



HAL
open science

Compressed 'CMB-lite' Likelihoods Using Automatic Differentiation

L Balkenhol

► **To cite this version:**

L Balkenhol. Compressed 'CMB-lite' Likelihoods Using Automatic Differentiation. 2024. hal-04847398

HAL Id: hal-04847398

<https://hal.science/hal-04847398v1>

Preprint submitted on 19 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

COMPRESSED 'CMB-LITE' LIKELIHOODS USING AUTOMATIC DIFFERENTIATION

L. BALKENHOL  ^{1,*}

¹Sorbonne Université, CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98 bis bd Arago, 75014 Paris, France

Version December 3, 2024

ABSTRACT

The compression of multi-frequency cosmic microwave background (CMB) power spectrum measurements into a series of foreground-marginalised CMB-only band powers allows for the construction of faster and more easily interpretable 'lite' likelihoods. However, obtaining the compressed data vector is computationally expensive and yields a covariance matrix with sampling noise. In this work, we present an implementation of the CMB-lite framework relying on automatic differentiation. The technique presented reduces the computational cost of the lite likelihood construction to one minimisation and one Hessian evaluation, which run on a personal computer in about a minute. We demonstrate the efficiency and accuracy of this procedure by applying it to the differentiable SPT-3G 2018 *TT/TE/EE* likelihood from the `candl` library. We find good agreement between the marginalised posteriors of cosmological parameters yielded by the resulting lite likelihood and the reference multi-frequency version for all cosmological models tested; the best-fit values shift by $< 0.1\sigma$, where σ is the width of the multi-frequency posterior, and the inferred parameter error bars match to within $< 10\%$. We publicly release the SPT-3G 2018 *TT/TE/EE* lite likelihood and a python notebook showing its construction on the `candl` website.

1. INTRODUCTION

Observations of the cosmic microwave background (CMB) anisotropies form a cornerstone of modern cosmology. Data from recent and contemporary experiments, such as the *Planck* satellite (Planck Collaboration et al. 2020a,c,b), the South Pole Telescope (SPT) (Carlstrom et al. 2011; Dutcher et al. 2021; Balkenhol et al. 2023; Pan et al. 2023; Ge et al. 2024), the Atacama Cosmology Telescope (ACT) (Kosowsky 2003; Aiola et al. 2020; Choi et al. 2020; Madhavacheril et al. 2024; Qu et al. 2024), and the BICEP/Keck experiments (Keating et al. 2003; Staniszewski et al. 2012; Ade et al. 2021) allow us to test the Λ (cosmological constant) (CDM) cold-dark-matter model across a wide range of angular scales in temperature and polarisation and search for new physics.

These experiments typically supply us with measurements of the CMB anisotropy power spectra at different observational frequencies binned into band powers; from N frequency channels $\geq N(N-1)/2$ multi-frequency spectra are constructed. At the likelihood level, one then fits for the common CMB signal based on a cosmological model and the varying foreground contamination in the multi-frequency band powers along with any systematic effects. The ensemble of foreground and systematic parameters is referred to as nuisance parameters. While, as their name suggests, the nuisance parameters are not of primary interest, they must be included in a likelihood analysis to obtain the correct results on cosmological parameters. However, nuisance parameters can be more numerous and hence slow down a Markov Chain Monte Carlo (MCMC) analysis by increasing the dimensionality of the parameter space to be explored and by slowing down the evaluation time of the likelihood due to the additional components in the data model.

While different methods have been proposed to compress the multi-frequency data into a shorter vector containing only the information of interest, the most common strategy for CMB power spectrum analyses is the construction of a so-called *lite* likelihood, which contains a set of foreground-marginalised CMB-only band powers. This allows for faster cosmological analyses primarily through the reduction of the number of nuisance parameters. The CMB-lite framework was first introduced by Dunkley et al. (2013) (hereafter D13) and has been applied to ACT (D13; Choi et al. (2020)), SPT (Calabrese et al. 2013), Planck (Planck Collaboration et al. 2016), and most recently BICEP/Keck data (Prince et al. 2024). Though lite likelihoods are advantageous once at hand, their construction is computationally expensive. As set out by D13, it typically involves sampling a high-dimensional parameter space ($\mathcal{O}(10^2)$) with an alternating Gibbs-sampling technique. This procedure naturally produces a covariance matrix with sampling noise, which must be reduced sufficiently or mitigated in another way to avoid adverse effects in subsequent cosmological analyses (Dodelson & Schneider 2013; Hartlap et al. 2007; Percival et al. 2022; Balkenhol & Reichardt 2022).

Recently, the application of automatic differentiation in cosmology has seen a surge in popularity (see Campaigne et al. (2023) for a series of example applications). In this work, we show that by using the JAX-friendly (Bradbury et al. 2018), differentiable CMB likelihood library `candl` (Balkenhol et al. 2023) we can construct lite likelihoods in circa one minute on a personal computer and eliminate the sampling noise in the covariance matrix. We apply this procedure to the SPT-3G 2018 *TT/TE/EE* data set (Balkenhol et al. 2023) and benchmark the performance of the lite likelihood against its multi-frequency counterpart. The lite likelihood and code to construct CMB-only band powers is made pub-

*lennart.balkenhol@iap.fr

licly available on the `candl` website.¹

We begin with background on the CMB-lite framework and the `candl` library in Section 2. In Section 3 we outline our technique for the construction of CMB-lite likelihoods and subsequently apply it to the SPT-3G 2018 $TT/TE/EE$ data set in Section 4. We share our conclusions in Section 5.

2. BACKGROUND

We briefly review the CMB-lite framework as introduced by D13. We formulate model band powers $C_b^{\mu\nu,\text{model}}(\phi)$ to describe the data, where the ϕ parameters contain CMB band powers C_b^{CMB} , as well as nuisance parameters θ , $\phi = (C_b^{\text{CMB}}, \theta)$. To be explicit, in this work we only consider band powers, i.e. binned power spectra, which we index with subscripts b and b' , and summation over repeated indices is implied. The nuisance parameters enter our frequency-dependent foreground model $C_b^{\mu\nu,\text{FG}}(\theta)$ and calibration factors $y^{\mu\nu}(\theta)$, such that

$$C_b^{\mu\nu,\text{model}}(\phi) = y^{\mu\nu}(\theta) \left(A^{\mu\nu} C_b^{\text{CMB}} + C_b^{\mu\nu,\text{FG}}(\theta) \right). \quad (1)$$

Here, $A^{\mu\nu}$ is sometimes referred to as the mapping-matrix (D13) or the design-matrix (Planck Collaboration et al. 2020c; Mocanu et al. 2019; Dutcher et al. 2021; Balkenhol et al. 2023); this matrix of zeros and ones connects the multi-frequency estimates of the same CMB band power in the data vector to the corresponding single element in C_b^{CMB} . More generally, we can think of this model as having multipole- and frequency-dependent additive components to the CMB signal and frequency-dependent multiplicative factors.

The CMB-lite framework seeks to determine the values of the CMB-only band powers and their covariance, marginalising over the nuisance parameters.² In order to do so one samples over ϕ , comparing the model prediction to the measured data $C_b^{\mu\nu,\text{data}}$, typically using a Gaussian likelihood

$$2 \ln \mathcal{L}^{\text{recon}}(C^{\mu\nu,\text{data}}|\phi) \propto \left(C_b^{\mu\nu,\text{model}}(\phi) - C_b^{\mu\nu,\text{data}} \right)^T \Sigma_{bb'}^{-1} \left(C_{b'}^{\mu\nu,\text{model}}(\phi) - C_{b'}^{\mu\nu,\text{data}} \right), \quad (2)$$

where Σ is the multi-frequency band power covariance matrix. We refer to this likelihood as the *reconstruction likelihood* in this work and its exploration as the *reconstruction process*. In D13, the reconstruction likelihood is analysed with an MCMC approach by alternating Gibbs sampling steps for the CMB band powers and the nuisance parameters. The resulting posterior distribution yields the values of the CMB-only band powers and their covariance, which are then used to build the lite likelihood used to explore cosmological models.

¹ <https://github.com/Lbalkenhol/candl>

² Note that in the CMB-lite framework the term *secondary parameters* is occasionally used to describe parameters that are marginalised over during the construction of the CMB-only band powers. So while the label nuisance parameters refers to everything non-cosmological, secondary parameters are only the subset of nuisance parameters not wanted in the final lite likelihood. We assume these are equivalent for the sake of simplicity here, though as we will see in Section 4 it can be beneficial to retain some select nuisance parameters in the lite likelihood.

When analysing the same cosmological model, good agreement between the marginalised posterior distributions obtained from the multi-frequency and lite likelihoods is observed. Still, small differences do arise, as during the reconstruction process one picks up fluctuations from noise and foreground contamination and assigns these to the CMB-only band powers. In the lite likelihood these may project differently onto cosmological parameters compared to the multi-frequency likelihood, where cosmological and nuisance parameters are explored simultaneously. Typically though, this realisation-dependent bias is below or comparable to the level of numerical noise expected in parameter estimation, which we take to be shifts in central values of marginalised posteriors by 10% of their width.

There are numerous advantages to the CMB-lite framework. First, it reduces the size of the data vector, speeding up likelihood evaluation. It also eases the exploration of the likelihood surface due to the reduction of the number of parameters to be sampled. Moreover, it combines the common information from the multi-frequency band powers in an optimal way, forming a minimum-variance combination of the data when the CMB signal is dominant. At the same time, uninteresting information is marginalised over and the resulting uncertainty is neatly captured in the covariance of the CMB-only band powers. This makes lite likelihoods also more intuitively interpretable. Crucially, no cosmological model enters the reconstruction likelihood (see equations 1 and 2). This means the reconstruction procedure does not introduce additional model-dependence beyond what already exists in the reference multi-frequency likelihood.

However, the framework is not without downsides. If cosmological parameters are strongly correlated with foreground parameters, constraints from the multi-frequency and lite likelihoods may differ non-negligibly. On a technical level, the MCMC analysis of the reconstruction likelihood is computationally expensive and leads to sampling noise in the covariance matrix estimate, which can cause adverse effects in the parameter constraints of the lite likelihood if not reduced sufficiently or otherwise mitigated (Dodelson & Schneider 2013; Hartlap et al. 2007; Percival et al. 2022; Balkenhol & Reichardt 2022). Of course, one is always free to run the reference multi-frequency likelihood when in doubt, though having to do so continuously diminishes the value of the CMB-lite framework.

In this work, we ameliorate the numerical challenges associated with the reconstruction procedure. To do so, we use the `candl` library of CMB likelihoods, which allows for easy, intuitive access to data from the SPT and ACT collaborations and, crucially, is written in a JAX-friendly way (Bradbury et al. 2018). JAX is a google-developed python package which exposes the code to an automatic differentiation algorithm that allows for the accurate and fast calculation of derivatives of `candl` likelihoods without relying on finite difference methods.³

3. TECHNIQUE

Thanks to `candl` and JAX the reconstruction likelihood, Equation 2, is differentiable, meaning we can eval-

³ <https://github.com/google/jax>

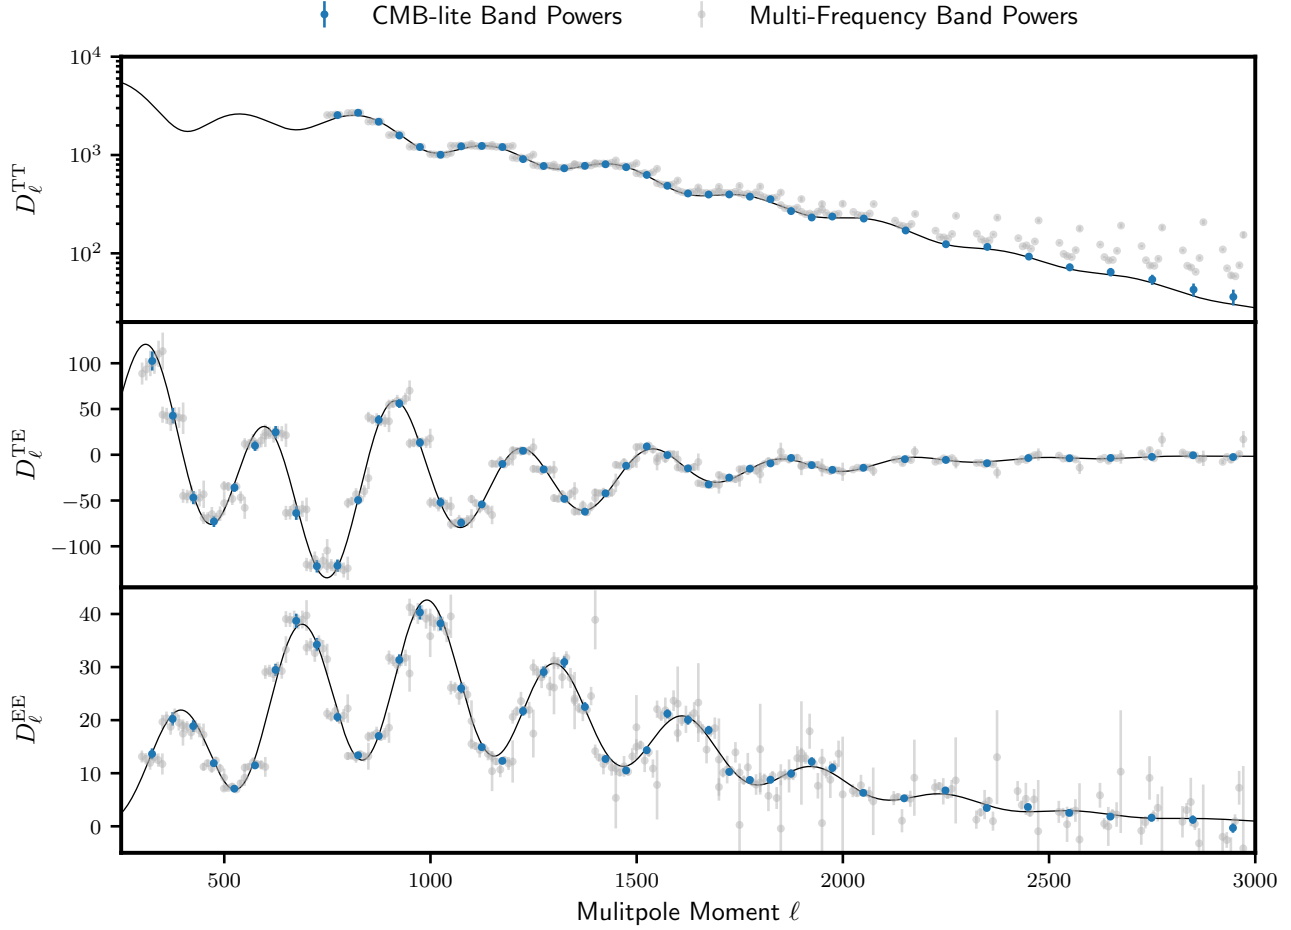


FIG. 1.— Foreground-marginalised CMB-only band powers (blue) based on the SPT-3G 2018 $TT/TE/EE$ multi-frequency data (grey). Though the construction of the CMB-only data vector does not require a cosmological model, the data visually follow the best-fit Λ CDM model of the multi-frequency likelihood shown in black. The CMB-lite band powers visibly remove the foreground contamination in the temperature power spectrum.

uate its first- and second-order derivatives

$$\left. \frac{\partial \mathcal{L}^{\text{recon}}}{\partial \phi_i} \right|_{\phi}, \quad \left(\frac{\partial^2 \mathcal{L}^{\text{recon}}}{\partial \phi_i \partial \phi_j} \right)_{\phi} \quad (3)$$

at any desired point ϕ quickly and accurately. Access to the gradient allows us to find the best-fit parameters of the reconstruction likelihood extremely fast using, e.g. the truncated Newton or Newton-Raphson methods (Nocedal & Wright 2006; Nash 1984). The Hessian is useful because evaluated at the best-fit point it is equal to the negative Fisher matrix and hence connects to the covariance of the CMB-only band powers (Heavens et al. 2014).

With access to these functions, we perform the reconstruction as follows:

1. Minimise the reconstruction likelihood, finding its best-fit point. This yields the CMB-only band powers.
2. Evaluate the Hessian of the reconstruction likelihood at the best-fit point. The negative inverse of this matrix serves as the covariance of the CMB-only band powers.

Hence, we obtain all necessary products for the lite likelihood from minimising the reconstruction likelihood once

and one evaluation of its Hessian. The above prescription greatly reduces the computational cost of the reconstruction procedure such that it can be performed on a personal laptop in about a minute. This means one can easily perform multiple reconstructions to assess robustness, using for example different foreground models or angular scale cuts. Moreover, this procedure eliminates numerical noise from the covariance matrix of CMB-only band powers. Though automatic-differentiation also has its limits, second-order derivatives are usually well behaved and we have not seen any suggestive artefacts in the application to follow in Section 4 (Campagne et al. 2023).

We note that while in principle the use of the Hessian of the reconstruction likelihood for the covariance of the CMB-only band powers does not capture non-Gaussian contributions, we did not detect significant biases due to this during the application to SPT data in Section 4. One can test for this effect by reconstructing a known scatter-free model spectrum and comparing parameter constraints for models of interest between the resulting lite and multi-frequency likelihoods. Should differences turn out to be sizeable, one may consider the methods put forward by Sellentin et al. (2014) or Arutjunjan et al. (2022) in order to improve the covariance estimate or re-

sort to MCMC sampling the reconstruction likelihood. For this the use of gradient-based sampling methods (e.g. Hamiltonian Monte Carlo, No U-Turns) is also promising as these tend to perform well in high-dimensionality scenarios (Hoffman & Gelman 2011).

4. APPLICATION TO DATA

In this section we apply the above procedure to the SPT-3G 2018 $TT/TE/EE$ likelihood implemented in `candl`. The minimisation of the reconstruction likelihood is carried out using a truncated Newton algorithm implemented in `scipy` (Virtanen et al. 2020; Nocedal & Wright 2006; Nash 1984). To benchmark the performance of the resulting lite likelihood and judge the success of the reconstruction procedure more generally we will consider three cosmological models: Λ CDM, Λ CDM+ N_{eff} , and Λ CDM+ A_L . We parametrise these models using the Hubble constant H_0 , the baryon and cold-dark-matter densities $\Omega_b h^2$ and $\Omega_c h^2$, respectively, the amplitude $\ln(10^{10} A_s)$ and spectral tilt n_s of the power spectrum of initial scalar fluctuations, the optical depth to reionisation τ , plus optionally either the effective number of relativistic degrees of freedom N_{eff} or the amplitude of the gravitational lensing effect on the power spectrum A_L . To calculate CMB spectra, we utilise the `CosmoPower` models trained for the analysis of the SPT-3G 2018 $TT/TE/EE$ data set on high-accuracy `CAMB` spectra (Balkenhol et al. 2023; Spurio Mancini et al. 2022; Piras & Spurio Mancini 2023; Lewis et al. 2000). To perform MCMC analyses of the multi-frequency and lite likelihoods in the aforementioned cosmological models we use `Cobaya` (Torrado & Lewis 2021) with a Metropolis-Hastings sampler. We use dedicated minimisers to obtain best-fit points, again using `scipy`'s truncated Newton algorithm. For cosmological analyses, we apply the same *Planck*-based Gaussian prior on τ centred on 0.054 with width 0.0074 as Balkenhol et al. (2023) in the lite likelihood (Planck Collaboration et al. 2020b).

The SPT-3G 2018 $TT/TE/EE$ multi-frequency likelihood has various noteworthy aspects for the reconstruction process. First, the window functions differ across frequencies owing to the instrumental beam and the filter strategy of the analysis, which as Prince et al. (2024) points out means that there exists no single set of uniquely defined CMB-only band powers to reconstruct; in principle, one would have to assign a CMB-only band power bin to each window function shape. However, since differences are small, we ignore them in the reconstruction procedure and use a combination of the multi-frequency window functions weighted by the band power covariance matrix in the lite likelihood.

Second, the multi-frequency likelihood involves the addition of a model- and frequency-dependent beam contribution to the band power covariance matrix. To avoid having to represent this operation at the level of the lite likelihood, one could reformulate the multi-frequency likelihood to account for the uncertainty of the beam measurement through modifications of the model spectra controlled by additional nuisance parameters to be marginalised over during the reconstruction process (see e.g. Henning et al. 2018). However, we find that accounting for this term in the reconstruction likelihood and simply ignoring it in the lite likelihood is sufficiently accurate for the data set at hand.

Third, we follow the suggestion of D13 and split the calibration module of the multi-frequency likelihood into two: an internal calibration, relative to the 150 GHz TT and EE auto-spectra, and an external, absolute calibration of all spectra controlled by new parameters T_{cal} and E_{cal} , which modify spectra like their frequency-dependent counterparts in the original likelihood. By keeping T_{cal} and E_{cal} in the lite likelihood we avoid introducing significant long-range correlations between the estimated CMB-only band powers, which ultimately makes it easier to interpret residuals of the data vector.

Fourth, the transformation of model spectra by the aberration effect due to the Earth's relative motion to the CMB depends on the derivative of the unbinned CMB power spectrum with respect to the multipole moment ℓ (Jeong et al. 2014). However, we are only interested in reconstructing CMB band powers, i.e. bins in multipole moments, and there exists no unique correspondence between the aberration effect at the band power and the unbinned CMB power spectrum level. Therefore, we consider aberration as part of the signal and aim to encapsulate it in our reconstruction, also accounting for this effect in the lite likelihood.

Lastly, we face a similar issue for the super-sample lensing effect (Manzotti et al. 2014). Though we could treat this phenomenon similarly to aberration, we would have to carry the parameter controlling the strength of the effect, the mean lensing convergence across the survey field, into the lite likelihood. Instead, we add a suitable contribution to the band power covariance matrix and therefore keep the number of nuisance parameters and

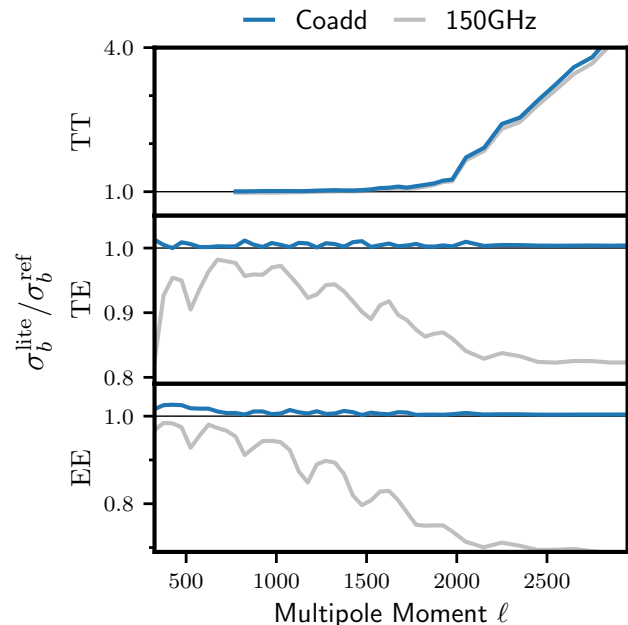


FIG. 2.— Ratio of the error bars of the lite likelihood to two reference cases (σ_b^{ref}): the coadd (blue) and the 150 GHz auto-spectrum (grey). As foregrounds are already marginalised over in the CMB-lite band powers, error bars are larger than the two reference cases at high ℓ for TT . For TE and EE we see that the lite likelihood improves on the 150 GHz auto-spectrum on all scales due to the combination of multi-frequency information. The beam and calibration uncertainties as well as the super-sample lensing effect cause an increase of $< 3\%$ of the lite error bars with respect to the coadd for TE and EE .

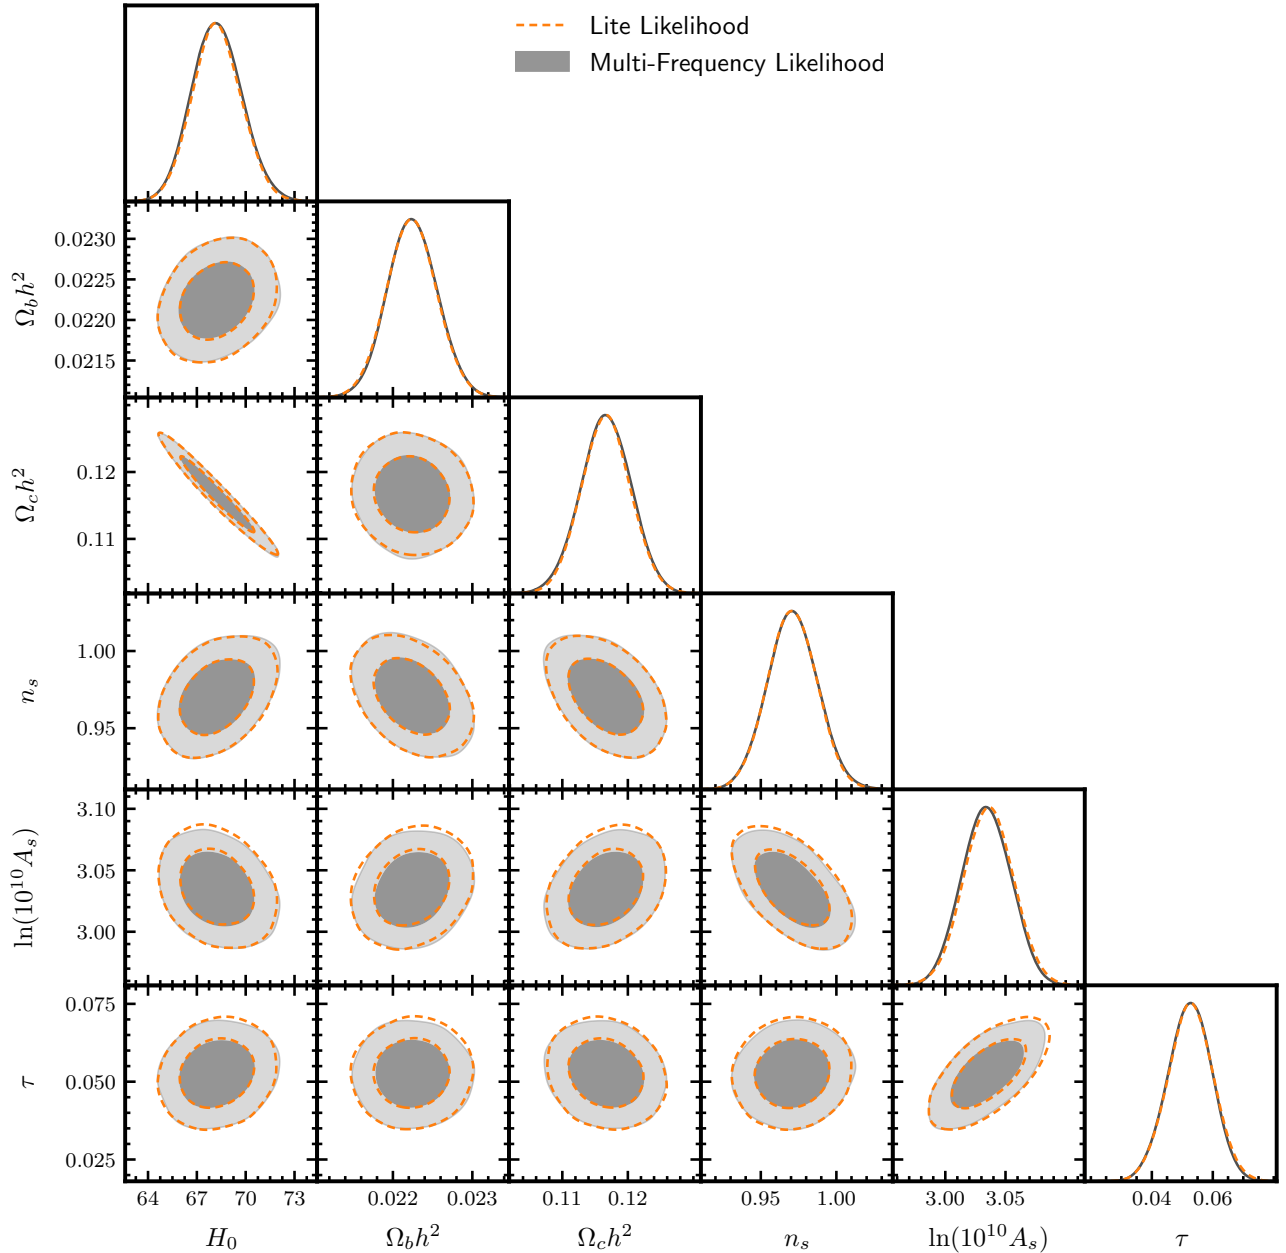


FIG. 3.— Marginalised posteriors for Λ CDM parameters obtained from the full multi-frequency (grey filled contours) and lite (orange dashed line contours) SPT-3G 2018 $TT/TE/EE$ likelihoods (68% and 95% confidence levels). The constraints match well, with only a small offset in $\ln(10^{10} A_s)$ of $< 10\%$ of the width of the multi-frequency posterior visually discernable.

transformations in the lite likelihood low.

Together, these points lead to a small systemic bias, which we quantify by generating scatter-free mock band powers for a set of fiducial parameters in Λ CDM. We perform the reconstruction and compare constraints on Λ CDM parameters between the multi-frequency and the lite likelihood. The best-fit values of the lite likelihood are within $< 0.01\sigma$ of the multi-frequency likelihood for all parameters, where σ refers to the width of the marginalised posteriors yielded by the multi-frequency likelihood. The error bars on cosmological parameters match to better than 1%. Hence any inaccuracies or biases in the reconstruction procedure are negligible compared to the level of uncertainty ascribed to MCMC anal-

yses and, as we demonstrate later on, to the realisation-dependent bias picked up by the reconstruction from noise and sample variance fluctuations when performed on real data.

We now perform the reconstruction procedure on the SPT-3G 2018 $TT/TE/EE$ data set. The resulting CMB-only band powers are shown in Figure 1. Like their multi-frequency counterpart, the band powers cover the multipole ranges $750 \leq \ell < 3000$ in TT and $300 \leq \ell < 3000$ in TE and EE with bin widths $\Delta\ell = 50$ for $\ell < 2000$ and $\Delta\ell = 100$ for $\ell > 2000$. Visually, we see that the power on small angular scales in temperature is reduced compared to the multi-frequency band powers as the foreground contamination is removed.

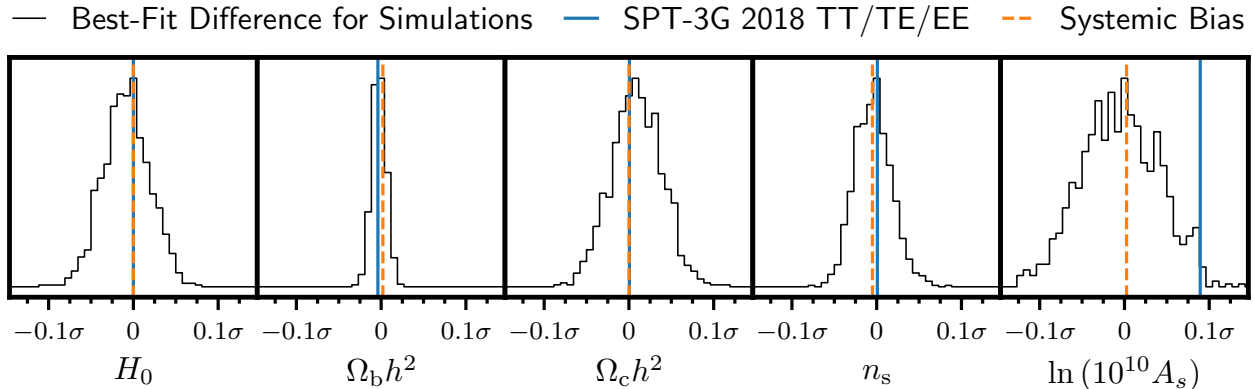


FIG. 4.— The difference in best-fit Λ CDM parameters of the lite likelihood compared to the reference multi-frequency likelihood for 1000 mock band power realisations (black histogram) normalised by the width of the constraints of the multi-frequency likelihood (1σ). We indicate the bias we observe for the SPT-3G 2018 $TT/TE/EE$ data as blue vertical lines. Dashed orange vertical lines indicate the systemic bias described in Section 4, which is subdominant. In general, with $\lesssim 0.1\sigma$, the size of the realisation-dependent scatter is at an acceptable level and even in bad cases comparable to the uncertainty inherent in an MCMC analysis. For the SPT-3G 2018 $TT/TE/EE$ data set, the only parameter with a noticeable shift is $\ln(10^{10}A_s)$, though the observed value is within 2.1 standard deviations of the mean and hence statistically normal.

In Figure 2 we compare the error bars of the CMB-only band powers to the ones of the *coadd* of multi-frequency spectra, i.e. the minimum-variance combination of the full data vector after foreground-subtraction. Since the covariance of the *coadd* only reflects the combination of multi-frequency information to reduce noise, this comparison allows us to quantify the information loss due to foregrounds and systematic effects at the band power level. Note that while the CMB-lite reconstruction does not require a cosmological model, this is not true for the *coadd* and we construct it using the best-fit point of the multi-frequency likelihood in Λ CDM. As expected, we see that the foreground contamination in temperature data leads to a gradual increase of the error bars in the lite likelihood with respect to the *coadd* towards small angular scales with the uncertainty on the final bin of the TT spectrum being a factor of 4.4 larger. For the polarisation spectra on the other hand, the lite likelihood’s error bars are inflated by $< 3\%$ more evenly across angular scales; this small increase reflects the uncertainties of the beam measurement and the calibration, as well as the expected size of super-sample lensing fluctuations for the survey field. Though neighbouring band powers in the lite likelihood remain mildly correlated due to flat-sky projection effects (Balkenhol et al. 2023), we do not introduce any significant new long-range correlations owing to our treatment of calibration.

We construct the lite likelihood based on the shown foreground-marginalised CMB-only band powers and their covariance. While we still need to account for the effects of aberration and absolute calibration at the likelihood level as mentioned above, we have reduced the number of transformations in the data model from 11 to two and the number of nuisance parameters from 33 to two (T_{cal} and E_{cal}). The length of the data vector has been reduced from 728 to 123 bins. The evaluation time of the lite likelihood is about a factor of four smaller than its multi-frequency counterpart and about half the evaluation time of the CosmoPower emulators.

The marginalised posteriors obtained from MCMC analyses of the two likelihoods in Λ CDM are shown in

Figure 3. The constraints match very well visually. This is also quantitatively true; we tabulate the difference in best-fit points and the ratio of parameter errors for all models considered in Table 1. All differences are below 0.1σ , where σ is the width of the marginalised parameter posteriors of the multi-frequency likelihood. Furthermore, the inferred parameter error bars match to within $< 10\%$ in all cases. We conclude that given the numerical noise present in MCMC analyses the lite likelihood performs well.

Still, we want to reassure ourselves that the shifts observed are compatible with the expectation from the realisation-dependent bias. Hence, we run 1000 simulations in which we draw Gaussian band power realisations from the multi-frequency band power covariance matrix based on a fiducial Λ CDM cosmology. We assign these band powers to the reference multi-frequency likelihood, minimise it, run the reconstruction procedure on them, minimise the resulting lite likelihood, and compare the obtained best-fit points, all within the Λ CDM model. We show the distribution of parameter shifts between the reference multi-frequency and lite likelihoods in Figure 4. The realisation-dependent bias is typically confined to within $\pm 0.1\sigma$, where σ is the width of the parameter posterior obtained from the multi-frequency likelihood, signalling the good robustness of the procedure. The specific reconstruction for the SPT-3G 2018 $TT/TE/EE$ data is remarkably bias free for H_0 , $\Omega_b h^2$, $\Omega_c h^2$ and n_s . The bias seen for $\ln(10^{10}A_s)$ is consistent with zero within two standard deviations of the distribution obtained from simulations and compatible with a statistical fluctuation. We further note that the minimum of the reconstruction likelihood, i.e. its value at the best-fit point, is consistent with the corresponding distribution obtained from simulations at 1.2 standard deviations from the mean. We conclude therefore that the small parameter shifts observed between the SPT-3G 2018 $TT/TE/EE$ multi-frequency and lite likelihoods are consistent with the expected size of statistical fluctuations and the reconstruction has been successful. We make the lite version of the likelihood publicly available

TABLE 1

Comparison of best-fit points and the width of marginalised posteriors between the SPT-3G 2018 $TT/TE/EE$ multi-frequency and lite likelihoods for Λ CDM, Λ CDM+ N_{eff} , and Λ CDM+ A_L models. The first column for each model lists the shift of best fit points obtained from dedicated minimiser runs normalised by the parameter error bars obtained from the MCMC analysis of the multi-frequency likelihood. The second column indicates the ratio of posterior widths. All observed shifts are $< 10\%$ the width of the posteriors obtained from the multi-frequency likelihood and the parameter errors match to within 10%, indicating good agreement between the two likelihoods.

Parameter	Λ CDM		Λ CDM+ N_{eff}		Λ CDM+ A_L	
	$\Delta\mu/\sigma^{\mu\nu}$	$\sigma^{\text{lite}}/\sigma^{\mu\nu} - 1$	$\Delta\mu/\sigma^{\mu\nu}$	$\sigma^{\text{lite}}/\sigma^{\mu\nu} - 1$	$\Delta\mu/\sigma^{\mu\nu}$	$\sigma^{\text{lite}}/\sigma^{\mu\nu} - 1$
H_0	0.000	-0.023	0.022	0.016	-0.016	0.014
$\Omega_b h^2$	-0.004	0.006	0.015	0.021	-0.008	-0.002
$\Omega_c h^2$	0.001	-0.021	0.020	0.059	0.016	0.014
n_s	0.001	-0.018	0.018	0.017	-0.009	0.026
$\ln(10^{10} A)$	0.089	0.020	0.090	0.002	0.082	0.017
τ	-0.003	0.026	-0.001	-0.008	0.000	0.003
N_{eff}	-	-	0.023	0.040	-	-
A_L	-	-	-	-	-0.019	-0.011

on the `candl` website.

5. CONCLUSIONS

In this work we have presented an implementation of the CMB-lite framework using automatic differentiation and have applied it to the differentiable SPT-3G 2018 $TT/TE/EE$ likelihood available in the `candl` library. The reconstruction runs in about a minute on a personal computer and yields a sampling noise-free covariance of the foreground-marginalised CMB-only band powers. The resulting lite likelihood performs well and recovers the parameter constraints of the reference multi-frequency likelihood for Λ CDM, Λ CDM+ N_{eff} , and Λ CDM+ A_L models to within good accuracy; best-fit points are offset by $< 0.1\sigma$, where σ is the width of the marginalised posterior of the multi-frequency likelihood, and the inferred parameter error bars match to within $< 10\%$ in all cases. The lite likelihood is publicly available on the `candl` website along with a python notebook demonstrating its construction.

More broadly, this work represents another successful application of automatic differentiation in cosmology. This technique continues to increase in popularity in the community and not only facilitates improvements of existing methodologies as is the case here, but also enables entirely novel analyses. This is a trend we expect to continue; as cosmological data continue to improve, so must our tools. Looking ahead, Simons Obser-

vatory (Simons Observatory Collaboration et al. 2019) and CMB-S4 data (CMB-S4 Collaboration et al. 2019) will have broad frequency coverage to confidently separate cosmological and foreground signals. This increases the number of multi-frequency spectra and hence also the advantages the CMB-lite framework can bring. At the same time, the low noise levels of these experiments will compel many robustness tests, which are efficiently carried out using the numerical techniques presented in this work.

L. B. is deeply grateful to Silvia Galli for her encouragement to pursue the work, helpful conversations, and useful suggestions. L. B. is also grateful to Karim Benabed for enlightening discussions and constructive suggestions. L. B. would like to thank the following people for providing comments on the manuscript: Silvia Galli, Karim Benabed, and Étienne Camphuis. This work uses JAX (Bradbury et al. 2018) and the scientific python stack (Hunter 2007; Jones et al. 2001; van der Walt et al. 2011). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101001897). This work has received funding from the Centre National d'Etudes Spatiales and has made use of the Infinity Cluster hosted by the Institut d'Astrophysique de Paris.

REFERENCES

- Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2021, *Phys. Rev. Lett.*, 127, 151301
- Aiola, S., Calabrese, E., Maurin, L., et al. 2020, *J. Cosmology Astropart. Phys.*, 2020, 047
- Arutjunjan, R., Schaefer, B. M., & Kreutz, C. 2022, arXiv e-prints, arXiv:2211.03421
- Balkenhol, L., Dutcher, D., Spurio Mancini, A., et al. 2023, *Phys. Rev. D*, 108, 023510
- Balkenhol, L. & Reichardt, C. L. 2022, *MNRAS*, 512, 4394
- Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: composable transformations of Python+NumPy programs
- Calabrese, E., Hlozek, R. A., Battaglia, N., et al. 2013, *Phys. Rev. D*, 87, 103012
- Campagne, J.-E., Lanasse, F., Zuntz, J., et al. 2023, *The Open Journal of Astrophysics*, 6, 15
- Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, *PASP*, 123, 568
- Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, *J. Cosmology Astropart. Phys.*, 2020, 045
- CMB-S4 Collaboration, Abazajian, K., Addison, G., et al. 2019, arXiv e-prints, arXiv:1907.04473
- Dodelson, S. & Schneider, M. D. 2013, *Phys. Rev. D*, 88, 063537
- Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, *J. Cosmology Astropart. Phys.*, 7, 25
- Dutcher, D., Balkenhol, L., Ade, P. A. R., et al. 2021, *Phys. Rev. D*, 104, 022003

- Ge, F., Millea, M., Camphuis, E., et al. 2024, arXiv e-prints, arXiv:2411.06000
- Hartlap, J., Simon, P., & Schneider, P. 2007, *A&A*, 464, 399
- Heavens, A. F., Seikel, M., Nord, B. D., et al. 2014, *MNRAS*, 445, 1687
- Henning, J. W., Sayre, J. T., Reichardt, C. L., et al. 2018, *ApJ*, 852, 97
- Hoffman, M. D. & Gelman, A. 2011, arXiv e-prints, arXiv:1111.4246
- Hunter, J. D. 2007, *Computing In Science & Engineering*, 9, 90
- Jeong, D., Chluba, J., Dai, L., Kamionkowski, M., & Wang, X. 2014, *Phys. Rev. D*, 89, 023003
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, *SciPy: Open source scientific tools for Python*, [Online; accessed 2014-10-22]
- Keating, B. G., Ade, P. A. R., Bock, J. J., et al. 2003, in *Polarimetry in Astronomy*. Edited by Silvano Fineschi. *Proceedings of the SPIE*, Volume 4843., 284–295
- Kosowsky, A. 2003, in the proceedings of the workshop on “The Cosmic Microwave Background and its Polarization,” *New Astronomy Reviews*, ed. S. Hanany & K. A. Olive (Elsevier)
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473
- Madhavacheril, M. S., Qu, F. J., Sherwin, B. D., et al. 2024, *ApJ*, 962, 113
- Manzotti, A., Hu, W., & Benoit-Lévy, A. 2014, *Phys. Rev. D*, 90, 023003
- Mocanu, L. M., Crawford, T. M., Aylor, K., et al. 2019, *J. Cosmology Astropart. Phys.*, 2019, 038
- Nash, S. G. 1984, *SIAM Journal on Numerical Analysis*, 21, 770
- Nocedal, J. & Wright, S. J. 2006, *Numerical Optimization* (New York, NY: Springer New York)
- Pan, Z., Bianchini, F., Wu, W. L. K., et al. 2023, *Phys. Rev. D*, 108, 122005
- Percival, W. J., Friedrich, O., Sellentin, E., & Heavens, A. 2022, *MNRAS*, 510, 3207
- Piras, D. & Spurio Mancini, A. 2023, *The Open Journal of Astrophysics*, 6, 20
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, *A&A*, 641, A1
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020b, *A&A*, 641, A6
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020c, *A&A*, 641, A5
- Planck Collaboration, Aghanim, N., Arnaud, M., et al. 2016, *A&A*, 594, A11
- Prince, H., Calabrese, E., & Dunkley, J. 2024, arXiv e-prints, arXiv:2403.00085
- Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., et al. 2024, *ApJ*, 962, 112
- Sellentin, E., Quartin, M., & Amendola, L. 2014, *MNRAS*, 441, 1831
- Simons Observatory Collaboration, Ade, P., Aguirre, J., et al. 2019, *J. Cosmology Astropart. Phys.*, 2019, 056
- Spurio Mancini, A., Piras, D., Alsing, J., Joachimi, B., & Hobson, M. P. 2022, *MNRAS*, 511, 1771
- Staniszewski, Z., Aikin, R. W., Amiri, M., et al. 2012, *Journal of Low Temperature Physics*, 167, 827
- Torrado, J. & Lewis, A. 2021, *J. Cosmology Astropart. Phys.*, 2021, 057
- van der Walt, S., Colbert, S., & Varoquaux, G. 2011, *Computing in Science Engineering*, 13, 22
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261

This paper was built using the Open Journal of Astrophysics \LaTeX template. The OJA is a journal which

provides fast and easy peer review for new papers in the **astro-ph** section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.