

Data compression and covariance matrix inspection: Cosmic shearTassia Ferreira^{1,2,3}, Tianqing Zhang,¹ Nianyi Chen¹, and Scott Dodelson^{1,4}

(LSST Dark Energy Science Collaboration)

¹*Department of Physics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15312, USA*²*PPG Cosmo, Universidade Federal do Espírito Santo, Vitória, ES 29075-910, Brazil*³*Laboratório Interinstitucional de e-Astronomia—LIneA, Rio de Janeiro, RJ 20921-400, Brazil*⁴*NSF AI Planning Institute for Physics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15312, USA*

(Received 2 November 2020; accepted 11 April 2021; published 28 May 2021)

Covariance matrices are among the most difficult pieces of end-to-end cosmological analyses. In principle, for two-point functions, each component involves a four-point function, and the resulting covariance often has hundreds of thousands of elements. We investigate various compression mechanisms capable of vastly reducing the size of the covariance matrix in the context of cosmic shear statistics. This helps identify which of its parts are most crucial to parameter estimation. We start with simple compression methods, by isolating and “removing” 200 modes associated with the lowest eigenvalues, then those with the lowest signal-to-noise ratio, before moving on to more sophisticated schemes like compression at the tomographic level and, finally, with the massively optimized parameter estimation and data compression (MOPED). We find that, while most of these approaches prove useful for a few parameters of interest, like Ω_m , the simplest yield a loss of constraining power on the intrinsic alignment (IA) parameters as well as S_8 . For the case considered—cosmic shear from the first year of data from the Dark Energy Survey—only MOPED was able to replicate the original constraints in the 16-parameter space. Finally, we apply a tolerance test to the elements of the compressed covariance matrix obtained with MOPED and confirm that the IA parameter A_{IA} is the most susceptible to inaccuracies in the covariance matrix.

DOI: [10.1103/PhysRevD.103.103535](https://doi.org/10.1103/PhysRevD.103.103535)**I. INTRODUCTION**

Cosmic shear is a weak lensing effect caused by the large-scale structure of the Universe and is an important tool for constraining cosmology. The most common way of obtaining information from cosmic shear is to use two-point functions and, as is often the case, this analysis assumes that the summary statistics have a Gaussian distribution, thus requiring a covariance matrix. For a two-point data vector of length N , the covariance matrix is a symmetric $N \times N$ matrix with $N \times (N + 1)/2$ individual elements that capture the auto and cross-correlation of the data vector. As the length of the data vector increases, the number of elements in the covariance matrix grows quadratically and becomes harder to evaluate. Compression schemes resolve this by significantly reducing the dimension of the matrix while still retaining relevant information about the parameters of interest, and also potentially speeding up computations. One way of accomplishing this is to use the massively optimized parameter estimation and data compression (MOPED), in which, if the noise in the data does not depend on the model parameters, then the Fisher matrix for both the full and compressed covariance

matrices coincides, and the compression is said to be lossless [1,2]. MOPED has been widely used in literature for a variety of applications, for example, CMB data [3], the redshift-space galaxy power spectrum and bispectrum [4], parameter-dependent covariance matrices [5], compression of the Planck 2015 temperature likelihood [6], weak lensing and galaxy clustering [7], and has been paired with a Gaussian process emulator to analyze weak lensing data [8].

We will focus on cosmic shear measurements from the Dark Energy Survey year 1 (DES Y1) release [9,10]; the data vector has 227 elements, varying with angular separation and different pairs of tomographic redshift bins. Since our parameter space consists of 16 free parameters, we can use MOPED to reduce the 25,878 independent elements of the covariance matrix, to only 136.

Apart from MOPED, we will be analyzing the covariance matrix with three other compression techniques: the first involves performing an eigenmode decomposition then discarding the modes associated with the lowest eigenvalues; the second approach removes those with the lowest signal-to-noise ratio. In order to obtain a compression competitive with MOPED in terms of shrinkage, i.e., about

10% of the original size, we remove, in both cases, 200 such modes.

Finally, the third method consists of a map-level compression [11], where linear combinations of the tomographic maps are used to retain as much information as possible. Compression of the tomographic bin pairs then considerably reduces the size of the data vector of the two-point functions. For example, we will see that most of the information in the four tomographic bins used by DESY1 can be compressed into a single linear combination of those bins, or one Karhunen-Loève (KL) mode. Therefore, instead of $(4 \times 5)/2$ two-point functions for each angular bin, we need include only one or two. For this purpose, the data vector for each tomographic bin will have the same length, and so the angular cuts to the dataset and covariance matrix will be different from the ones used in the aforementioned DESY1 paper. The chosen covariance matrix has a dimension of 190×190 . With one KL mode, we can compress the shear data vector down to 10% of its original size, yielding 190 independent elements for the covariance matrix of the new data vector.

A number of codes compute covariance matrices analytically; here we test two, which, despite being different, produce compatible constraints. It must follow that the parts most relevant to parameter estimation are similar to each other. Establishing which regions of the covariance matrix are most informative is the first step towards building a tool for comparing these matrices without the need for a full cosmological analysis. In this work, we use compression schemes to look for an indication of which elements should be considered for covariance matrix comparison.

In Sec. II, we start by describing the dataset and the covariance matrices used. We then proceed to review each compression scheme and apply them to a DESY1 cosmic shear, demonstrating how well they reproduce the constraints obtained with the full covariance matrix. We follow this by showing that compression can be a helpful tool to compare two different covariance matrices, in Sec. III. Our tolerance test is described in Sec. IV, where we investigate the change in parameter constraints resulting from the addition of noise separately to elements and eigenvalues of the covariance matrix. Finally, our conclusions are summarized in Sec. V.

II. METHODS

A. DES cosmic shear: Data and analysis

In this section, we introduce the data and covariance matrices that are used in this work. Our tests are carried out using cosmic shear statistics $\xi_{\pm}(\theta)$, focusing on the results from DESY1. The data are divided into four tomographic redshift bins spanning the interval $0.20 < z < 1.30$, which yields ten bin-pair combinations, each one containing 20 angular bins between 2.5 and 250 arcmin. We thus begin

with 200 data points per statistic, giving 400 in total. We then apply the angular cuts described in [10], which remove the scales most sensitive to baryonic effects; this leaves 167 points for $\xi_{+}(\theta)$ and 60 for $\xi_{-}(\theta)$, resulting in 227 data points corresponding to the aforementioned 227×227 covariance matrix.

Table I shows the 16-parameters varied and the priors placed on them. Since cosmic shear is not sensitive to most of these, their constraints are largely dominated by the priors used. As such, throughout, we will only be showing constraints on three of those: the matter density parameter, Ω_m , the amplitude of matter fluctuations, $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$, and the amplitude of the intrinsic alignment, A_{IA} .

To perform cosmological parameter inference we use the CosmoSIS [12–19] pipeline, while employing the MultiNest [20] sampler to explore the parameter space, with 1000 livepoints, efficiency set to 0.05, tolerance to 0.1 and constant_efficiency set to True.

The covariance matrices are the following:

- (i) the full covariance matrix (FCM) used in the DESY1 analysis, which includes non-Gaussian effects and supersample variance; generated by Cosmo-like [21];
- (ii) one containing only the Gaussian part, which we will refer to as the Gaussian covariance matrix

TABLE I. List of the priors used in the analysis for parameter constraints (\mathcal{U} denotes flat in the given range and \mathcal{G} is Gaussian with mean equal to its first argument and dispersion equal to its second). For the cosmological parameters, we fix $w = -1.0$, $\Omega_k = 0.0$ and $\tau = 0.08$. The astrophysical parameters are associated with the intrinsic alignment, they follow the relation $A_{IA}(z) = A_{IA}[(1+z)/1.62]^\eta$. Lastly, for systematics we have m^i corresponding to the shear calibration and Δz^i for the source photo- z shift, with $i = [1, 4]$ in both cases.

Parameter	Prior
Cosmological	
Ω_m	$\mathcal{U}(0.1, 0.9)$
$\log A_s$	$\mathcal{U}(3.0, 3.1)$
H_0 (km s ⁻¹ Mpc ⁻¹)	$\mathcal{U}(55, 91)$
Ω_b	$\mathcal{U}(0.03, 0.07)$
$\Omega_c h^2$	$\mathcal{U}(0.0005, 0.01)$
n_s	$\mathcal{U}(0.87, 1.07)$
Astrophysical	
A_{IA}	$\mathcal{U}(-5, 5)$
η_{IA}	$\mathcal{U}(-5, 5)$
Systematic	
m^i	$\mathcal{G}(0.012, 0.023)$
Δz^1	$\mathcal{G}(-0.001, 0.016)$
Δz^2	$\mathcal{G}(-0.019, 0.013)$
Δz^3	$\mathcal{G}(0.009, 0.011)$
Δz^4	$\mathcal{G}(-0.018, 0.022)$

(GCM); generated by the same code used to analyze the KiDS-450 survey [22,23].

Thus, throughout, the covariance labels FCM and GCM differ for several reasons: first, they are two independent codes and, second, although the code for the KiDS-450 survey does contain all the functionality in *Cosmolike*, we ran the GCM with the simplest settings in order to accentuate the differences. The ensuing discrepancies help us assess various validation techniques. Where not otherwise stated, the analysis and constraints will be performed on the FCM.

Figure 1 shows the projected cosmological constraints for the FCM and the GCM, using the same data vector and angular cuts. The 68% confidence level (CL) constraints are as follows: for the FCM: $\Omega_m = 0.306^{+0.018}_{-0.023}$, $S_8 = 0.784^{+0.054}_{-0.06}$ and $A_{IA} = 0.852^{+0.359}_{-0.233}$; and for the GCM: $\Omega_m = 0.309^{+0.017}_{-0.023}$, $S_8 = 0.787^{+0.051}_{-0.058}$ and $A_{IA} = 0.948^{+0.329}_{-0.22}$. This shows that the variations we introduced to the calculation of the two matrices are measurable in the parameter constraints.

B. Eigenvalues

Let us start with the easy task of analyzing the eigenvalues of the covariance matrix. Each eigenvalue is associated with a linear combination of the data vector, or a *mode*.

The idea is to remove the contribution of the lowest eigenvalues, since these are usually attributed to numerical noise and, as such, contain the least amount of information.

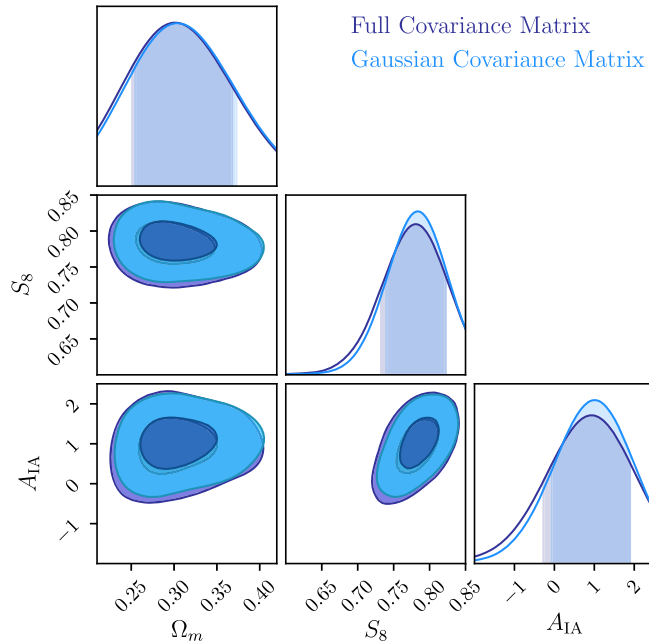


FIG. 1. Constraints on cosmological parameters Ω_m and S_8 and intrinsic alignment parameter A_{IA} for two covariance matrices produced for cosmic shear. The purple curve is for the FCM while the blue is for the GCM. In the 16-dimensional parameter space, the volume of the posterior is about 22% larger for the former.

The highest eigenvalues, on the other hand, are said to be the most informative [24]. The procedure is simple, we first diagonalize the covariance matrix in order to calculate its eigenvalues then sort them into increasing order. Setting the lowest eigenvalues to zero would result in a nonpositive definite (NPD) matrix, so we replace them instead with lower values (9 orders of magnitude lower), thus removing their effective contribution; we then transform back to the original basis and perform a cosmological analysis with the new covariance matrix.

For the purpose of reducing the covariance matrix to about 10% of its original size, we follow the procedure above to discard the 200 eigenmodes with the lowest eigenvalues. The results reported in Fig. 2 show a loss of constraining power on two of the three parameters shown. This is consistent with the fact that we are removing about 90% of the information contained in the covariance matrices. However, it is inconsistent with the conjecture that the modes with lowest eigenvalues are irrelevant, in fact, constraints on S_8 for the FCM are $0.779^{+0.044}_{-0.46}$, whereas, for the new covariance matrix, we obtain $0.725^{+0.076}_{-0.083}$, showing an increase in the errors of almost 77%. It is then clear that this method is incompatible with a 10% reduction, and so we must look for a different way of ordering the modes.

C. Signal-to-noise ratio

Instead of looking only at the “noise”—or the eigenvalues of the covariance matrix—a better way to assess




FIG. 2. Constraints on cosmological parameters Ω_m , S_8 and the intrinsic alignment parameter A_{IA} for the original covariance matrix (in purple) and for the two new covariance matrices obtained in Sec. II B (in blue) and Sec. II C (in magenta).

the importance of modes is to consider the signal as well. We can define the expected signal-to-noise ratio (SNR) as

$$\left(\frac{S}{N}\right)^2 = T_i C_{ij}^{-1} T_j, \quad (1)$$

where T_i is the predicted theoretical signal for the i^{th} data point, given a fiducial cosmology, and C_{ij} is the covariance matrix. Repeated indices are summed in all cases, throughout this work. If C_{ij} were diagonal, then the eigenvectors would simply be the T_i s themselves, and not a linear combination of them, and we could estimate the SNR squared expected in each mode by just computing T_i^2/C_{ii} , with ii denoting the diagonal element i . Then we could throw out the modes with the lowest SNR. Since this is not the case here, we have to first diagonalize C_{ij} and then order the values. We write the expected SNR squared as

$$\left(\frac{S}{N}\right)^2 = \frac{v_i^2}{\lambda_i}, \quad (2)$$

where λ_i are the eigenvalues of the covariance matrix, which is diagonalized with the unitary matrix U , and the eigenvectors are

$$v_i \equiv U_{ij}^T T_j, \quad (3)$$

with the superscript T denoting the transpose. From a naive point of view, this makes it clear which modes should be kept and which should be dropped; modes v_i for which $(v^2/\lambda)_i$ is small can be discarded. As we will later see, however, it is not as simple as that.

After obtaining the SNR for the covariance matrix, we reduce the 200 lowest values to 7 orders of magnitude lower, which is equivalent to increasing the noise (or decreasing the signal) of these modes. We then obtain a new covariance matrix with the corresponding modified

SNR values. The parameter constraints for this method are shown in Fig. 2, where we note that only Ω_m is well constrained (in agreement with those obtained with the original covariance matrix to within a 2σ interval). The constraining power on A_{IA} and S_8 , on the other hand, is weakened, which suggests that the modes removed do indeed carry relevant information for these parameters.

We can investigate this loss by tweaking our understanding of which modes carry information. The “signal” that these modes are ordered by is the amplitude of the data points. The parameters, however, are sensitive to the shape as well as the amplitude. To address this, we can identify the SNR for each parameter individually by taking

$$\left(\frac{\partial S/\partial p_\alpha}{N}\right)^2 = \frac{(\partial v_i/\partial p_\alpha)^2}{\lambda_i}, \quad (4)$$

where $\partial/\partial p_\alpha$ is the derivative with respect to each parameter. The importance of this procedure is illustrated in Fig. 3, which shows the normalized SNR for a given mode on the x -axis against the SNR for Ω_m , S_8 and A_{IA} . The shaded region shows the 200 modes excluded in the previous analysis, where we see the presence of low SNR modes that contain information about the parameters. This is particularly true for the intrinsic alignment parameter A_{IA} , which seems to explain the poor constraints shown in Fig. 2. As a result, simply cutting on raw SNR loses constraining power.

On the other hand, as [25] argues, removing the modes with the highest SNR is recommended in order to obtain a bias-free inference (another way would be to use a non-Gaussian likelihood). In light of that, we followed the same procedure used for removing the modes with the lowest SNR, but instead set the 200 highest modes to values several orders of magnitude lower. This yielded weaker constraints for not only for S_8 and A_{IA} , but also for Ω_m . We believe that this divergence was due to the large quantity of




FIG. 3. Scatter plot for the relation between the signal to noise (SNR) for each parameter (y-axis) against that for the full set of parameters (x-axis). The derivatives are shown with respect to Ω_m (blue circle), for S_8 (orange \times) and for the intrinsic alignment parameter A_{IA} (green triangle). The purple rectangle spreads until the two hundred lowest values of SNR, which corresponds to the values that are modified for parameter constraints, as described in Sec. II C.

modes removed for our analysis and does not, in any way, invalidate the findings of the aforementioned work.

D. Tomographic compression

The tomographic compression method of this section is based on a Karhunen-Lo  ve (KL) decomposition for the shear power spectrum suggested by [11] and later applied to real space two-point function in [26] for the CFHTLenS survey. Its implementation consists of finding the eigenmode—in this case, a linear combination of the convergence in different tomographic bins—with most of the signal-to-noise ratio contribution to the power spectrum, and then transforming the two-point function of this eigenmode into real space. This is not the most general compression method for the two-point function in real space, since the weight is dependent on the multipole ℓ . However, as found in [26], it is effective on the real space data, nonetheless.

Before diving into the derivation, it is worth summarizing the results. With COSMOSIS, we calculate the shear angular power spectrum C_ℓ of the convergence κ^i , where $i = [1, 4]$ for the four tomographic bins probed by DESY1 with a fiducial cosmology at the best-fit parameters. We thus have $4 \times 5/2 = 10$ pairs of bins for which we can compute spectra. The left plot in Fig. 4 shows the diagonal elements of the signal part, S_ℓ , and of the noise part, N_ℓ , of the spectrum. The right-hand panel shows the signal-to-noise ratio for the KL-transformed eigenmodes, which we call D_ℓ , ranging from $\ell = 10$ to $\ell = 2500$. That is, we identify a mode as $b_{\ell m} = r_i \kappa_{\ell m}^i$, where r_i is the weight factor on the i^{th} tomographic bin. We can see that the first KL mode contains most of the SNR contribution to the power spectrum. However, if we want to recover more information, we also should include the second and the cross mode between the first and second KL-mode.

With the total power spectrum $C_\ell = S_\ell + N_\ell$, we calculate the Karhunen-Lo  ve (KL) modes for each ℓ (so we drop the ℓ subscript) via a general eigenvalue problem,

$$C^{ij} e_p = \lambda_p N^{ij} e_p. \quad (5)$$

The index p in e_p corresponds to the p^{th} KL-mode of C . Using Cholesky decomposition, $N = LL^T$, we express the new observable as $b_p = e_p L^{-1} \kappa$. And we find that $E_\ell = [e_1, e_2, \dots]^T$ is a transformation of basis so that the shear signal is diagonalized. We can now calculate the power spectrum D_ℓ for the new uncorrelated observable $b_{\ell m}$,

$$D_\ell = \langle b_{\ell m} b_{\ell m}^T \rangle = E_\ell L^{-1} C_\ell L^{-1} E_\ell^T = \Lambda_\ell, \quad (6)$$

where $\Lambda_\ell = \text{diag}[\lambda_1, \lambda_2, \dots]$. If we denote $E_\ell N^{-1}$ as R_ℓ and further write $U_\ell^{ij} = R_\ell^i R_\ell^j$, where i and j are the indices for the tomographic bin-pairs, we have the compression in terms of one simple linear combination,

$$D_\ell = R_\ell^i C_\ell^{ij} R_\ell^j = U_\ell^{ij} C_\ell^{ij}, \quad (7)$$

with U_ℓ^{ij} being the weight we will use to compress the two-point functions. We note that these KL-modes $b_{\ell m}^p$ are uncorrelated, so that their power spectrum $D_\ell^{pp'}$ is a diagonal matrix whose entries are $1 + \text{SNR}$ of the corresponding eigenmodes. This allows us to compress ten tomographic bin-pairs to one, or two, by taking only the modes with the highest SNR.

We want, however, to eventually compress the two-point function data vector of DESY1, which is measured in the real space tomographic bin pair ij and related to the angular power spectrum C_ℓ via




FIG. 4. *Left*: Shear power spectrum of the FCM. Solid lines are diagonal elements of the signal matrix S_ℓ , and dashed lines are the diagonal elements of noise matrix N_ℓ . *Right*: Signal-to-noise ratio matrix D_ℓ of the first to fourth KL-modes of the power spectrum on the left panel.

$$\begin{aligned}\xi_{+}^{ij}(\theta) &= \int \frac{\ell d\ell}{2\pi} J_0(\ell\theta) \mathcal{C}^{ij}(\ell), \\ \xi_{-}^{ij}(\theta) &= \int \frac{\ell d\ell}{2\pi} J_4(\ell\theta) \mathcal{C}^{ij}(\ell).\end{aligned}\quad (8)$$

In order to use linear combinations of all the tomographic bins, we need to ensure that the combination is ℓ -independent, that is to say, the transformed two-point correlation function, $\tilde{\xi}_{\pm}(\theta)$, can be directly calculated from other two-point functions. In fact, Fig. 5 shows that the $U^{ij}(\ell)$ are generally ℓ -independent, except for low ℓ s, due to the existence of cosmic variance. Therefore, we have

$$\begin{aligned}\tilde{\xi}_{\pm}(\theta) &= \int \frac{\ell d\ell}{2\pi} J_{0/4}(\ell\theta) D(\ell) \\ &= \int \frac{\ell d\ell}{2\pi} J_{0/4}(\ell\theta) U_{\ell}^{ij} \mathcal{C}^{ij}(\ell) \\ &= \bar{U}^{ij} \xi_{\pm}^{ij}(\theta),\end{aligned}\quad (9)$$

where \bar{U}^{ij} is the average U_{ℓ}^{ij} given by

$$\bar{U}^{ij} = \frac{\int_{\ell_{\min}}^{\ell_{\max}} d\ell (2\ell + 1) U_{\ell}^{ij}}{\int_{\ell_{\min}}^{\ell_{\max}} d\ell (2\ell + 1)}.\quad (10)$$

We make a more conservative angular cut than the one discussed in [9], making sure that both $\xi_{\pm}(\theta)$ are uniform in regard to tomographic combinations. We consider an angular scale for ξ_{+} from $7.195'$ to $250.0'$, and for ξ_{-} from $90.579'$ to $250.0'$. Therefore, for the purpose of exploring the KL-transform, the raw data vector has a length of 190. By shrinking ten tomographic combinations for each angle into 1 KL-mode, the data vector is reduced to length 19, and so the number of elements in the covariance matrices has a compression of 99%.

In Fig. 5, we plot the normalized KL-eigenmodes $e_{\ell}^p N^{-\frac{1}{2}}$ and their corresponding weights, $U_{\ell}^{ij} = R_{\ell}^i R_{\ell}^j$. Modes with increasing ℓ are plotted in increasing opacity of the color. While the KL-modes do vary by a slight amount for different ℓ , their sensitivity to it is not very significant since they converge for higher ℓ to their weighted average, which we represent with the dashed black lines. For the first KL-mode, the tomographic bins with higher redshift are weighted more than those with low redshift. This is also shown in the right panel by the weight on tomographic combination that the combination of bin 3 and bin 4 carries most of the weight in the signal-to-noise ratio. This agrees with the fact that low-redshift galaxies are less affected by lensing than high-redshift galaxies, as indicated in the left panel of Fig. 4.

We ran the likelihood analysis as detailed in Sec. IIA with the first KL-mode and the first two KL-modes with their cross-correlation mode, which correspond to a 10-to-1 and 10-to-3 compression, respectively, and show the parameter constraints on the $\Omega_m - S_8 - A_{\text{IA}}$ plane in Fig. 6. We do not include the third and fourth KL mode because they contain considerably less signal to noise. We can see that the first KL-mode is generally not sufficient to recover the information in the data vector. Since the first two modes contain most of the SNR contribution at a map level, we were able to recover the Ω_m constraints. However, information about the $S_8 - A_{\text{IA}}$ combination is clearly lost. This could be due to the fact that the SNR-prioritized modes are not the sensitive direction for these parameters, as was also the case in Fig. 3. Indeed, the $S_8 - A_{\text{IA}}$ plane shows a strong correlation between these two parameters. This likely explains why the constraints for S_8 widened: the KL-modes fail to break the degeneracy on A_{IA} , which is mostly present in the modes that are insensitive to cosmic shear and are discarded in the compression process.




FIG. 5. *Left*: Column vectors of the matrix R_{ℓ} , or $e_{\ell}^p N^{-\frac{1}{2}}$, for compressing the shear power spectrum \mathcal{C}_{ℓ} . *Right*: Transformation on tomographic bin combination U_{ij} constructed from the KL-eigenmodes. For both plots, the dashed black lines are the weighted average of each mode. The lightest shade represents $\ell = 10$ and the increment is $\Delta\ell = 10$ for each darker shade.




FIG. 6. Cosmological constraints marginalized over all 16 parameters for the 190×190 FCM and that compressed using the first KL-mode and the first two KL-modes.

E. Applying MOPED

The MOPED compression scheme takes place at the two-point level, with the compressed data vector containing linear combinations of the many two-point functions. In principle, this requires only N_p linear combinations of the two-point functions, where N_p is the number of free parameters, and each mode, or linear combination, contains all the information necessary about the parameter of interest.

For each parameter p_α that is varied, one captures a single linear mode,

$$y_\alpha = U_{ai} D_i, \quad (11)$$

where D_i are the data points, and the coefficients are defined as

$$U_{ai} \equiv \frac{\partial T_j}{\partial p_\alpha} C^{-1}_{ji}, \quad (12)$$

with T_j being the theoretical prediction for the data point D_j for a fiducial cosmology. An illustration of the matrix U_{ai} is shown in Fig. 7, showing the weighting vectors for parameters Ω_m , S_8 and A_{IA} .

The now much smaller dataset $\{y_\alpha\}$, which contains N_p data points, carries its own covariance matrix, from which χ^2 can be computed for each point in parameter space. Propagating through shows that this covariance matrix is related to the original C_{ij} via

$$C_{\alpha\beta} = U_{ai} C_{ij} U_{j\beta}, \quad (13)$$

which also happens to be identical to the Fisher matrix of our likelihood. This compression was first suggested by Tegmark *et al.* (1997) for a single parameter only. The nontrivial extension to multiple parameters, where the full Fisher matrix is reproduced with the compressed data, is the MOPED algorithm [1]. One difference here is that our weighing vector given by Eq. (12) does not carry the normalizing factor of Eq. (11) in [1]. In our case, the covariance matrix is 227×227 , while the number of parameters needed to specify the model is only 16, so $C_{\alpha\beta}$ is a 16×16 matrix. We have apparently capture, from the initial set of $(227 \times 228)/2 = 25,878$ independent elements of the covariance matrix, a small subset (only 136) of linear combinations that really matter. If two covariance matrices give the same set of $C_{\alpha\beta}$, it should not matter whether any of the other thousands of elements differ from one another.

Ultimately, what matters is how well the likelihood does at extracting parameter constraints. Since most analyses assume a Gaussian likelihood, this boils down to how well the contours in parameter space agree when computing χ^2 using two different covariance matrices.

Figure 8 compares the constraints obtained for the compressed covariance matrix and dataset with results from the full one. The two curves agree extremely well for the parameters shown: Ω_m , S_8 and A_{IA} . This is also true for all the other cosmological and intrinsic alignment




FIG. 7. An illustration of the 227 values of the weights corresponding to Ω_m , S_8 and A_{IA} used for compressing the covariance matrices. Note the similarity of the weighting vectors for S_8 and A_{IA} , and that the largest values correspond to the last 60 elements, i.e., those that we will use to compress the part of the covariance matrix that holds information for ξ_- .




FIG. 8. Constraints on cosmological parameters Ω_m and S_8 and for the intrinsic alignment parameter A_{IA} for the original covariance matrix, FCM, (in purple) and for the compressed one (in blue).

parameters, where their mean values agree at the 1σ confidence level. While the volume of the whole constrained parameter space does increase by about 13%, the constraints for most parameters are less than 4% broader, which shows that the information loss is negligible.

III. COMPARISON OF COVARIANCE MATRICES

Armed with this information about compression, we now set out to compare the two covariance matrices, the GCM and the FCM, described in Sec. IIA.

A. Element-by-element comparison

We begin by performing an element-by-element comparison between the two covariance matrices. If there were only a single data point, then the covariance matrix would be one number and comparing two covariance matrices to try to understand why they give different constraints would be as simple as comparing these two numbers. The simplest generalization is then to do an element-by-element comparison. We make a scatter plot of the elements of the two matrices in the bottom panel of Fig. 9, where we can see that the elements of the FCM are, in general, larger than the GCM's, with many of the off-diagonal elements differing by 2 orders of magnitude or more. In some ways, this is useful and reassuring, as it aligns with what we see in the parameter constraints, in Fig. 1: larger elements in the covariance matrix translates to less constraining power.

The limitation of this method is that it remains unclear which of the differences are driving the final discrepancies




FIG. 9. In both plots, the red points refer to the diagonal elements, and the color bar varies according to the number of elements in one hexagonal bin, where the darkest blue color corresponds to only one element, and the brightest yellow shade to 2000. *Top*: Scatter plot of the ratio of the elements of the GCM and the FCM vs the FCM value. For illustrative purposes, we draw a black, horizontal line at $GCM/FCM = 1$. *Bottom*: Density of the scatter plot of the positive elements of the GCM and the FCM, with the black line showing $FCM = GCM$.

in parameter constraints. This difficulty is an outgrowth of the increasing size of the datasets and hence the growing number of elements of the covariance matrix that any two codes are likely to disagree on. This element-by-element comparison, however, would prove much more practical if we fewer elements to compare. Towards that end, we turn to compressed covariance matrices.

B. Compressed matrices comparison

Since we have shown that, out of all the compression schemes shown here, the only one capable of reproducing the original parameter constraints was MOPED, that is what we will be using in this section.




FIG. 10. Results for the covariance matrices compressed following the procedure described in Sec. IIE, with the red points corresponding to the diagonal elements. *Top*: One-to-one scatter of the ratio of the elements of the GCM and the FCM, over elements of the FCM. The black horizontal line is drawn at GCM/FCM = 1. *Bottom*: One-to-one scatter of the elements of the compressed matrices, with the black line describing FCM = GCM.

We compress both covariance matrices using the same $U_{\alpha,i}$ (we also tried using different U 's for each and obtained similar results).

Figure 10 shows a one-to-one scatter plot of the compressed elements, which, as expected, exhibits a similar behavior to that observed in Fig. 9, with the elements of the FCM being larger than those of the GCM. Here, however, the ratio of the diagonal elements is closer to 1, with a fractional difference of up to 17%, as compared to 26% with the original matrices. Perhaps even more importantly, there are much fewer points on this plot, since MOPED reduces the number of elements that need to be compared. These figures provide a greater insight into the relevant elements for parameter estimation: the dispersion is largely damped, and most of the elements are within 25% of each other, which explains what we see in the parameter constraints. Figure 11 shows the correlation matrix for the




FIG. 11. The upper right and lower left plots display the correlation matrix for the GCM and the FCM respectively, and the difference between them, Δr_{ij} , is shown on the lower right.

GCM and the FCM, and the difference between the normalized off-diagonal elements. The small differences suggest that the root of the slightly looser constraints obtained with the GCM is the larger diagonal elements of the MOPED-reduced covariance matrix. That is, a problem that initially required inspecting hundreds of thousands of elements is reduced to one involving only 16.

IV. TOLERANCE OF THE COMPRESSED MATRICES

Now that we have shown that we are indeed able to compress the covariance matrix into a much simpler and considerably smaller one, our next step is to analyze the amount of error the elements can tolerate while reproducing compatible parameter constraints.

In the next two sections we test two different ways of perturbing the covariance matrix: first we consider an error to the elements themselves, then we follow a similar procedure to study the effects of introducing an error to the eigenvalues of the compressed covariance matrix. In both cases the perturbation is drawn in the following manner: consider that we want to test the impact of an error $x\%$; this can either be an increase or a decrease in the original element, or eigenvalue, as what we care about most is not whether the parameter constraints will be larger, but rather how different they are. For this error to be random, but centered at our desired percentage, we draw a δ , for each new element/eigenvalue, from a Gaussian distribution, $\mathcal{G}(0, \frac{x}{100})$ and calculate the new value to be

$$C_{\alpha\beta}^{\text{new}} = (1 + \delta)C_{\alpha\beta}^{\text{old}}, \quad (14)$$




FIG. 12. An error plot showing the changes to the constraints for Ω_m , S_8 and A_{IA} for errors added at 5%, 10%, 15%, 25%, 30%, 35%, 40% and 45% of the original elements (in purple) and eigenvalues (in green) of the compressed covariance matrix. The blue rectangle covers the 68% CL interval obtained for the FCM, and the darker blue vertical line shows the mean value for the respective parameter.

where, for the eigenvalue, we replace $C_{\alpha\beta}$ with λ_i . This analysis is done only for the FCM, with errors ranging from 5%–45%, and for 50 realizations of the perturbed matrices.

One of the concerns that arises when modifying the covariance matrix is that the resulting one has to be positive definite (PD). For this reason, in each section we also describe the steps taken to ensure this. An intelligent way of guaranteeing PD would be to perturb the log of the covariance matrix. The issue, however, is how to introduce an error to the log matrix that would be similar to what we expect to see in the original covariance matrix. In

$$C_{\alpha\beta}^{\text{new}} = e^{(1+\delta)(\log C^{\text{old}})_{\alpha\beta}}, \quad (15)$$

the value of $C_{\alpha\beta}^{\text{new}}$ is not necessarily within $\delta\%$ of $C_{\alpha\beta}^{\text{old}}$. Introducing a 10% error, for example, in such a matrix, results in a perturbed covariance matrix with some of its elements differing by several orders of magnitude from the original one. A safer procedure would then be to perturb the log of its eigenvalues, but, since we have a section dedicated to perturbations to the eigenvalues themselves, we deemed this would be repetitive.

A. Modifying the elements

Once we generate new values for each independent element, following Eq. (14), we check for positive definiteness. Since the resulting matrix is, more often than not, not PD, we correct this by identifying the smallest negative eigenvalue and adding it to the diagonal [27]. We check that, although doing this largely increases the values of the diagonal elements, less than 40% have a standard deviation of more than twice the original perturbation.

The constraints for Ω_m , S_8 and A_{IA} are shown in Fig. 12, in purple, where the blue rectangle spans over the constraints for the unchanged compressed covariance matrix. The relative change in size for the 68% CL interval is mostly $> 10\%$ for the cosmological parameters; on the

other hand, for the intrinsic alignment parameter A_{IA} , the mean values are more than 1σ away from the original one and the loss in constraining power goes up to $\sim 30\%$.

B. Modifying the eigenvalues

Another way of introducing error to the covariance matrix is to perturb its eigenvalues. For a symmetric matrix, we have

$$C = Q\Lambda Q^{-1}, \quad (16)$$

where $\Lambda = \lambda I$, with λ being the eigenvalues and I the identity matrix; and Q is a square matrix whose columns are composed of the eigenvectors of $C_{\alpha\beta}$. The eigenvalues are then perturbed as described in Eq. (14), and the error, δ is drawn from $\mathcal{G}(0, \frac{x}{100})$, with the requirement that $|\delta| < 1$. We then have $\lambda^{\text{new}} > 0$, and thus the perturbed covariance matrix associated with these new eigenvalues is PD.

The results for this method are also plotted in Fig. 12, in green. Despite the results following the same tendency as those of the last section, we find that about 80% of the elements of the perturbed covariance matrices are within 10% of their original value.

V. CONCLUSION

In this work, we set out to explore different ways of compressing, comparing and analyzing covariance matrices, giving particular emphasis to the MOPED compression scheme. We started by looking at the parameter constraints of two 227×227 covariance matrices, the FCM and the GCM, generated for DESY1 cosmic shear measurements, and saw that, although some of their elements differed by several orders of magnitude, they generated similar constraints. It was clear, then, that not all elements contribute equally to the parameter constraints, and we needed to

employ increasingly complicated methods to try and locate the most relevant parts of the covariance matrix.

The first step was then to analyze the eigenvalues. We began with the hypothesis that modes associated with the highest eigenvalues carry most information, as such, those with the lowest eigenvalues would contribute less to parameter estimation. Using this notion to compress the covariance matrix we “removed” the lowest 200 eigenvalues, by setting them to several orders of magnitude lower. While the loss in constraining power for Ω_m was only around 20%, we saw an increase of about 77% in the size of the constraints for S_8 , and more than 100% for A_{IA} . Next, we moved on to the signal-to-noise ratio, and, using a similar procedure adopted for the eigenvalues, we “removed” the modes with the lowest SNR. The results were similar to those obtained with the eigenvalue cuts and showed us that these modes did not contribute significantly to constraining some cosmological parameters, like Ω_m , however constraints on the intrinsic alignment parameters, and even S_8 were more affected. This is consistent with the fact that the IA parameters are more sensitive to low SNR scales in cosmic shear, and it shows us that we need to look at the SNR per parameter before making any cuts, so that we do not lose important information for the parameters that we want to constrain.

The next step was to shrink the covariance matrix by applying a tomographic compression, where we decompose the shear angular power spectrum into KL modes, then we look for modes with the highest SNR and use them to compress the shear data vector. We thus go from ten tomographic bin combinations to only one or two. The resulting covariance matrix, for one mode, is then reduced from 190×190 to 19×19 or 59×59 , showing a reduction of about 99% or 91%, respectively. We show, however, that one mode is not sufficient for constraining the parameters of our model, with the results being similar to our previous tests involving SNR: the constraints for Ω_m , for example, are reproduced with the first and second KL-mode, but this is not the case for the IA parameters. Since essential information of IA parameters is contained in low SNR KL-mode, the high KL-modes failed to break the degeneracy of $A_{IA} - S_8$ correlation, resulting in wider S_8 constraints.

Finally, we applied MOPED, which uses linear combinations of the data vector. By transforming the data vector and covariance matrix with a weighting vector that is parameter dependent, we were able to reduce the 227×227 matrix to a 16×16 matrix. We show that the cosmological analysis using this compressed matrix reproduced similar constraints to the DESY1 analysis, for an uncompressed covariance matrix. We also showed a comparison of the elements of the compressed covariance matrix for the FCM and the GCM and found that the new elements show reasonable agreement, with their correlation matrices being very similar, and the diagonal elements showing a percentage difference of less than 15%.

Given these results, we successfully show that MOPED is the only compression scheme, out of the ones considered in this work, capable of capturing all the relevant information required to reproduce reliable parameter constraints for the 16 parameters of interest. It is worth noting here that compression does not automatically speed up the computation for parameter inference if the projection to form the MOPED coefficients is done at each sampled point. For speed gains, a way to compute the theoretical MOPED coefficients rapidly needs to be found, e.g., using Gaussian processes to generate the compressed theory [8].

When looking at the one-to-one element comparison of the FCM and the GCM, in Fig. 9, the region of large variance suggests that there could be considerable differences in the parameter constraints. We see, however, in Fig. 1, that this is not the case. This becomes clearer when comparing the elements of the compressed covariance matrices, where, while they do follow the same tendency as the full comparison, only a smaller portion of the elements display a greater dispersion. In this sense, one of our most important results is in the ability of using MOPED to compare different matrices.

One last step was taken to analyze the error tolerance of the compressed FCM. We adopted two ways of doing this, by introducing error taken from a Gaussian distribution for 5%–45% of the original 1) element and, 2) eigenvalue of the compressed covariance matrix. For the latter, we checked that only about 20% of elements of the resulting, perturbed, covariance matrix showed errors within the expected value, while the vast majority had only about a 10% error. In both cases, however, the results were similar: for the cosmological parameters Ω_m and S_8 , the 2σ constraints changed by about 7%, on average, while for the intrinsic alignment parameter A_{IA} , the constraints were up to 30% larger. Finally, we highlight the increasing shift, in the mean values of A_{IA} , to about 32% smaller than those obtained with the uncompressed FCM; while for the cosmological parameters this was only about 5%, in general.

ACKNOWLEDGMENTS

The authors wish to thank Sukhdeep Singh, Hung-jin Huang, Patricia Larsen, Benjamin Joachimi, Mike Jarvis, and Rossana Ruggeri for useful discussions. This paper has undergone internal review by the LSST Dark Energy Science Collaboration. The internal reviewers were David Alonso, Jonathan Blazek, and Alan F. Heavens. The D. E. S. C. acknowledges ongoing support from the Institut National de Physique Nucléaire et de Physique des Particules in France; the Science and Technology Facilities Council in the United Kingdom; and the Department of Energy, the National Science Foundation, and the LSST Corporation in the United States. DESC uses resources of the IN2P3 Computing Center (CC-IN2P3–Lyon/Villeurbanne—France) funded by the Centre

National de la Recherche Scientifique; the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231; STFC DiRAC HPC Facilities, funded by UK BIS National E-infrastructure capital grants; and the UK particle physics grid, supported by the GridPP Collaboration. This work was performed in part under DOE Contract No. DE-AC02-76SF00515. T.F. also acknowledges financial support from CAPES and

FAPES. The contributions are listed below. T.F. contributed to the manuscript, led the analysis for the eigenvalues, SNR and MOPED, as well as the comparison of the compressed covariance matrices. T.Z. contributed to the manuscript, participated, and contributed substantially to all analysis, and led the analysis for the tomographic compression. N.C. contributed to the compression code. S.D. proposed the project, the analyses, led the discussions and also contributed to writing and editing the manuscript.

-
- [1] A. F. Heavens, R. Jimenez, and O. Lahav, Massive lossless data compression and multiple parameter estimation from galaxy spectra, *Mon. Not. R. Astron. Soc.* **317**, 965 (2000).
 - [2] M. Tegmark, A. N. Taylor, and A. Heavens, Karhunen-Loève eigenvalue problems in cosmology: How should we tackle large data sets?, *Astrophys. J.* **480**, 22 (1997).
 - [3] A. Zablocki and S. Dodelson, Extreme data compression for the CMB, *Phys. Rev. D* **93**, 083525 (2016).
 - [4] D. Gualdi, M. Manera, B. Joachimi, and O. Lahav, Maximal compression of the redshift-space galaxy power spectrum and bispectrum, *Mon. Not. R. Astron. Soc.* **476**, 4045 (2018).
 - [5] A. F. Heavens, E. Sellentin, D. de Mijolla, and A. Vianello, Massive data compression for parameter-dependent covariance matrices, *Mon. Not. R. Astron. Soc.* **472**, 4244 (2017).
 - [6] H. Prince and J. Dunkley, Data compression in cosmology: A compressed likelihood for planck data, *Phys. Rev. D* **100**, 083502 (2019).
 - [7] R. Ruggeri and C. Blake, Compressing combined probes: Redshift weights for joint lensing and clustering analyses, *Mon. Not. R. Astron. Soc.* **498**, 2948 (2020).
 - [8] A. Mootoovaloo, A. F. Heavens, A. H. Jaffe, and F. Leclercq, Parameter inference for weak lensing using Gaussian processes and MOPED, *Mon. Not. R. Astron. Soc.* **497**, 2213 (2020).
 - [9] M. A. Troxel *et al.* (DES Collaboration), Dark energy survey year 1 results: Cosmological constraints from cosmic shear, *Phys. Rev. D* **98**, 043528 (2018).
 - [10] T. M. C. Abbott *et al.* (DES Collaboration), Dark energy survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing, *Phys. Rev. D* **98**, 043526 (2018).
 - [11] D. Alonso, Science-driven 3D data compression, *Mon. Not. R. Astron. Soc.* **473**, 4306 (2018).
 - [12] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski, CosmoSIS: Modular cosmological parameter estimation, *Astron. Comput.* **12**, 45 (2015).
 - [13] Antony Lewis, Anthony Challinor, and Anthony Lasenby, Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models, *Astrophys. J.* **538**, 473 (2000).
 - [14] D. Kirk, A. Rassat, O. Host, and S. Bridle, The cosmological impact of intrinsic alignment model choice for cosmic shear, *Mon. Not. R. Astron. Soc.* **424**, 1647 (2012).
 - [15] M. Kilbinger, K. Benabed, J. Guy, P. Astier, I. Tereno, L. Fu, D. Wraith, J. Coupon, Y. Mellier, C. Balland, F. R. Bouchet, T. Hamana, D. Hardin, H. J. McCracken, R. Pain, N. Regnaul, M. Schultheis, and H. Yahagi, Dark-energy constraints and correlations with systematics from CFHTLS weak lensing, SNLS supernovae Ia and WMAP5, *Astron. Astrophys.* **497**, 677 (2009).
 - [16] C. Howlett, A. Lewis, A. Hall, and A. Challinor, CMB power spectrum parameter degeneracies in the era of precision cosmology, *J. Cosmol. Astropart. Phys.* **04** (2012) 027.
 - [17] S. Bridle and L. King, Dark energy constraints from cosmic shear power spectra: Impact of intrinsic alignments on photometric redshift requirements, *New J. Phys.* **9**, 444 (2007).
 - [18] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, Revising the Halofit model for the nonlinear matter power spectrum, *Astrophys. J.* **761**, 152 (2012).
 - [19] R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman, Stable clustering, the halo model and non-linear cosmological power spectra, *Mon. Not. R. Astron. Soc.* **341**, 1311 (2003).
 - [20] F. Feroz, M. P. Hobson, and M. Bridges, MultiNest: An efficient and robust Bayesian inference tool for cosmology and particle physics, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
 - [21] E. Krause and T. Eifler, COMOLIKE—Cosmological likelihood analyses for photometric galaxy surveys, *Mon. Not. R. Astron. Soc.* **470**, 2100 (2017).
 - [22] F. Köhlinger *et al.*, KiDS-450: The tomographic weak lensing power spectrum and constraints on cosmological parameters, *Mon. Not. R. Astron. Soc.* **471**, 4412 (2017).
 - [23] B. Joachimi, C. A. Lin, M. Asgari, T. Tröster, C. H. *et al.*, KiDS-1000 methodology: Modelling and inference for joint weak gravitational lensing and spectroscopic galaxy clustering analysis, *Astron. Astrophys.* **646**, A129 (2020).

-
- [24] M. S. Vogeley and A. S. Szalay, Eigenmode analysis of galaxy redshift surveys. I. Theory and methods, *Astrophys. J.* **465**, 34 (1996).
- [25] A. J. Louca and E. Sellentin, The impact of signal-to-noise, redshift, and angular range on the bias of weak lensing 2-point functions, *Open J. Astrophys.* **3** (2020).
- [26] E. Bellini, D. Alonso, S. Joudaki, and L. V. Waerbeke, Shear shear: Weak lensing with one mode, *Open J. Astrophys.* **2** (2019).
- [27] Ke-Hai Yuan and Wai Chan, Structural equation modelling with near singular covariance matrices, *Computat. Stat. Data Anal.* **52**, 4842 (2008).