



Universiteit  
Leiden  
The Netherlands

## Extremely expensive likelihoods: a variational-Bayes solution for precision cosmology

Rizzato, M.; Sellentin, E.

### Citation

Rizzato, M., & Sellentin, E. (2023). Extremely expensive likelihoods: a variational-Bayes solution for precision cosmology. *Monthly Notices Of The Royal Astronomical Society*, 521(1), 1152-1161. doi:10.1093/mnras/stad638

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/3717334>

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

# Extremely expensive likelihoods: a variational-Bayes solution for precision cosmology

Matteo Rizzato<sup>1★</sup> and Elena Sellentin<sup>1,2</sup>

<sup>1</sup>*Leiden Observatory, Leiden University, Huygens Laboratory, Niels Bohrweg 2, NL-2333 CA Leiden, The Netherlands*

<sup>2</sup>*Mathematical Institute, Leiden University, Snellius Building, Niels Bohrweg 1, NL-2333 CA Leiden, The Netherlands*

Accepted 2023 February 20. Received 2023 January 8; in original form 2022 March 9

## ABSTRACT

We present a variational-Bayes solution to compute non-Gaussian posteriors from extremely expensive likelihoods. Our approach is an alternative for parameter inference when Markov chain Monte Carlo (MCMC) sampling is numerically prohibitive or conceptually unfeasible. For example, when either the likelihood or the theoretical model cannot be evaluated at arbitrary parameter values, but only previously selected values, then traditional MCMC sampling is impossible, whereas our variational-Bayes solution still succeeds in estimating the full posterior. In cosmology, this occurs e.g. when the parametric model is based on costly simulations that were run for previously selected input parameters. We demonstrate the applicability of our posterior construction on the KiDS-450 weak lensing analysis, where we reconstruct the original KiDS MCMC posterior at 0.6 per cent of its former numerical posterior evaluations. The reduction in numerical cost implies that systematic effects which formerly exhausted the numerical budget could now be included.

**Key words:** methods: analytical – methods: data analysis – methods: statistical.

## 1 INTRODUCTION

Modern cosmological surveys are designed to probe the evolved Universe with superb data precision. These surveys scout for faint signatures of new physics on top of the standard model of cosmology, or even a breakdown of the standard model itself. The precision of upcoming cosmological data sets promises the most competitive constraints on the cosmic law of gravity, the equation of state of dark energy, and dark matter (Durrer & Maartens 2008). Within the framework of a cosmological model, these data sets are also used to constrain the number of neutrino species and the neutrino masses (Audren et al. 2013; Sellentin & Durrer 2015; Hagstotz et al. 2021).

During the next decade, the forefront of cosmological surveys will be composed of the Euclid satellite (Laureijs et al. 2011) and the ground-based ‘Legacy Survey of Space and Time’ (LSST Science Collaboration et al. 2009). These surveys take over from a successful series of contemporary surveys, such as Kilo Degree Survey (Mandelbaum et al. 2018), the Dark Energy Survey (Abbott et al. 2016), and the Hyper Suprime Cam on Subaru (Mandelbaum et al. 2018). A major challenge for these surveys is that the improvement in data quality has to be matched by an equal improvement in data analysis accuracy: insufficient data analysis techniques will lose information actually contained in the data, and introduce incorrect assumptions, thereby causing biases and uncertainties larger than nominally quoted, and certainly larger than those resulting from more adequate analysis techniques.

In consequence, the field of cosmology is currently entering a phase where theoretical modelling of the data and likelihood calculations both become prohibitively expensive to compute. For

this paper, we therefore assume that the aim is to infer parameters from data which are so precise or so complex that only few model evaluations at the required precision are possible. In particular, we assume that it is impossible to evaluate the theoretical model and the likelihood  $\approx 10^5 - 10^6$  times, which is a typical number of evaluations employed by current cosmological analyses using Markov chain Monte Carlo (MCMC).

In the most extreme case, parameter inference will have to start from e.g. 50 highly accurate simulations run for different cosmological models. In this case, not only is the number of possible posterior evaluations fixed to 50, but the positions where the posterior can be evaluated are additionally pre-determined: it can only be computed at the parameter values chosen as input parameters for the simulations. In contrast, an MCMC run would require that the posterior can be evaluated at any arbitrary position throughout parameter space.

Put briefly, we are searching for a posterior construction technique that succeeds despite computing the posterior at few, potentially pre-determined fixed parameter values.

To this end, we propose in this paper a variational-Bayes solution, which fits a flexible distribution to few evaluations of the likelihood, thereby yielding an approximate posterior. The essential assumption enabling the speed-up is that the posterior be fairly smooth, though not necessarily unimodal.

Our paper is structured as follows. In Section 2, we present the class of variational-Bayes methods. These methods fit a parametric distribution to evaluations of the posterior at few points. Accordingly, the quality of the resulting fit is determined by a suitable choice of the flexible distribution. In Section 3, we therefore set up our own distribution to be fitted. The result from our variational-Bayes technique is compared in Section 4 to the publicly available MCMC from the KiDS-450 tomographic weak gravitational lensing analysis (Hildebrandt et al. 2017). We finally summarize the benefits of the

\* E-mail: [matteo.rizzato@iap.fr](mailto:matteo.rizzato@iap.fr)

method while paving the way to follow up studies in Section 5. In spirit, our variational-Bayes approach is similar to Amendola & Gómez-Valent (2020), which uses a different (sometimes divergent) variational distribution.

## 2 VARIATIONAL-BAYES METHODS

Variational-Bayes methods construct a posterior by fitting a candidate distribution with free shape parameters to a few calls of the target posterior (MacKay 2002).

In this paper, we denote the model parameters of physical interest by  $\mathbf{p}$  and the data by  $\mathbf{x}$ , the sought posterior distribution of the parameters as constrained by the data is  $\mathcal{P}(\mathbf{p}|\mathbf{x})$ . Bayes theorem relates  $\mathcal{P}(\mathbf{p}|\mathbf{x})$  to the likelihood  $\mathcal{L}(\mathbf{x}|\mathbf{p})$  and to the parameters' prior  $\pi(\mathbf{p})$  as

$$\mathcal{P}(\mathbf{p}|\mathbf{x}) \propto \mathcal{L}(\mathbf{x}|\mathbf{p})\pi(\mathbf{p}). \quad (1)$$

The aim of parameter inference is to calculate this posterior, usually quoting the maximum a posteriori (MAP) estimate of best-fitting parameters and plotting the marginal distributions of parameter pairs. For precision cosmology with expensive likelihoods, the usual MCMC sampling to this end can become unfeasible.

Variational-Bayes methods can compute posteriors at a lower numerical cost, as they provide an analytical approximation to the posterior. The approximation is computed based on a few posterior evaluations only. The approximating distribution is usually known as the *variational distribution*  $\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P})$  and it is parametrized by a set of shape parameters  $\mathbf{P}$ , also referred to as ‘variational parameters,’ whose values have to be determined such that the variational distribution approximates the wanted posterior in the sense of

$$\mathcal{P}(\mathbf{p}|\mathbf{x}) \approx \mathcal{V}(\mathbf{p}|\mathbf{x}, \hat{\mathbf{P}}), \quad (2)$$

where  $\hat{\mathbf{P}}$  are the best-fitting shape parameters. For precision cosmology, it is obviously crucial that the posterior approximation indicated in equation (2) is of convincing quality. Achieving this quality constitutes the primary challenge to be addressed in this paper.

Technically, the key to a successful variational-Bayes solution is a variational distribution that is so flexible in shape that it can accurately represent arbitrarily formed posteriors. Accordingly, we develop in this paper a variational-Bayes method where the variational distribution  $\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P})$  is a generalization of the non-Gaussian ‘DALI’ distribution presented in Sellentin, Quartin & Amendola (2014, hereafter S14) and in Sellentin (2015, hereafter S15). DALI distributions form a series and are positive definite and non-divergent at any order. They are based on a Taylor expansion, and can hence be proven to converge towards the true distribution. This Taylor expansion can be truncated at any order, but in practice truncating at the second order is almost always sufficient (Sellentin & Schäfer 2016). Accordingly, we will here build upon the second-order DALI approximation obtained in equation 15 of S14. The adoption of DALI is a deliberate choice: Fitting any arbitrary multivariate polynomial instead can lead to regions in parameter space where the posterior probability takes (meaningless) negative values, or regions of infinite probability, far away from the peak. DALI avoids these problems by construction as detailed in S14.

## 3 SETTING UP THE VARIATIONAL DISTRIBUTION

In this section, we set up the variational distribution  $\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P})$  whose shape parameters  $\mathbf{P}$  will be fitted to a few posterior samples.

We start by reviewing in Section 3.1 the DALI formalism as developed in S14 and S15, and then discuss how to optimize it for use as a variational distribution in Section 3.2. This derivation is somewhat involved and readers might want to skip directly to Section 4 where the performance of the posterior reconstruction technique is demonstrated on the KiDS-450 data.

### 3.1 DALI review

We shortly review the DALI formalism as published in S14 and S15.

Let us consider a Gaussian data set  $\mathbf{x}$  with expectation value  $\boldsymbol{\mu}$  and covariance matrix  $\mathbf{C}$ . The parameters to be inferred are  $\mathbf{p}$ , and the mean  $\boldsymbol{\mu}(\mathbf{p})$  is to be determined. For this paper it suffices to use a constant  $\mathbf{C}$ , for parameter dependent  $\mathbf{C}$  see S15.

DALI forecasts the posterior of parameters  $\mathbf{p}$ , i.e. it predicts the posterior *before* actual data are taken. To achieve these forecasts, DALI follows the line of thought of Fisher matrix forecasts (Tegmark, Taylor & Heavens 1997) by computing the expected log-posterior by averaging out the not yet collected data. Accordingly, both DALI and Fisher matrix forecasts compute

$$\mathcal{D}(\mathbf{p}) \propto \exp(\log \mathcal{P}(\mathbf{p}|\mathbf{x})), \quad (3)$$

where the angular brackets denote the average over the data  $\mathbf{x}$ . While Fisher forecasts are limited to Gaussian posterior approximations, DALI yields the actual non-Gaussian posteriors.

We now define the deviation from the maximum a posteriori point (MAP)  $\hat{\mathbf{p}}$  as

$$\boldsymbol{\Delta} = \mathbf{p} - \hat{\mathbf{p}}. \quad (4)$$

The DALI formalism at second order provides the following series expansion  $\mathcal{D}(\mathbf{p})$  to the average parameter posterior in equation (3):

$$\begin{aligned} \mathcal{D}(\mathbf{p}) \propto \exp \left[ -\frac{1}{2} \boldsymbol{\mu}_{,\alpha}^t \mathbf{C}^{-1} \boldsymbol{\mu}_{,\beta} \Delta_{\alpha} \Delta_{\beta} \right. \\ \left. - \frac{1}{2} \boldsymbol{\mu}_{,\alpha}^t \mathbf{C}^{-1} \boldsymbol{\mu}_{,\beta\gamma} \Delta_{\alpha} \Delta_{\beta} \Delta_{\gamma} \right. \\ \left. - \frac{1}{8} \boldsymbol{\mu}_{,\alpha\beta}^t \mathbf{C}^{-1} \boldsymbol{\mu}_{,\gamma\delta} \Delta_{\alpha} \Delta_{\beta} \Delta_{\gamma} \Delta_{\delta} \right]. \end{aligned} \quad (5)$$

Partial derivatives are indicated by commas,  $\boldsymbol{\mu}_{,\alpha} \equiv \partial \boldsymbol{\mu} / \partial p_{\alpha}$ ,  $\boldsymbol{\mu}_{,\alpha\beta} \equiv \partial^2 \boldsymbol{\mu} / \partial p_{\alpha} \partial p_{\beta}$  and we imply Einsteinian summation over repeated indices. Note that the first summand in equation (5) corresponds to the Fisher approximation. The remaining two summands appear simultaneously in the second order DALI approximation and jointly guarantee positive definiteness, finiteness, and convergence towards the true posterior.

In this paper, Greek indices run from 1 to  $d = \dim[\mathbf{p}]$ . The dimension of the data is  $D = \dim[\mathbf{x}]$ . We also introduce calligraphic indices, such as  $\mathcal{I}$ , in order to denote compound indices for ordered couples  $(\alpha\beta)$  within the list

$$\{(11), (12), \dots, (\alpha\beta)_{\beta \geq \alpha}, \dots, (dd)\}. \quad (6)$$

Therefore, calligraphic indices take values  $\mathcal{I} = 1, \dots, s$  and  $s \equiv d(d+1)/2$ .

### 3.2 Variational inference with DALI

DALI, as reviewed in Section 3.1, handles parameter forecasts where data have not yet been collected, and hence, the parameters are not yet *inferred*. Variational inference does, however, handle actual data and actual inference. In this section, we therefore modify the DALI formalism for use in variational inference, such that actual data can be used. In other words, the aim is to modify  $\mathcal{D}(\mathbf{p})$  of equation (5) such that it can be used as variational distribution  $\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P})$ .

We begin by noting that the components of  $\boldsymbol{\mu}_\alpha$ ,  $\boldsymbol{\mu}_{\alpha\beta}$  and  $\mathbf{C}$  in equation (5) carry physical units and have hence (conventional but otherwise) arbitrary values. We remove this arbitrariness by whitening the data, the advantage being that this whitening eliminates  $\mathbf{C}$  as degenerate free shape parameters we would otherwise need to fit for.

From this point on, we will often implicitly imply the dependence on the measured data set  $\mathbf{x}$  in order to simplify the notation.

We introduce the Cholesky decomposition of  $\mathbf{C}^{-1}$  in terms of a lower triangular matrix  $\mathbf{L}$  as  $\mathbf{L}\mathbf{L}^\dagger = \mathbf{C}^{-1}$  and define the vectors

$$\mathbf{V}_\alpha \equiv \mathbf{L}^\dagger \boldsymbol{\mu}_\alpha, \quad \mathbf{M}_{\alpha\beta} \equiv \mathbf{L}^\dagger \boldsymbol{\mu}_{\alpha\beta}, \quad (7)$$

and write equation (5) as

$$\begin{aligned} \mathcal{V}(\mathbf{p}|\mathbf{P}) \propto \exp \left[ -\frac{1}{2} \mathbf{V}_\alpha^\dagger \mathbf{V}_\beta \Delta_\alpha \Delta_\beta \right. \\ \left. -\frac{1}{2} \mathbf{V}_\alpha^\dagger \mathbf{M}_{\beta\gamma} \Delta_\alpha \Delta_\beta \Delta_\gamma \right. \\ \left. -\frac{1}{8} \mathbf{M}_{\alpha\beta}^\dagger \mathbf{M}_{\gamma\delta} \Delta_\alpha \Delta_\beta \Delta_\gamma \Delta_\delta \right]. \end{aligned} \quad (8)$$

The scalar products between the vectors  $\mathbf{V}_\alpha$  and  $\mathbf{M}_{\alpha\beta}$  are invariant under the orthogonal group: while any orthogonal matrix simultaneously applied to these vectors would affect their components, their scalar products will remain invariant. Hence, we remove the dependence on the vector orientation by expressing the scalar products in terms of the vector lengths  $V_\alpha$ ,  $M_{\alpha\beta}$  and of the cosine of the angle between them, as follows:

$$\mathbf{V}_\alpha^\dagger \mathbf{V}_\beta = V_\alpha V_\beta \cos(\theta_{\alpha\beta}), \quad (9)$$

$$\mathbf{V}_\alpha^\dagger \mathbf{M}_{\beta\gamma} = V_\alpha M_{\beta\gamma} \cos(\theta_{\alpha;\beta\gamma}), \quad (10)$$

$$\mathbf{M}_{\alpha\beta}^\dagger \mathbf{M}_{\gamma\delta} = M_{\alpha\beta} M_{\gamma\delta} \cos(\theta_{\alpha\beta;\gamma\delta}). \quad (11)$$

The aim is then to construct a variational distribution  $\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P})$  such that

$$\mathcal{V}(\mathbf{p}|\mathbf{x}, \mathbf{P}) \propto \exp[-\mathcal{Y}(\boldsymbol{\Delta}, \mathbf{x}, \mathbf{P})], \quad (12)$$

$$\mathcal{Y}(\boldsymbol{\Delta}, \mathbf{x}, \mathbf{P}) = \frac{1}{2} \mathbf{B}^\dagger(\boldsymbol{\Delta}, \mathbf{x}, \mathbf{P}) \mathbf{B}(\boldsymbol{\Delta}, \mathbf{x}, \mathbf{P}), \quad (13)$$

where  $\mathcal{Y}(\boldsymbol{\Delta}, \mathbf{x}, \mathbf{P})$  has to be a quadratic form in order to guarantee a normalizable, positive definite posterior.

To advance towards the definition of a vector  $\mathbf{B}$  as in equation (13), we define the vectors  $\boldsymbol{\Delta}^\vee$  and  $\boldsymbol{\Delta}^\mathfrak{m}$  as

$$\boldsymbol{\Delta}^\vee \equiv \{\Delta_1 V_1, \dots, \Delta_d V_d\}, \quad (14)$$

$$\begin{aligned} \boldsymbol{\Delta}^\mathfrak{m} \equiv \{\Delta_1 \Delta_1 M_1, 2\Delta_1 \Delta_2 M_2, \dots, \\ 2(1 - \delta_{\alpha\beta}) \Delta_\alpha \Delta_\beta M_{(\alpha\beta)}, \dots, \Delta_d \Delta_d M_d\}, \end{aligned} \quad (15)$$

and the matrices  $\mathbf{A}^\vee$ ,  $\mathbf{A}^{\vee\mathfrak{m}}$  and  $\mathbf{A}^\mathfrak{m}$ , whose components are

$$A_{\alpha\beta}^\vee \equiv \cos(\theta_{\alpha\beta}), \quad (16)$$

$$A_{\alpha\mathcal{I}}^{\vee\mathfrak{m}} \equiv \cos(\theta_{\alpha;\mathcal{I}}), \quad (17)$$

$$A_{\mathcal{I}\mathcal{J}}^\mathfrak{m} \equiv \cos(\theta_{\mathcal{I}\mathcal{J}}). \quad (18)$$

In this notation, the negative logarithm of the variational distribution from equation (12) is

$$\mathcal{Y}(\boldsymbol{\Delta}, \mathbf{P}) = \frac{1}{2} (\boldsymbol{\Delta}^\vee)^\dagger \mathbf{A}^\vee \boldsymbol{\Delta}^\vee + \frac{1}{2} (\boldsymbol{\Delta}^\vee)^\dagger \mathbf{A}^{\vee\mathfrak{m}} \boldsymbol{\Delta}^\mathfrak{m} + \frac{1}{8} (\boldsymbol{\Delta}^\mathfrak{m})^\dagger \mathbf{A}^\mathfrak{m} \boldsymbol{\Delta}^\mathfrak{m}. \quad (19)$$

However, equation (19) is not automatically a quadratic form, if the matrices  $\mathbf{A}^\vee$ ,  $\mathbf{A}^{\vee\mathfrak{m}}$  and  $\mathbf{A}^\mathfrak{m}$  are independent of each other. Accordingly, the necessary relation between these matrices still has to be enforced, such that equation (19) is a valid candidate for our log-posterior. We hence relate these matrices to each other by jointly decomposing them as follows.

The Fisher matrix  $\mathbf{F}$  is

$$F_{\alpha\beta} = \mathbf{V}_\alpha^\dagger \mathbf{V}_\beta = V_\alpha V_\beta \cos(\theta_{\alpha\beta}) = V_\alpha V_\beta A_{\alpha\beta}^\vee, \quad (20)$$

and has to be positive definite. Consequently,  $\mathbf{A}^\vee$  has to be positive definite as well. We therefore Cholesky-decompose it

$$\mathbf{A}^\vee = \mathbf{L}^\vee (\mathbf{L}^\vee)^\dagger, \quad L_{\alpha\beta}^{\vee>\alpha} = 0, \quad L_{\alpha\alpha}^\vee \geq 0. \quad (21)$$

Given the scalar products of equation (16), we have the constraint

$$A_{\alpha\alpha}^\vee = 1. \quad (22)$$

To enforce this condition, we request the rows  $\{\mathbf{v}_\alpha\}$  of the matrix  $\mathbf{L}^\vee$  to be vectors of unitary norm.

As for  $\mathbf{A}^\mathfrak{m}$ , we point out that if it is positive definite, then we will have

$$\mathbf{A}^\mathfrak{m} \text{ positive definite} \Leftrightarrow \forall \boldsymbol{\Delta}^\mathfrak{m} \in \mathbb{R}^s, \quad (\boldsymbol{\Delta}^\mathfrak{m})^\dagger \mathbf{A}^\mathfrak{m} \boldsymbol{\Delta}^\mathfrak{m} \geq 0. \quad (23)$$

Hence, we see that  $\mathbf{A}^\mathfrak{m}$  must indeed be positive definite, otherwise the scalar product with  $\boldsymbol{\Delta}^\mathfrak{m}$  will pick up a minus sign, which will cause the posterior to diverge at large distances from the peak. We therefore Cholesky-decompose it too, yielding

$$\mathbf{A}^\mathfrak{m} \equiv \mathbf{L}^\mathfrak{m} (\mathbf{L}^\mathfrak{m})^\dagger. \quad (24)$$

This allows us to implement the prior constraints

$$L_{\mathcal{I}\mathcal{J}>\mathcal{I}}^\mathfrak{m} = 0, \quad L_{\mathcal{I}\mathcal{I}}^\mathfrak{m} > 0. \quad (25)$$

Additionally, we again demand the rows  $\{\mathbf{m}_\mathcal{I}\}$  of  $\mathbf{L}^\mathfrak{m}$  to be vectors of unitary norm, which ensures that

$$A_{\mathcal{I}\mathcal{I}}^\mathfrak{m} = 1. \quad (26)$$

We now turn to the decomposition of  $\mathbf{A}^{\vee\mathfrak{m}}$  which is a rectangular matrix. Decomposing this matrix correctly is crucial to yield a positive-definite and normalizable posterior. We hence review the QR-decomposition of rectangular matrices in Appendix B.

We introduce rectangular orthogonal matrices  $\mathbf{Q} \in \mathbb{R}^{n \times p}$ ,  $np$ , which satisfy the relation  $\mathbf{Q}^\dagger \mathbf{Q} = \mathbf{1}_{p \times p}$ . For  $n = p$ , we obtain the elements of the orthogonal group  $\mathbf{O}(n)$ . Furthermore, we define the following matrices:

$$\mathbf{1}_{d_1 \times d_2} \in \mathbb{R}^{d_1 \times d_2}, \quad d_1 \geq d_2, \quad [\mathbf{1}_{d_1 \times d_2}]_{ij} \equiv \begin{cases} \delta_{ij} & \text{for } i \leq d_1, \\ 0 & \text{elsewhere.} \end{cases} \quad (27)$$

Matrices defined by equation (27) will be the usual identity matrices only if square. If rectangular, additional zero-rows appear, e.g. we have

$$\mathbf{1}_{3 \times 2} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}. \quad (28)$$

We now introduce an auxiliary dimension  $D$  and set  $D \geq sd$ : this allows the  $s$  vectors  $\{\mathbf{m}_\mathcal{I}\}$  to be linearly independent, such that the matrix of their relative angles  $\mathbf{A}^\mathfrak{m}$  is as general as can be. We gather the  $d$  vectors  $\{\mathbf{V}_\alpha\}$  as the columns of a matrix  $\mathbf{V} \in \mathbb{R}^{D \times d}$  that we QR-decompose as

$$\mathbf{V} = \mathbf{Q}^\vee \mathbf{R}^\vee, \quad \mathbf{Q}^\vee \in \mathbf{O}(D \times d), \quad \mathbf{R}^\vee \in \mathbb{R}^{d \times d}, \quad (29)$$

$\mathbf{R}^\vee$  being upper triangular. This decomposition is unique whenever  $\mathbf{R}_{ii}^\vee > 0$ . If we define the normalization matrices  $\mathbf{N}^\vee$ ,  $\mathbf{N}^\mathfrak{m}$  as

$\mathbf{N}_{ij}^v = \delta_{ji}/V_i$  and  $\mathbf{N}_{\mathcal{I}\mathcal{J}}^m = \delta_{\mathcal{I}\mathcal{J}}/M_{\mathcal{I}}$ , respectively, we can then write the matrix  $\mathbf{A}^v$  from equation (16) as

$$\mathbf{A}^v = \mathbf{N}^v \mathbf{V}^t \mathbf{V} \mathbf{N}^v = \mathbf{N}^v (\mathbf{R}^v)^t (\mathbf{Q}^v)^t \mathbf{Q}^v \mathbf{R}^v \mathbf{N}^v = \left(\tilde{\mathbf{R}}^v\right)^t \tilde{\mathbf{R}}^v, \quad (30)$$

$$\tilde{\mathbf{R}}^v \equiv \mathbf{R}^v \mathbf{N}^v. \quad (31)$$

However, since the Cholesky decomposition for a positive definite matrix is unique, we conclude, by comparing equation (21) and equation (30), that  $\mathbf{L}^v = \left(\tilde{\mathbf{R}}^v\right)^t$  and in particular

$$\mathbf{V} \mathbf{N}^v = \mathbf{Q}^v (\mathbf{L}^v)^t. \quad (32)$$

We can perform similar calculations for  $\mathbf{A}^m$  from equation (18) and the matrix  $\mathbf{M}$  obtained from the vectors  $\{M_{\mathcal{I}}\}$  leading to

$$\mathbf{M} \mathbf{N}^m = \mathbf{Q}^m (\mathbf{L}^m)^t, \quad \mathbf{Q}^m \in \mathcal{O}(D \times s). \quad (33)$$

We hence decompose the matrix  $\mathbf{A}^{vm}$  from equation (17) as

$$\mathbf{A}^{vm} = \mathbf{N}^v \mathbf{V}^t \mathbf{M} \mathbf{N}^m = \mathbf{L}^v \left[ (\mathbf{Q}^v)^t \mathbf{Q}^m \right] (\mathbf{L}^m)^t \equiv \mathbf{L}^v \mathbf{R}^{vm} (\mathbf{L}^m)^t. \quad (34)$$

The rectangular orthogonal matrices  $\mathbf{Q}^v$  and  $\mathbf{Q}^m$  can be parametrized as (Shepard, Brozell & Gidofalvi 2015)

$$\mathbf{Q}^v = \tilde{\mathbf{Q}}^v \mathbf{1}_{D \times d}, \quad \tilde{\mathbf{Q}}^v \in \mathcal{O}(D), \quad (35)$$

$$\mathbf{Q}^m = \tilde{\mathbf{Q}}^m \mathbf{1}_{D \times s}, \quad \tilde{\mathbf{Q}}^m \in \mathcal{O}(D), \quad (36)$$

providing the following structure to  $\mathbf{R}^{vm}$

$$\mathbf{R}^{vm} = \mathbf{1}_{D \times d}^t \mathbf{Q}^{vm} \mathbf{1}_{D \times s}, \quad (\mathbf{Q}^{vm})^t \in \mathcal{O}(D), \quad (37)$$

as  $\mathbf{Q}^{vm} = (\tilde{\mathbf{Q}}^v)^t \tilde{\mathbf{Q}}^m$ . Finally, by virtue of equation (34), we can write  $\mathcal{Y}(\mathbf{\Delta}, \mathbf{P})$  as a quadratic form defining the vector  $\mathbf{B}(\mathbf{\Delta}, \mathbf{P})$  as

$$\mathbf{B}(\mathbf{\Delta}, \mathbf{P}) \equiv \mathbf{Q}^v (\mathbf{L}^v)^t \mathbf{\Delta}^v + \frac{1}{2} \mathbf{Q}^m (\mathbf{L}^m)^t \mathbf{\Delta}^m, \quad \mathbf{B} \in \mathbb{R}^D. \quad (38)$$

We can simplify the expression for  $\mathbf{B}(\mathbf{\Delta}, \mathbf{P})$  by noticing that  $\mathcal{Y}(\mathbf{\Delta}, \mathbf{P})$  has a symmetry given by the group  $\mathcal{O}(D)$

$$\mathbf{B} \xrightarrow{\mathcal{O}(D)} \mathbf{H} \mathbf{B}, \quad \mathbf{H} \in \mathcal{O}(D) \quad (39)$$

$$\mathcal{Y}_{\mathbf{H}} = (\mathbf{H} \mathbf{B})^t (\mathbf{H} \mathbf{B}) = \mathbf{B}^t \mathbf{B} = \mathcal{Y}. \quad (40)$$

In line with Section 3.2, we want to reduce the number of actual degrees of freedom (dofs) by gauging this symmetry via e.g.  $\mathbf{Q}^m \equiv \mathbf{1}_{D \times s}$ , hence leading to

$$\mathbf{R}^{vm} = \mathbf{1}_{D \times d}^t (\tilde{\mathbf{Q}}^v)^t \mathbf{1}_{D \times s}, \quad (41)$$

$$\mathbf{B} = \mathbf{Q}^v (\mathbf{L}^v)^t \mathbf{\Delta}^v + \frac{1}{2} \mathbf{1}_{D \times s} (\mathbf{L}^m)^t \mathbf{\Delta}^m. \quad (42)$$

In general, the identification above does not gauge all the degrees of freedom associated with the transformation in equation (39) as  $\mathbf{Q}^m \equiv \mathbf{1}_{D \times s}$  fixes  $\dim[\mathcal{O}(d \times D)] = Ds - s(s+1)/2$  dofs. and  $\dim[\mathcal{O}(d \times D)] \leq \dim[\mathcal{O}(D)]$  for  $s \leq D$ .

At this point, it is not possible to simplify any further our formalism without any information on the dimension  $D = \dim[B]$ . We set  $D \equiv s$  as it minimizes the number of shape parameters to fit for while guaranteeing the positive definiteness of  $\mathbf{A}^v$  and  $\tilde{\mathbf{Q}}^v \in \mathcal{O}(s)$ . Under this assumption, we have that

$$\tilde{\mathbf{Q}}^v \in \mathcal{O}(s), \quad (43)$$

$$\mathbf{R}^{vm} = \mathbf{1}_{s \times d}^t (\tilde{\mathbf{Q}}^v)^t. \quad (44)$$

and the gauge  $\mathbf{Q}^m \equiv \mathbf{1}_{s \times s}$  fixes all the degrees of freedom associated with the symmetry from equation (39).

We then parametrize the matrix  $\tilde{\mathbf{Q}}^v$  as

$$\tilde{\mathbf{Q}}^v \equiv \prod_{q=1}^d \mathbf{H}_{v^{(q)}}, \quad (45)$$

with  $\mathbf{H}_{v^{(q)}} \in \mathbb{R}^{s \times s}$  being the Householder reflector of order  $q$ . Such operator is determined by a vector  $v^{(q)} \in \mathbb{R}^{s-q}$  (Shepard et al. 2015). The components of the vectors  $\{v^{(q)}\}$  are chosen to be among the shape parameters of our formalism.

### 3.2.1 Loss function and dimensionality of the problem

We finally summarize our strategy. The variational distribution we designed is

$$\mathcal{V}(\mathbf{p}|\mathbf{P}) \propto \exp[-\mathcal{Y}(\mathbf{\Delta}, \mathbf{P})], \quad (46)$$

$$\mathcal{Y}(\mathbf{\Delta}, \mathbf{P}) = \frac{1}{2} \left\| \tilde{\mathbf{Q}}^v \mathbf{1}_{s \times d} (\mathbf{L}^v)^t \mathbf{\Delta}^v + \frac{1}{2} (\mathbf{L}^m)^t \mathbf{\Delta}^m \right\|^2, \quad (47)$$

$$\mathbf{\Delta} = \mathbf{p} - \hat{\mathbf{p}}, \quad \tilde{\mathbf{Q}}^v \equiv \prod_{q=1}^d \mathbf{H}^{(q)}, \quad (48)$$

and it is parametrized in terms of the following vector of shape parameters

$$\mathbf{P} = \left\{ \left\{ \mathbf{v}_{\alpha} \right\}_{\alpha=1, \dots, d-1}, \left\{ \mathbf{m}_{\mathcal{I}} \right\}_{\mathcal{I}=1, \dots, s-1}, \left\{ \mathbf{v}^{(q)} \right\}_{q=1, \dots, d}, \right. \\ \left. \times \left\{ V_{\alpha} \right\}_{\alpha=1, \dots, d}, \left\{ M_{\mathcal{I}} \right\}_{\mathcal{I}=1, \dots, s}, \hat{\mathbf{p}} \right\}. \quad (49)$$

In total, for a parameter space of dimension  $d = \dim[\mathbf{p}]$ , the associated number of variational parameters is

$$N_{\text{shape}} \equiv \dim[\mathbf{P}] = \frac{1}{8} (10d + 7d^2 + 6d^3 + d^4). \quad (50)$$

In Appendix A, we summarize how different parts of the variational distribution from equation (46) depend on the shape parameters. Starting from a set of randomly sampled points,

$$\left\{ \mathcal{P}_i; \mathbf{p}_i \right\}_{i=1, \dots, n_s}, \quad \mathcal{P}_i = \mathcal{P}(\mathbf{p}_i|\mathbf{x}), \quad (51)$$

we identify the optimal shape parameters  $\mathbf{P}$  such that  $\mathcal{V}(\mathbf{p}|\mathbf{P}) \approx \mathcal{P}(\mathbf{p})$  by minimizing a quadratic loss function,

$$\Phi(\mathbf{P}) = \frac{1}{n_s} \sum_{i=1}^{n_s} \phi^2(\mathcal{P}_i, \mathcal{V}_i), \quad \mathcal{V}_i \equiv \mathcal{V}(\mathbf{p}_i|\mathbf{P}). \quad (52)$$

Depending on the application, different definitions can be chosen for the scalar function  $\phi$ , which we call cost. In this paper, we choose

$$\phi(\mathcal{P}_i, \mathcal{V}_i) = \left( \frac{|\mathcal{V}_i - \mathcal{P}_i|}{|\mathcal{V}_i| + |\mathcal{P}_i|} \right)^{\frac{1}{2}}. \quad (53)$$

## 4 (RE-)CONSTRUCTION OF THE KIDS-450 POSTERIOR

In this section, we apply our variational-Bayes approach on the KiDS-450 analysis of Hildebrandt et al. (2017), which provides a public MCMC chain<sup>1</sup> consisting of  $\sim 4 \times 10^5$  accepted samples (after burn-in removal) from a seven-dimensional posterior. When accounting also for rejected points and including the burn-in, the provided chain

<sup>1</sup>[http://kids.strw.leidenuniv.nl/cs2016/MCMC\\_README.html](http://kids.strw.leidenuniv.nl/cs2016/MCMC_README.html)

counts  $\sim 3 \times 10^6$  samples, which represents the true numerical cost of this MCMC run.

We split the samples in a test set of  $N_{\text{test}}$  samples and a smaller fitting set of  $N_{\text{fit}}$  samples. The variational-Bayes algorithm fits the variational distribution only to the fitting set. The test set would in a realistic application not be available and is here used only to evaluate the deviation between our variational-Bayes posterior and the KiDS-450 MCMC posterior.

We compare the two posteriors by a symmetrized version of a percentual deviation, known as SMAPE (Symmetrized Mean Absolute Percentage Error). At each point  $\mathbf{p}_i$  of the test set, the final variational-Bayes posterior  $\mathcal{V}(\mathbf{p}_i|\mathbf{x}, \hat{\mathbf{P}})$  is compared to the MCMC posterior  $\mathcal{P}(\mathbf{p}_i|\mathbf{x})$  by the SMAPE, defined as

$$\text{SMAPE}_i \equiv \frac{|\mathcal{V}(\mathbf{p}_i|\mathbf{x}, \hat{\mathbf{P}}) - \mathcal{P}(\mathbf{p}_i|\mathbf{x})|}{|\mathcal{V}(\mathbf{p}_i|\mathbf{x}, \hat{\mathbf{P}})| + |\mathcal{P}(\mathbf{p}_i|\mathbf{x})|}. \quad (54)$$

This SMAPE can take values between 0 (exact agreement) and 1 (maximally discrepant). It is similar to a percentual deviation, only that it correctly reports a non-zero deviation also when either of the two compared quantities can take zero values. The SMAPE will therefore detect it when either of the two posteriors vanishes while the other does not.

The local SMAPE at each point can then be averaged into a single global estimate:

$$\widehat{\text{SMAPE}} = \frac{1}{N} \sum_{i=1}^N \text{SMAPE}_i w_i, \quad (55)$$

where we introduced the weight

$$w_i = \max(\mathcal{P}(\mathbf{p}_i|\mathbf{x}), \mathcal{V}(\mathbf{p}_i|\mathbf{x}, \hat{\mathbf{P}})). \quad (56)$$

The introduction of this weight causes points of high posterior value to dominate the global SMAPE estimate. Simultaneously, points in low-density regions have a lower impact. In this way, the global SMAPE correctly reports the quality of posterior reconstruction within a few standard deviations from the peak, rather than reporting the quality of reconstruction in the posterior tails.

#### 4.1 Two-dimensional posteriors

We begin with the simpler task of reconstructing six 2D marginals derived from the otherwise seven-dimensional KiDS-450 posterior. The MCMC marginals were obtained by binning the KiDS chain. A few points were then randomly selected and provided to the variational-Bayes algorithm to fit.

The fitter updates the values of the shape parameters  $\mathbf{P}$  and requires a loss function  $\phi$  to be minimized. We chose equation (55) as the loss function, noting that this choice can be exchanged.

Figs 1 and 2 display the original MCMC posterior in blue, and our variational-Bayes posterior in red. In Fig. 1, we considered  $N_{\text{fit}} \approx 3N_{\text{shape}}$  (45 fitting points). In Fig. 2, we tested the variational-Bayes approach under the extreme condition of providing the minimally possible number of fitting points, namely  $N_{\text{fit}} = N_{\text{shape}} = 14$ .

In both these cases, the number of posterior evaluations is orders of magnitude smaller than for a typical MCMC chain. We mark the two posterior peaks with a red or blue star and indicate the 68 per cent, 90 per cent, and 95 per cent contours. The position of the fitting points is indicated by black crosses. The green background colour map reflects the local value for the SMAPE from equation (54).

In Fig. 1, we employ (roughly) three times more fitting points than shape parameters to be optimized. In all studied cases, this led to the

variational-Bayes solution being virtually indistinguishable from the original MCMC posterior.

As the accuracy of the variational-Bayes solution must decrease for decreasing number of fitting points, we study in Fig. 2 the limiting case of  $N_{\text{fit}} = N_{\text{shape}}$ . As expected, the quality of the variational-Bayes solution then becomes highly case dependent: it may or may not be an accurate representation of the true posterior.

While the purpose of the current work is to formally obtain the expression in equation (46) and prove its applicability to realistic use cases, we defer to future work a thorough study on the robustness of the reconstruction accuracy against the minimum number of training samples and their distribution.

#### 4.2 Seven-dimensional posterior

In this section, we apply the variational-Bayes approach to the construction of the joint seven-dimensional posterior of the KiDS-450 analysis (Hildebrandt et al. 2017). According to equation (50), the number of shape parameters to fit for is  $N_{\text{shape}} = 608$ .

In higher dimensions, we found that the speed of the fitting procedure can be vastly accelerated if the fitter is aided in the following two points.

First, the slowest shape parameters to estimate are the coordinates of the posterior peak  $\hat{\mathbf{p}}$ . As  $\hat{\mathbf{p}}$  enters the variational distribution through  $\Delta = \mathbf{p} - \hat{\mathbf{p}}$ , each time the fitter updates the coordinates of the peak, all other shape parameters lead to a momentarily worse fit and have to be updated too. As there are many more shape parameters other than the peak coordinates, this leads to a slow-down of the fitter which can be prevented if the peak coordinates are estimated first. When providing the coordinates of the peak known from the MCMC chain, we therefore observed a speed-up of the fitter.

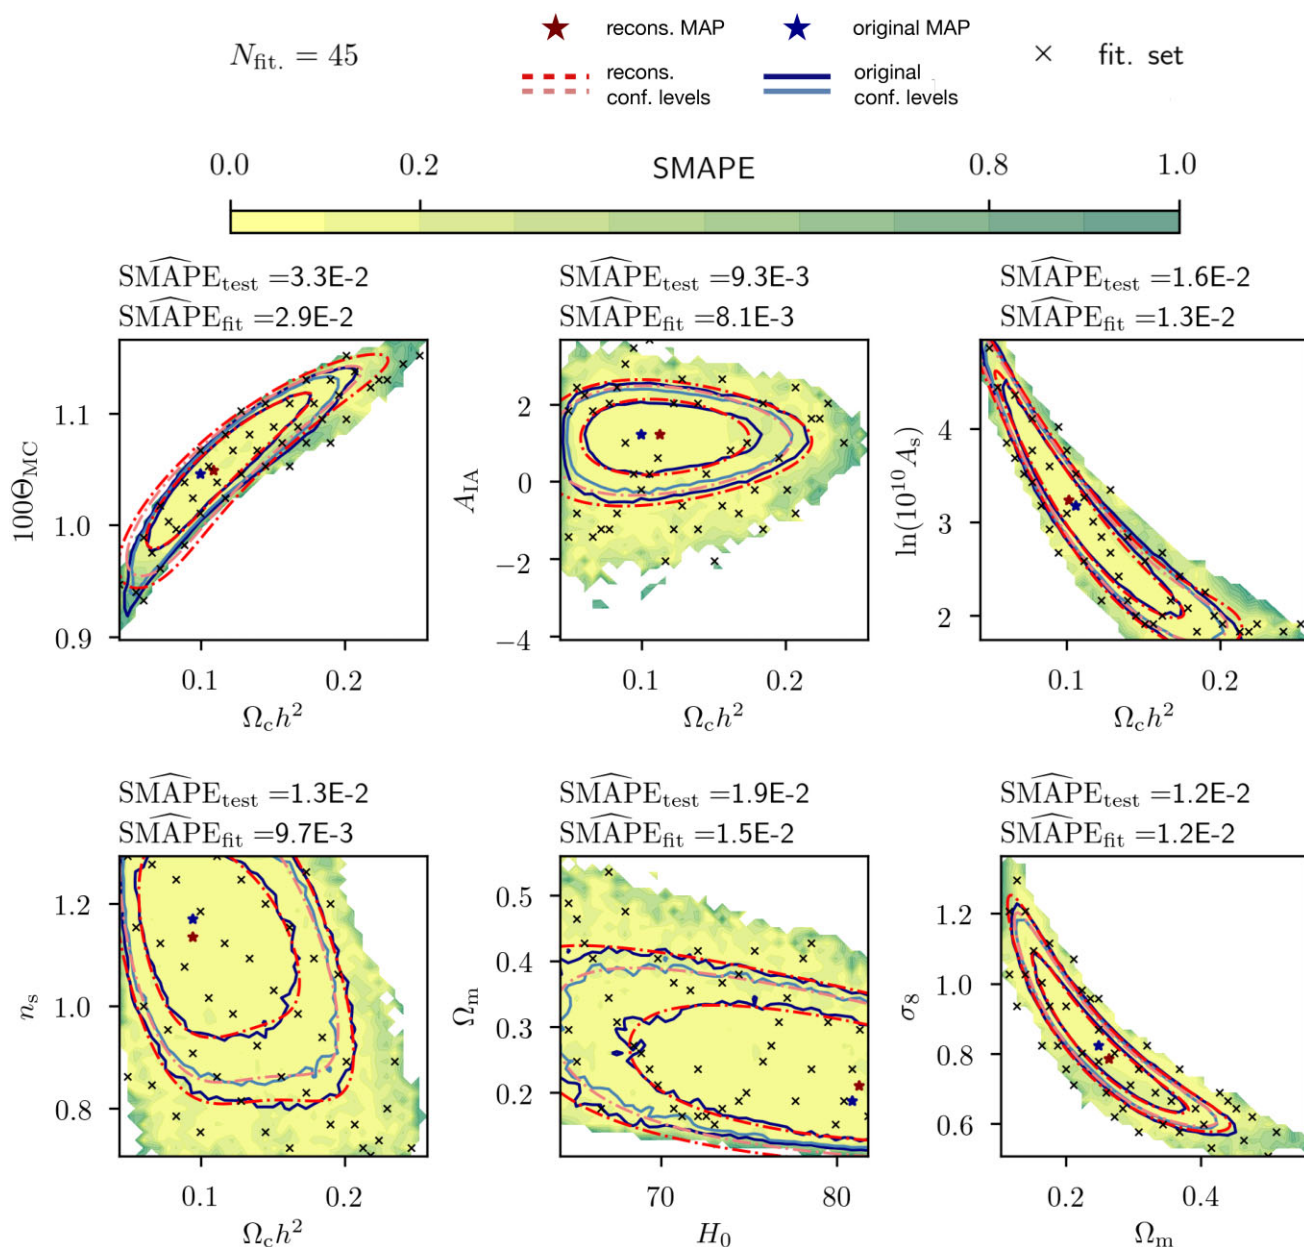
Secondly, the fitter's speed can be accelerated when providing it with extra information in order to alleviate the burden resulting from the curse of dimensionality. The curse of dimensionality describes that the corners of an  $n$ -dimensional space take up ever larger fractions in comparison to the volume of an  $n$ -sphere. In higher dimensions, the fitter will thus encounter ever larger volumes about which it has no information. More fitting points must then be provided. However, we here argue that those additional fitting points do not need to stem from an evaluation of the likelihood; the remote corners of the  $n$ -dimensional space have negligible probability anyhow, making an evaluation of the expensive likelihood unnecessary. To assist the fitter in handling the curse of dimensionality, we thus provide it with artificial points where we set the posterior density to zero.

The triangular plot in Fig. 3 shows the result offered by the variational-Bayes approach for the seven-dimensional KiDS-450 posterior. We employed  $N_{\text{fit}} = 17\,800$  posterior points randomly selected from the KiDS-450 MCMC chain. These are points where the actual likelihood code was run. We augment these by 8900 points where we set the posterior density to zero. This set-up allows us to reach a final score of  $\widehat{\text{SMAPE}}_{\text{test}} = 0.015$  on the selected samples. The number of likelihood evaluations used in our model is 0.6 per cent of the total original MCMC posterior samples.

## 5 DISCUSSION AND CONCLUSIONS

### 5.1 Discussions

We have formally derived a novel variational-Bayes solution to compute non-Gaussian posteriors from extremely expensive likelihoods. We explored its geometrical properties and applied it to a realistic use case. We made a first attempt in evaluating the reconstruction

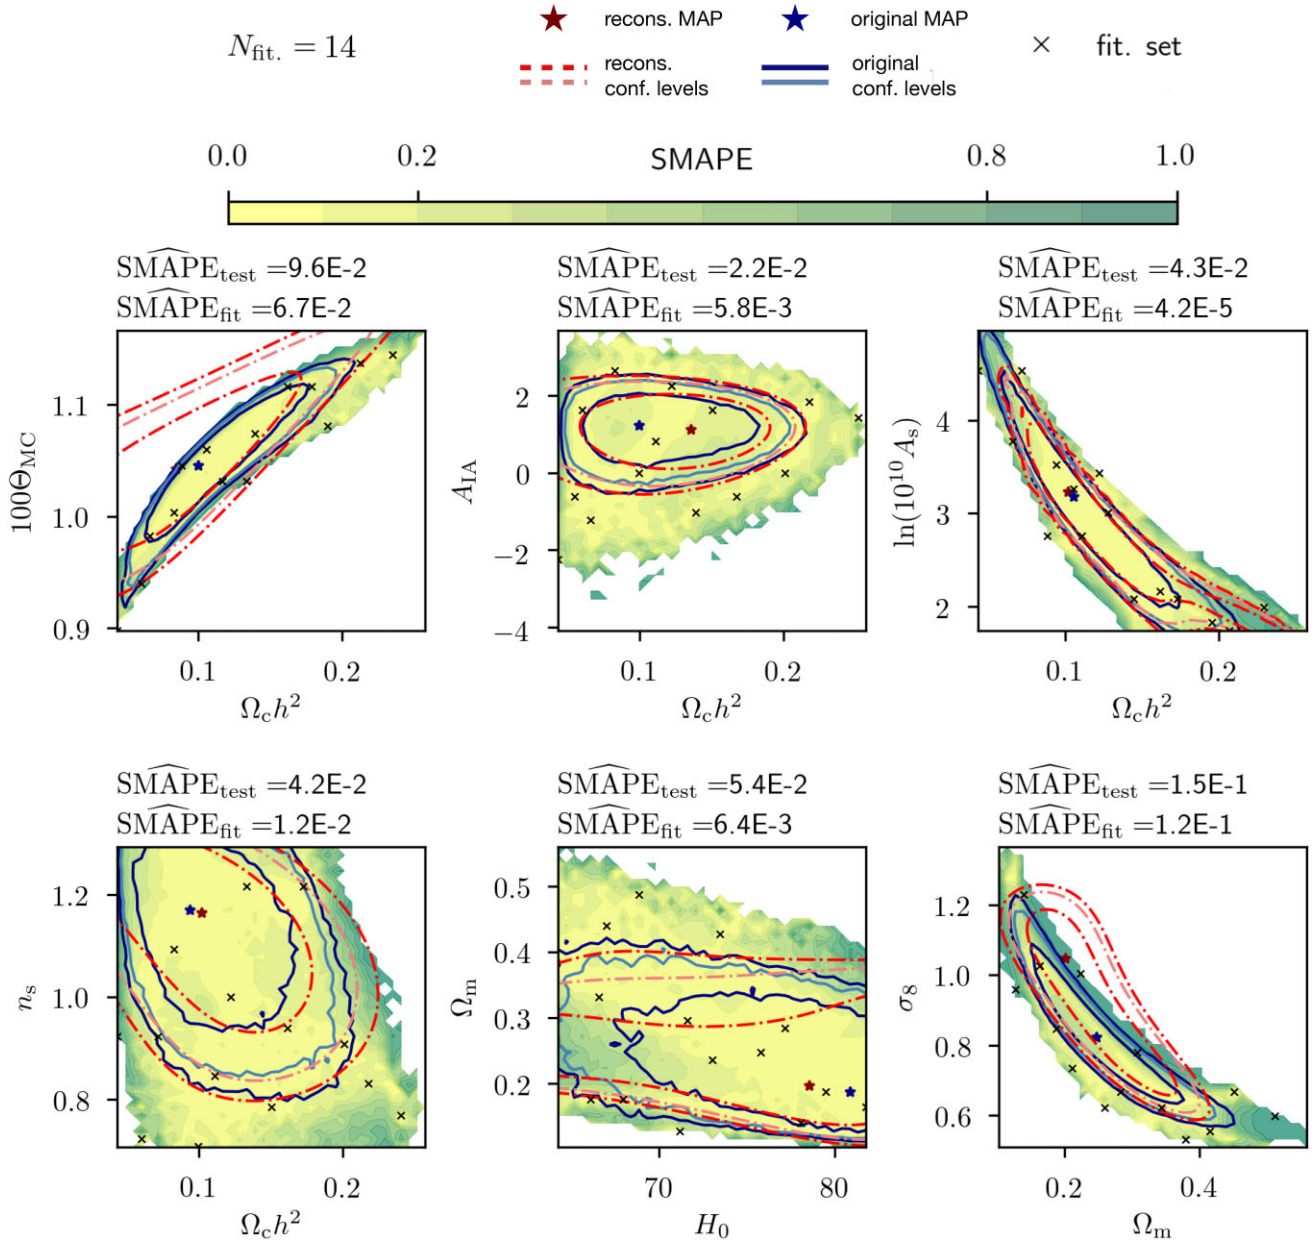


**Figure 1.** Constructing two-dimensional KiDS-450 marginals from  $N_{\text{fit}} = 45$  posterior samples only. The points used for construction are indicated by black crosses, while the 68 per cent, 90 per cent and 95 per cent confidence levels obtained via the variational-Bayes posterior are depicted in shades of red. The confidence levels referring the original KiDS-450 posterior, are depicted in shades of blue, these being obtained by marginalizing the original MCMC chain of  $\sim 3 \times 10^6$  (accepted and rejected) samples. The dark red and blue stars, respectively, indicate the MAP for the reconstructed and the original posterior. The yellow and green background map indicates that differences between the two approaches only affect the posterior tails.

performance of the method proposed optimizing the parameters of the model against subsamples of the original MCMC chains. In this final section, we pave the way to future studies aiming at validating the method via suited performance metrics and statistical tests.

In the figures presented as support to our analysis, it is easy to judge the quality of the variational-Bayes posterior, as it can be compared to the original MCMC posterior. In a realistic set-up, the MCMC posterior would not be available for comparison, and the quality of the variational-Bayes posterior has to be judged from the few likelihood evaluations of the fitting set only. One summary statistic that can be evaluated from the fitting set without comparison to an MCMC run is the global SMAPE from equation (55), which we

indicate in our figures. We found that values below 0.02 can indicate an accurate posterior reconstruction, but counter examples exist. For example, Fig. 2 evidences that global SMAPE values of  $10^{-4}$  can still yield inaccurate multimodal posteriors. In general, performance metrics evaluated on the fitting set do not assess how well the model generalizes out-of-sample. Accordingly, we found the most reliable intrinsic judgement of fitting quality to be a Leave-N-Out approach. We divide the fitting set into  $N$  random subsets with fixed cardinality. We then repeat the variational-Bayes optimization multiple times, randomly leaving out one of the  $N$  subsets, being subsequently employed as test set for performance evaluation. If all fits yield the same posterior, then the entire fitting set is likely to include



**Figure 2.** Stress-testing the variational-Bayes solution by providing the absolute minimum number of posterior samples: the variational distribution has 14 free parameters, hence we provided only  $N_{\text{fit}} = 14$  points. Comparing the original confidence levels (shades of blue) to the ensuing variational-Bayes ones (shades of red), we find a mixture of good and insufficient reconstructions of the true posterior shape. This illustrates how rapidly the variational distribution converges towards the true posterior. The dark red and blue stars, respectively, indicate the MAP for the reconstructed and the original posterior.

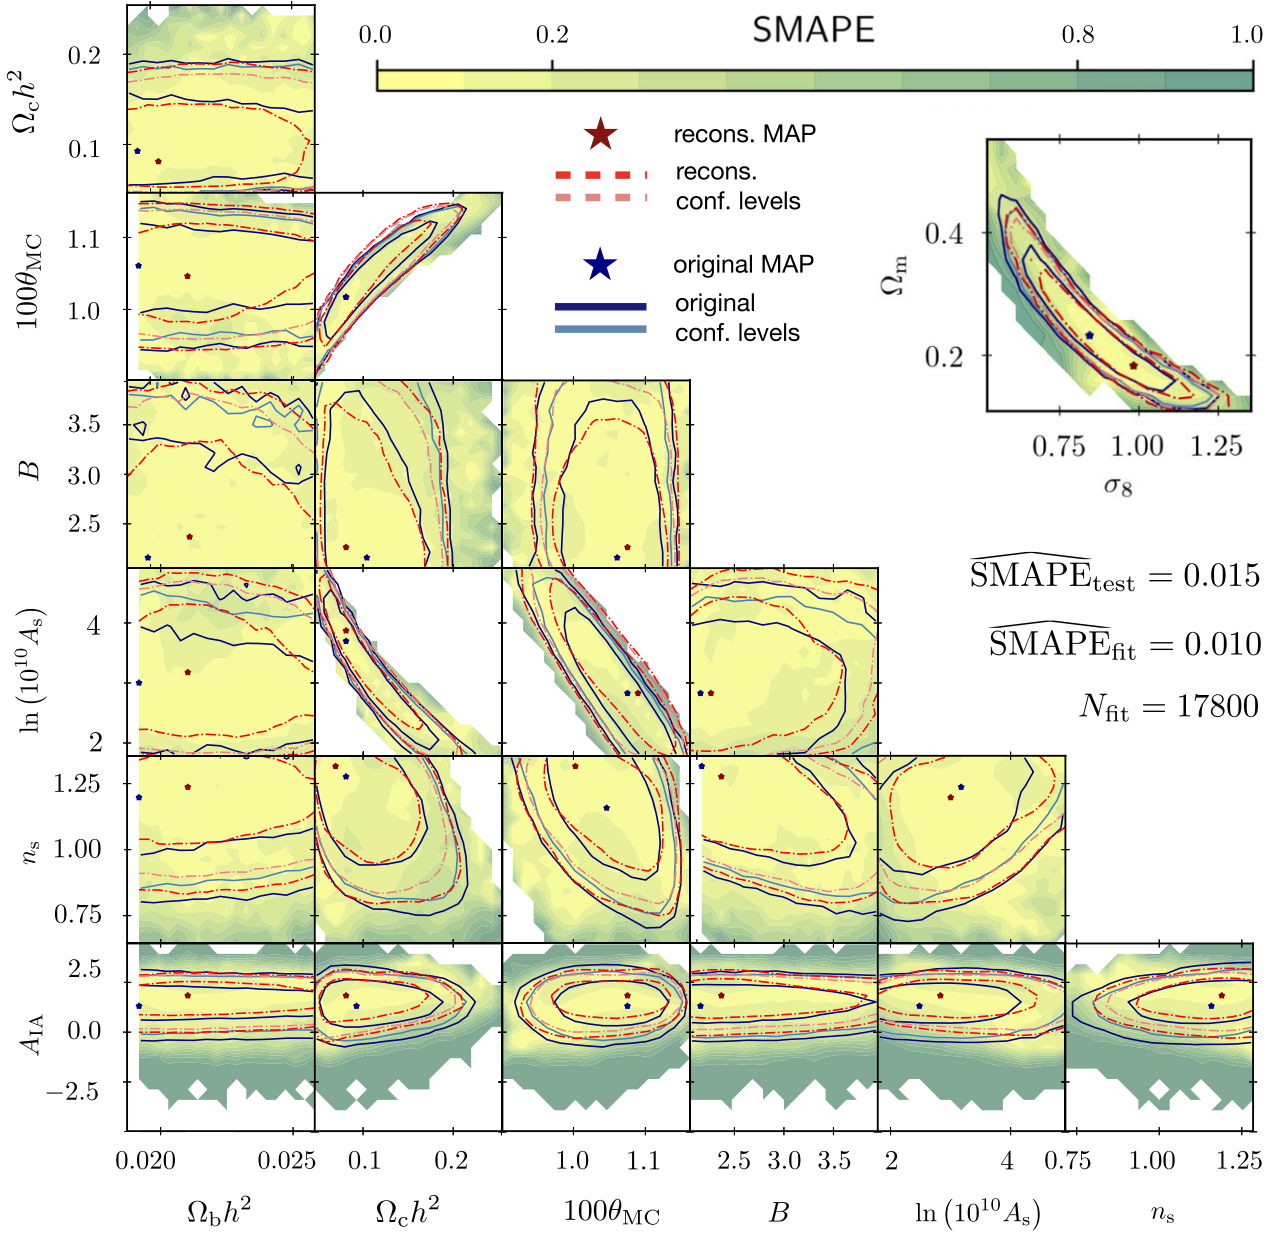
sufficiently many samples to guarantee a successful variational-Bayes posterior. In contrast, if leaving points out yields a drastically different posterior, then insufficiently many likelihood evaluations are available. If possible, more samples should then be added in the area where differences in the estimated posteriors are seen. In Fig. 4, we employ the Leave-N-out verification to validate the reconstruction of the marginal posterior  $\{\sigma_8, \Omega_m\}$  in Fig. 1. For this test, 10 random test sets of 10 points were obtained from an initial set of 55 points. 45 being points sufficient for the optimization, we obtained 10 extremely similar reconstructions, four of which are proposed in Fig. 4.

Throughout our work, we relied on the global SMAPE as metric for evaluating the divergence between reconstructed and original

posterior. This is a general purpose metric and for future work we advise to relate it to relevant summary statistics such as the mean and standard deviation of the projected marginal distributions. Furthermore, follow up studies should aim at generalizing the preliminary analysis carried out in this section to identify the minimum number of samples which ensures satisfactory reconstruction performance regardless the distribution of the fitting points.

## 5.2 Conclusions

We have developed a variational-Bayes solution building upon the work of S14 and S15, with the aim of constructing posteriors when MCMC cannot be run, either due to run times (extremely expensive




**Figure 3.** Bottom left: Corner plot showing the two-dimensional marginals for all the primary parameters considered in the original KiDS-450 tomographic weak gravitational lensing analysis. Top right: Marginal distribution for the derived parameters  $\{\sigma_8, \Omega_m\}$ . In each subplot, we compare the marginal distribution from the full MCMC in light blue (solid line) and our variational-Bayes solution in red (dot-dashed line). The variational-Bayes solution approximates the full seven-dimensional posterior, although only its marginals are visible here. The number of shape parameters is  $N_{\text{shape}} \approx 600$ . The number of posterior samples employed for fitting is  $N_{\text{fit}} = 17\,800$ . We also specified the posterior to be zero at 8900 further positions in order to deal with the curse of dimensionality. This means that we set the posterior to zero in the vast and empty corners of the seven-dimensional space. The variational-Bayes solution accordingly only required 0.6 per cent of the MCMC’s computational cost.

likelihoods), or because the physical model can only be evaluated at discrete locations in parameter space (simulation-based inference).

A variational-Bayes approach fits a so-called variational distribution to a set of posterior samples, thereby determining suitable values for the tunable shape-parameters of the variational distribution. The

best-fitting variational distribution is taken as representation of the posterior.

We applied our algorithm on the KiDS-450 analysis (Hildebrandt et al. 2017) that originally used  $\sim 3 \times 10^6$  MCMC (accepted and rejected) samples (including burn-in) to construct their posteriors.



**Figure 4.** Illustration of the Leave-N-out test to judge the quality of the variational-Bayes posterior. Out of an initial set of 55 points (black crosses), each panel of the first row selected randomly  $N_{\text{fit}} = 45$  points, with the discarded 10 points being greyed out. As the posterior shape barely varies between the panels, using all 55 points for a joint reconstruction will yield a high-quality posterior. In the lower row, only 22 points were randomly chosen. The Leave-N-out test then indicates that these were too few points to trust the posterior reconstruction, as each panel displays a differently shaped posterior.

In comparison, we constructed accurate 2D marginals of the KiDS posterior from 14 to 45 samples, and the joint seven-dimensional posterior from  $\approx 18\,000$  samples.

A future aim is to minimize the number of required samples even further.

## ACKNOWLEDGEMENTS

MR thanks Guilhem Lavaux for support with the GSL. We thank Simon Portegies Zwart for fruitful discussions and sharing of numerical infrastructure. We convey our gratitude to Doogesh Kodi Ramanah for constructive suggestions to improve the manuscript. This work has made use of the Horizon Cluster hosted by Institut d’Astrophysique de Paris. We thank Stephane Rouberol for maintenance and running of this computing cluster.

## DATA AVAILABILITY

The data underlying this article are available at [http://kids.strw.leidenuniv.nl/cs2016/MCMC\\_README.html](http://kids.strw.leidenuniv.nl/cs2016/MCMC_README.html)

## REFERENCES

- Abbott T. et al., 2016, *Phys. Rev. D*, 94, 022001  
 Amendola L., Gómez-Valent A., 2020, *MNRAS*, 498, 181  
 Audren B., Lesgourgues J., Bird S., Haehnelt M. G., Viel M., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 026  
 Durrer R., Maartens R., 2008, *General Relat. Gravitation*, 40, 301  
 Hagstotz S., de Salas P. F., Gariazzo S., Gerbino M., Lattanzi M., Vagnozzi S., Freese K., Pastor S., 2021, *Phys. Rev. D*, 104, 123524  
 Hildebrandt H. et al., 2017, *MNRAS*, 465, 1454  
 LSST Science Collaboration et al., 2009, preprint (arXiv:0912.0201)  
 Laureijs R. et al., 2011, preprint (arXiv:1110.3193)

- MacKay D. J. C., 2002, *Information Theory, Inference and Learning Algorithms*. Cambridge Univ. Press, Cambridge  
 Mandelbaum R., Miyatake H., Hamana T., Oguri M., Simet M., Utsumi Y., 2018, *Publ. Astron. Soc. Japan*, 70, S25  
 Sellentin E., 2015, *MNRAS*, 453, 893 (S15)  
 Sellentin E., Durrer R., 2015, *Phys. Rev. D*, 92, 063012  
 Sellentin E., Schäfer B. M., 2016, *MNRAS*, 456, 1645  
 Sellentin E., Quartin M., Amendola L., 2014, *MNRAS*, 441, 1831 (S14)  
 Shepard R., Brozell S. R., Gidofalvi G., 2015, *J. Phys. Chem. A*, 119, 7924  
 Tegmark M., Taylor A., Heavens A., 1997, *ApJ*, 480, 22

## APPENDIX A: NUMBER OF SHAPE PARAMETERS

The shape parameters entering the variational distribution in equation (46) are as follows:

$$\text{dof} [\hat{\boldsymbol{p}}] = d, \quad (\text{A1})$$

$$\text{dof} [\mathbf{L}^{\text{v}}] = \frac{d(d-1)}{2}, \quad (\text{A2})$$

$$\text{dof} [\mathbf{\Delta}^{\text{v}}] = d, \quad (\text{A3})$$


$$\text{dof} [\tilde{\mathbf{Q}}^{\text{v}}] = \frac{d(2s-d-1)}{2}, \quad (\text{A4})$$

$$\text{dof} [\mathbf{L}^{\text{m}}] = \frac{s(s-1)}{2}, \quad (\text{A5})$$

$$\text{dof} [\mathbf{\Delta}^{\text{m}}] = s. \quad (\text{A6})$$

Equations (A2) and (A6) result from requiring the columns of  $\mathbf{L}^{\text{v}}$  and  $\mathbf{L}^{\text{m}}$  to have unitary norm. Summing all up, the total number of shape parameters is

$$N_{\text{shape}} = \frac{1}{8} (10d + 7d^2 + 6d^3 + d^4). \quad (\text{A7})$$



**Figure A1.** Total number number of shape parameters  $N_{\text{shape}}$  as a function of the dimension  $d$  over two different ranges.

We show how equation (A7) scales with the dimension  $d$  in Fig. A1.

## APPENDIX B: QR-DECOMPOSITION OF RECTANGULAR MATRICES

We briefly describe the QR-decomposition for real-valued rectangular, rather than square matrices. The QR-decomposition factorizes an

$m \times n$  matrix  $\mathbf{A}$  with  $m > n$  into the product  $\mathbf{A} = \mathbf{Q}\mathbf{R}$ , such that the matrix  $\mathbf{Q}$  is  $m \times m$  and orthogonal, and the matrix  $\mathbf{R}$  is a rectangular  $m \times n$  upper-triangular matrix. This implies the last  $m - n$  rows of the matrix  $\mathbf{R}$  are zero. Square orthogonal matrices form a group, with the identity being  $\mathbf{I} = \mathbf{Q}^T\mathbf{Q}$ .

The effect of the last rows in  $\mathbf{R}$  being zero is that many elements of  $\mathbf{Q}$  become irrelevant, as they will be multiplied by zero. This irrelevance can be denoted more explicitly by partitioning the matrices as

$$\mathbf{Q} = (\mathbf{Q}_1 \ \mathbf{Q}_2), \quad \mathbf{R} = \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix}. \quad (\text{B1})$$

Here,  $\mathbf{R}_1$  is  $n \times n$  upper triangular, and  $\mathbf{0}$  denotes the matrix of zero rows of  $\mathbf{R}$ . The matrix  $\mathbf{Q}_2$  is  $m \times (m - n)$  and denotes the irrelevant columns of  $\mathbf{Q}$  being multiplied by zero.  $\mathbf{Q}_1$  is a rectangular  $m \times n$  matrix. We then have

$$\mathbf{A} = (\mathbf{Q}_1 \ \mathbf{Q}_2) \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix} = \mathbf{Q}_1\mathbf{R}_1. \quad (\text{B2})$$

For  $m \times n$  orthogonal matrices such as  $\mathbf{Q}_1$ , we have  $\mathbf{Q}_1^T\mathbf{Q}_1 = \mathbf{I}_{n \times n}$ . In consequence, the matrix  $\mathbf{A}^T\mathbf{A}$  will then have the decomposition  $\mathbf{A}^T\mathbf{A} = \mathbf{R}_1^T\mathbf{Q}_1^T\mathbf{Q}_1\mathbf{R}_1$ , implying that  $\mathbf{R}_1$  is the Cholesky factor of  $\mathbf{A}^T\mathbf{A}$  as the two orthogonal matrices yield the  $n \times n$  identity matrix.

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