.02 and the rate of change in the speed of sound $|s|_{\text{max}} \gg |u|_{\text{max}}$. These dips are exactly reproduced with the linear parametrization of the reconstruction of $u(\tau)$. They were already listed during previous searches using Planck 2013 and 2015 data (Hu & Torrado, 2015; Torrado et al., 2017). In particular, the mode corresponding at $\tau_i \approx -400$ was identified faintly in (Torrado et al., 2017). These modes are present in a broad multipole ℓ range ($\ell \approx [100-2000]$), fitting some structures in the temperature and polarization data.

Modes at early conformal time τ_i (-800, -1000) (see the second row of Figure 2.8) are found, but are more poorly constrained and with worse $\Delta \chi^2$ with respect to Λ CDM. Mode A.3 at $\tau_i \approx -800$ shows similar characteristics to the modes at a late conformal time (small amplitude and same behaviour in the EFT parameter s). It fits an apparent oscillating structure of the temperature C_ℓ at the first acoustic peak. This mode is also found with the linear parametrization and in previous searches in (Hu & Torrado, 2015; Torrado et al., 2017). The modes at -800 and -1000, with a larger amplitude, have $|s|_{\text{max}} \approx |u|_{\text{max}}$. They slightly differ from the modes found in previous studies. The main reason is that $u(\tau)$ reconstructed using the logarithmic parametrization differs from the linear one at high values of τ_i . The modes at -1000 try to fit the characteristic $\ell \approx 20 - 40$ structure of the CMB temperature angular power spectrum. As identified in (Torrado et al., 2017), this kind of features impose

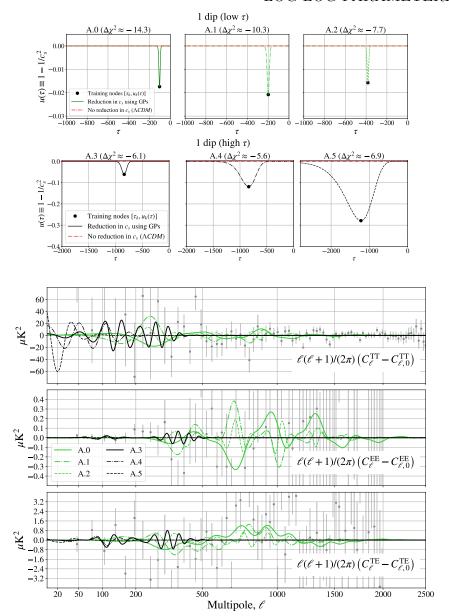


Figure 2.8: One single dip, one training node. **Top:** Different profiles $u(\tau)$ for the 6 maxima a posteriori when only 1 training node is used (and consequently only one dip is visible). The reconstruction is done following the logarithmic parametrization explained in subsection 2.3.1. We found a principal MAP and 5 other fits when the multimodal posterior distribution is further analysed (see, for example, Figure 2.6, where other peaks in the posterior distribution are visible). **Bottom:** Differences in the CMB temperature (TT), E-polarization (EE) and cross-correlated power spectra (TE) between the MAP to the Planck 2018 data and the featureless Λ CDM baseline model for the reconstructed speed of sound profiles $u(\tau)$ A.0 - A.5 shown above. Notice how these profiles fit small deviations from Λ CDM at low and high multipoles ℓ . The same color and line-style correspondence between the $u(\tau)$ profiles and the differences in the CMB spectra has been used.

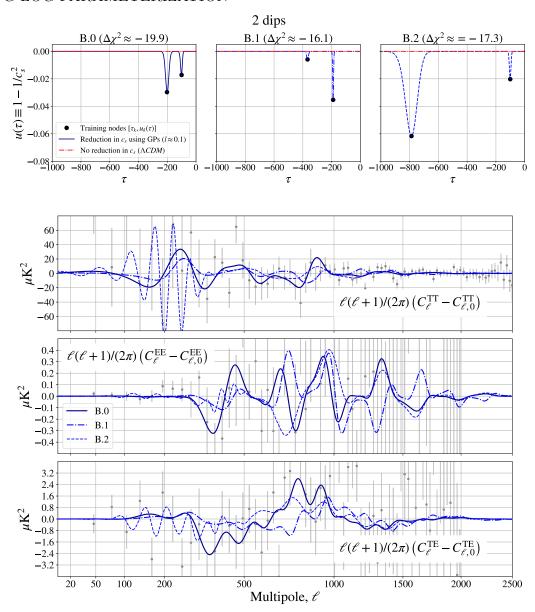


Figure 2.9: 2 non-overlapping dips, 2 training nodes. **Top:** Different profiles $u(\tau)$ for the 3 maxima a posteriori when only 2 training nodes are used and only two clearly different dips are observed. The reconstruction is done following the logarithmic parametrization explained in subsection 2.3.1. We found a principal best fit and 2 other fits when the multimodal posterior distribution is further studied (see, for example, Figure 2.6, where other peaks in the posterior distribution are visible). **Bottom:** Differences in the CMB temperature (TT), E-polarization (EE) and cross-correlated power spectra (TE) between the best fit to the Planck 2018 data and the featureless Λ CDM baseline model for the reconstructed speed's of sound profile $u(\tau)$ shown above. Notice how these profiles fit small deviations from Λ CDM at low and high multipoles ℓ . The same color and line-style correspondence between the $u(\tau)$ profiles and the differences in the CMB spectra has been used.

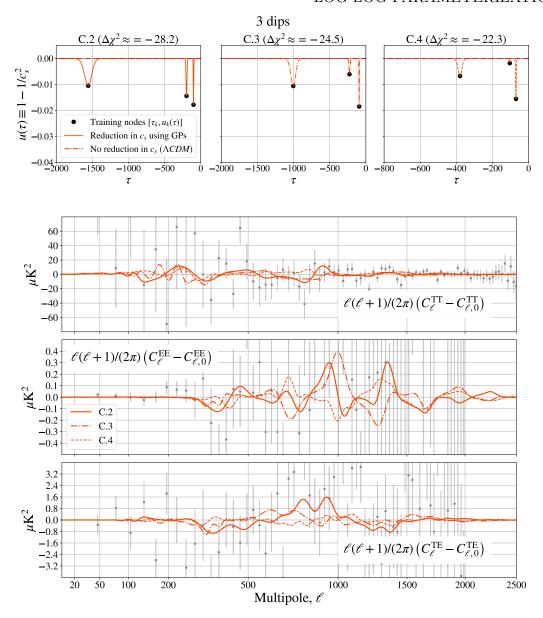


Figure 2.10: 3 non-overlapping dips, 3 training nodes. **Top:** Different profiles $u(\tau)$ for the 3 maxima a posteriori when only 3 training nodes are used and three differentiated dips are observed. The reconstruction is done following the logarithmic parametrization explained in subsection 2.3.1. We found 3 MAP when the corresponding multimodal posterior distribution is further studied (see, for example, Figure 2.6, where other peaks in the posterior distribution are visible). **Bottom**: Differences in the CMB temperature (TT), E-polarization (EE) and cross-correlated power spectra (TE) between the best fit to the Planck 2018 data and the featureless Λ CDM baseline model for the reconstructed speed's of sound profile $u(\tau)$ shown above. Notice how these profiles fit small deviations from Λ CDM at low and high multipoles ℓ . The same color and line-style correspondence between the $u(\tau)$ profiles and the differences in the CMB spectra has been used.

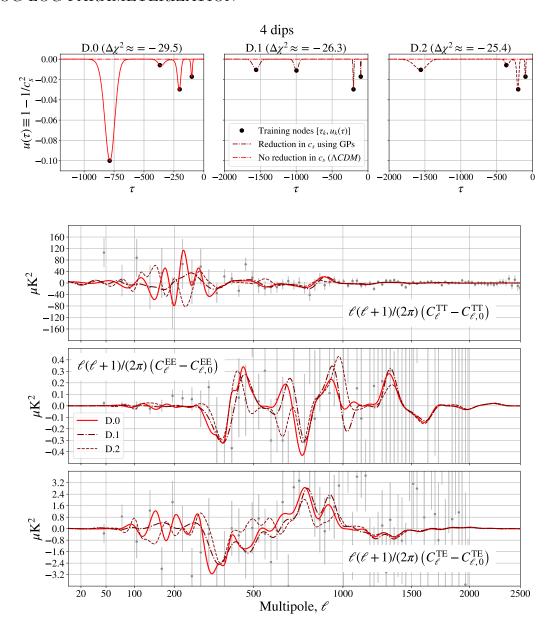


Figure 2.11: 4 non-overlapping dips, 4 training nodes. **Top:** Different profiles $u(\tau)$ for the 3 maxima a posteriori when only 4 training nodes are used and four differentiated dips are observed. The reconstruction is done following the logarithmic parametrization explained in subsection 2.3.1. In this case, we observe how possible reductions at $\tau \approx -100, -200, -400$ and -800 can consecutively take place. **Bottom:** Differences in the CMB temperature (TT), E-polarization (EE) and cross-correlated power spectra (TE) between the best fit to the Planck 2018 data and the featureless Λ CDM baseline model for the reconstructed speed's of sound profile $u(\tau)$ shown above. Notice how these profiles fit small deviations from Λ CDM at low and high multipoles ℓ . The same color and line-style correspondence between the $u(\tau)$ profiles and the differences in the CMB spectra has been used.

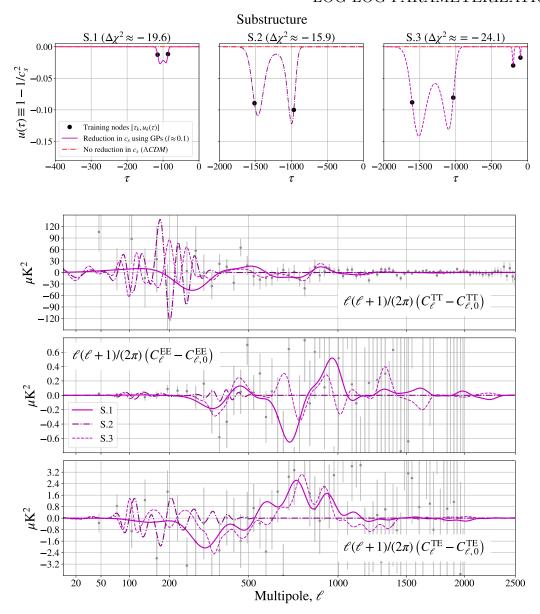


Figure 2.12: Profiles with substructure (2 training nodes and 4 training nodes). **Top:** Different profiles $u(\tau)$ for 2 maxima a posteriori when only 2 or training nodes are used and the profiles of $u(\tau)$ show some grade of substructure. The reconstruction is done following the logarithmic parametrization explained in subsection 2.3.1. We found 2 fits (one an earlier conformal time and a another one at late conformal time), when the corresponding multimodal posterior distributions are further studied (see, for example, Figure 2.6, where other peaks in the posterior distribution are visible). **Bottom:** Differences in the CMB temperature (TT), E-polarization (EE) and cross-correlated power spectra (TE) between the best fit to the Planck 2018 data and the featureless Λ CDM baseline model for the reconstructed speed's of sound profile $u(\tau)$ shown above. Notice how these profiles fit small deviations from Λ CDM at low and high multipoles ℓ . The same color and line-style correspondence between the $u(\tau)$ profiles and the differences in the CMB spectra has been used.

a tighter upper limit on the scalar-to-tensor ratio r, although these are still within the analogous bounds in Λ CDM.

2.6.2 Two dips (denoted by "B")

In this case, the inflaton would suffer two consecutive reductions of the speed of sound (due to, for instance, two consecutive turns in the field space). In the previous study using Planck 2015 (Torrado et al., 2017), it was pointed out that, a priori, the features due to single reductions of the sound speed that do not overlap can in principle co-exist. These combinations would be modes at early τ_i with another late mode (i.e: -1000 and -100). However, the results of the reconstruction using two training nodes show a richer picture (see Figure 2.9). We have identified that the modes at -100 and -200 can result from an overlapping feature that is preferred by the data, and thus, it is the overall MAP for the 2-nodes reconstruction. The dips at -200 and -400 can also co-exist (mode B.1), with a worse $\Delta \chi^2$. These two overlapping features fit TT, TE and EE structure across a large range of ℓ . On the other hand, there is a possible combination of the modes at -800 and -100 (see the first row of Figure 2.9). This $u(\tau)$ profile includes the fitting of the apparent oscillations around the first acoustic peak and small deviations across the rest of the multipole scale. All of these combinations of modes fulfill $|s|_{\text{max}} \gg |u|_{\text{max}}$.

2.6.3 Three dips (denoted by "C")

When the profile of the inflaton's reduction of the speed of sound $u(\tau)$ is reconstructed using 3 training nodes, we find more complex profiles. The usual dips at late τ_i around -100 and -200 combine to mild and small reductions at early conformal times around -1000 and -1500 (see modes C.2 and C.3 of subsection 2.6.3, respectively). The combinations that are preferred by the data are those whose earlier training node τ_i is placed around $u_i \approx -0.01$. These small dips at early τ_i were not found alone during the search using one training node (modes A shown in Figure 2.8). However, these modes at early conformal times are loosely constrained and the confidence intervals are large. Overall, these profiles fits the data very similarly as the single standing modes A.0 and A.1, as well as the combination B.0. Furthermore, we have also verified that the modes A.0 and A.2 can co-exist (contrary to the case of reconstruction with 2 training nodes, where the combination -100 and -400 was not found) if a very small mode is added close to -100. All these profiles presented in Figure 2.10 are in the limit $|s|_{\text{max}} \gg |u|_{\text{max}}$. Finally, the 2-dip profiles explained above are also reproduced when we run with 3 training nodes.

2.6.4 Four dips (denoted by "D")

We find similar profiles for $u(\tau)$ as the ones for 3 training nodes, adding one extra dip and finding the remaining possible combination of nodes at late τ_i around -100, -200 and -400 with modes at earlier times at -800, -1000 or -1500 (see modes D.0, D.1 and D.3 of Figure 2.11). Thus, we have corroborated that the modes A.0, A.1, A.2

and A.4 can, in principle, co-exist. Still, the small modes at early $\tau_i \approx -1000, -1500$ are less likely to show up in the posterior, as they are poorly constrained given the data. It is worth mentioning that we can mostly reproduce all the different profiles found during the search using up to three nodes when four training nodes are used. In this case, one, two or three training nodes are placed in such a way the corresponding profile looks very similar to the cases A, B or C (either the training node is placed close to $u_i \approx 0$ or close to the previous training mode itself).

2.6.5 Reductions with substructure (denoted by "S")

Apart from concatenations of transient reductions in the speed of sound, we have also observed some possible fits which show more complicated feature patterns according to the data. These are profiles of $u(\tau)$ that do not clearly show full dips but have some kind of substructure (see the upper row of Figure 2.12, profile S.1). We find a sub-structured maxima a posteriori at a late conformal time (centred around $\tau_i \approx -100$), which resembles the mode A.0 but with two small sub-reductions. Similarly to A.0, the limit of the EFT functions is $|s|_{\text{max}} \gg |u|_{\text{max}}$.

Motivated by the loose constraints of the training nodes in the range of early conformal times $-800 < \tau_i < -3000$, we have launched a GPs reconstruction using two training nodes, which are restricted to remain in the range $-3500 < \tau_i < -990$, obtaining the profile S.2 of Figure 2.12. This profile have the particularity that $|s|_{\text{max}} \ll |u|_{\text{max}}$. This profile tries to fit not only the structure of the CMB TT angular power spectrum around $\ell \approx 20-40$ but also the apparent structure of the first acoustic peak in the TT and TE data. To answer the question if it is possible the substructure mode S.2 to co-exist with any of the modes at late conformal time τ_i , we have relaunched the GPs reconstruction with 4 training nodes, constraining the two early ones in the range $-3500 < \tau_1 < \tau_2 < -990$. When this constraint is imposed, the later training nodes τ_3 and τ_4 are placed clearly around -100 and -200. The resulting profile (see mode S.3 in Figure 2.12) can fit the CMB data at low ℓ but also in a broader range similarly to the case of the mode A.0, increasing the statistical significance $\Delta \chi^2$. In this case, the EFT limit is reverted to $|s|_{\text{max}} \gg |u|_{\text{max}}$ due to the narrow mode at -100.