

Extreme data compression while searching for new physics

Alan F. Heavens^{1*}, Elena Sellentin², and Andrew H. Jaffe¹

¹Imperial Centre for Inference and Cosmology (ICIC), Imperial College, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, U.K.

²Leiden Observatory, Leiden University, Huygens Laboratory, Niels Bohrweg 2, NL-2333 CA Leiden, The Netherlands.

Accepted 17 August 2020. Received 4 August 2020; in original form 15 June 2020

ABSTRACT

Bringing a high-dimensional dataset into science-ready shape is a formidable challenge that often necessitates data compression. Compression has accordingly become a key consideration for contemporary cosmology, affecting public data releases, and reanalyses searching for new physics. However, data compression optimized for a particular model can suppress signs of new physics, or even remove them altogether. We therefore provide a solution for exploring new physics *during* data compression. In particular, we store additional agnostic compressed data points, selected to enable precise constraints of non-standard physics at a later date. Our procedure is based on the maximal compression of the MOPED algorithm, which optimally filters the data with respect to a baseline model. We select additional filters, based on a generalised principal component analysis, which are carefully constructed to scout for new physics at high precision and speed. We refer to the augmented set of filters as MOPED-PC. They enable an analytic computation of Bayesian evidences that may indicate the presence of new physics, and fast analytic estimates of best-fitting parameters when adopting a specific non-standard theory, without further expensive MCMC analysis. As there may be large numbers of non-standard theories, the speed of the method becomes essential. Should no new physics be found, then our approach preserves the precision of the standard parameters. As a result, we achieve very rapid and maximally precise constraints of standard and non-standard physics, with a technique that scales well to large dimensional datasets.

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

1 INTRODUCTION

It may seem obvious that more data will lead to better constraints on the free parameters of a physical theory from a data set. In reality, however, there may be considerable scope to reduce the size of the dataset without compromising precision on the parameters. Data compression can be a valuable tool, especially in the cosmological context of large datasets, when relatively few parameters are to be inferred, and may be necessary, for likelihood-free inference (Alsing et al. 2019; Leclercq 2018) or when data covariance matrices have to be estimated numerically (Heavens et al. 2017).

The possibility of radical data compression is easily evidenced by Gaussian-distributed data points whose expectation value μ only is to be inferred: if there are N data points, $x_1, \dots, x_N$, each drawn from the same Gaussian distribution, then these can be compressed to a single datum $\bar{x} = \sum_i x_i/N$, which contains the same information on μ as the original N points. Data compression which conserves the data's information content may be said to be optimal or lossless, and lossy if information is lost during the compression¹. The MOPED algorithm (Heavens, Jimenez & Lahav 2000)

draws on a physical theory to induce such a lossless data compression. It uses the theory to radically compress the original data from their original size down to the number of parameters of the physical theory.

However, what if the physical theory used by MOPED to compress the data is not the correct one? It then runs the risk of ‘compressing away’ the information needed to study extended theories. In this paper we therefore tackle the challenge of inferring physics from compressed data sets whose original uncompressed data are much larger than the number of free parameters of the theory. We target jointly the most precise constraints on the physical parameters as possible, alongside studying whether the entire theory is supported by the data or not. It might well be possible that for no value of the free parameters does the theory fit the data satisfactorily, and this may motivate the search for alternative theories. It is important to realize that the stage of data compression might be the last instance where a challenge to the theory is possible: the compression may suppress, or even remove, signals of new physics, as it only strives to preserve the information content of the standard theory.

Additionally, a computational bottleneck arises: if a theory is abandoned, a replacement is needed to explain the data. If there exist no compelling theoretical reasons which single out an unambiguous replacement theory, then iterative testing over a huge theory space begins. This has, e.g., been evidenced by the track record

¹ Lossless here has a restricted meaning of no loss of precision in the parameters, and is a much weaker demand than requiring that the original dataset can be reconstructed from the compressed data.

of cosmology, where a discrepancy between the posteriors of the Planck satellite (Planck Collaboration 2018) and CFHTLenS (Leauthaud et al. 2017) triggered a decade-lasting series of proposed theory amendments, which continues to last into the era of the Kilo Degree Survey (KiDS; Hildebrandt, et al. 2020) and the Dark Energy Survey (DES; Abbott et al. 2018).

With the launch of ESA’s Euclid satellite (Amendola et al. 2018), and the ‘Legacy Survey of Space and Time’ (LSST; LSST Science Collaboration, et al. 2009) starting operations in the near future, a continuing trend of proposing and testing non-standard theories is foreseeable. In light of the rapidly increasing data size, non-strategic testing of theories is however rapidly becoming numerically prohibitive. To address this, we provide a principled, largely analytical, strategy of testing whether a data set shows any evidence of being incompatible with a standard theory, while also compressing it such that it maintains its information content, and facilitating inference of standard and non-standard physics. Both the analytical solutions and the compression are anticipated to facilitate reanalyses of Euclid- and LSST-sized datasets upon their public releases.

We base our study on the MOPED (‘Massive Optimised Parameter Estimation and Data compression’) extreme data compression algorithm, published in Heavens, Jimenez & Lahav (2000) and Reichardt, Jimenez & Heavens (2001). MOPED is a linear compression algorithm, which sifts the initial data through a set of filters, whereupon the filters output compressed data points. MOPED constructs these filters by assuming the data are well fitted by a parametric theory with M free parameters. It compresses from initially N data points down to M data points, i.e. one datum per parameter. It is optimal in the sense that the Fisher matrix of the compressed data is identical to the Fisher matrix of the original dataset in the case that 1) the data are Gaussian-distributed, 2) the information comes from the expectation value of the data, not its covariance, 3) the fiducial parameter set (used to define the filter functions) is the correct set of values. Condition (3) is clearly not possible to guarantee a priori, but one can perform a sub-optimal analysis based on an initial guess, determine maximum-likelihood parameters, and compress again with these as the fiducial model, iterating as needed. Practice has found extremely rapid convergence, and often no iteration at all is required (Heavens, Jimenez & Lahav 2000).

This approach is enormously powerful if the theory used for compression is the correct model for the data. The purpose of this paper is to include an investigation of whether the data to be compressed actually indicate any evidence of disfavoured the compressing theory. If the data show no sign of disfavoured the theory, then MOPED is complete.

What happens if the theory requires extension? MOPED by design produces too few data points to test alternative theories. After compression, falsifying the theory is extremely difficult, as the information incompatible with the theory will largely have been dropped during compression. Likewise, due to having compressed down to M data points, a theory with more than M parameters is difficult to constrain. If the dataset has been compressed in order to ease its public release, then the compression to M data

points directly hinders public reanalyses to study extended models.

In this paper, we therefore define a set of additional compression vectors - basis functions that can investigate extensions of the baseline physics model. Any number of these components can be added, up to the size of the original dataset, and the amplitudes of these modes represent additional parameters to be determined. Although allowing extra flexibility, the price seems to be a potentially large increase in the dimension of the parameter space and a consequently very challenging sampling problem. However, by carefully constructing these modes, we can define a linear model for the extra parameters, based on a generalised principal component analysis (gPCA). The posterior distribution of the additional parameters can then be determined analytically, with almost no additional computational cost. Bayes factors can be computed analytically to determine if an extended model is favoured, and the highest amplitude modes can indicate in which data directions the model needs extending. Furthermore, extended physical models can be expanded in this basis set, allowing their parameters to be inferred at minimal extra cost, and in particular without performing a higher-dimensional Monte Carlo Markov Chain (MCMC) sampling. Our technique allows not only extreme compression of the data, but simultaneously permits model comparison, and additionally creates analytical shortcuts which replace otherwise numerical brute-force approaches in inference.

Throughout the paper, we will refer to the theory on which the MOPED compression is based as the *baseline physics model*. We refer to our extension of the baseline theory, in the gPCA modes, for falsification purposes as the *agnostically extended model*. The extension is agnostic as it introducing the *maximum* number of possible additional parameters, and may be used to test the correctness of the baseline physics model. This provides the most complete study of deviations from the baseline theory. Finally, we can linearly expand *specific extended models* in the agnostic extension basis set. These specific models may be physically or astronomically motivated and introduce fewer extra parameters to capture potential deviations from the standard theory. These models can be tested analytically against the baseline model using Bayesian model comparison techniques. As mentioned, all of this can be achieved via simple matrix operations, without running more, potentially very challenging, MCMC chains. Finally, our agnostic parameters are introduced such that they do not weaken the constraints on the primary parameters of the baseline physical model: this is of central importance, as the introduction of additional nuisance parameters usually inflates the uncertainty of the parameters of primary interest.

The layout of the paper is as follows. In section 2 we review the MOPED radical data compression algorithm, which is designed around a particular (baseline) model. In section 3 we introduce a basis of additional orthogonal vectors as ‘agnostic’ extensions to the baseline theory, and show that the mode amplitudes can be determined analytically without further MCMC studies. In section 4 we show how specific physics-driven extensions to the baseline theory may be tested analytically, by projecting their signal derivatives onto the extended modes. In section 5 we consider a detailed illustrative example of determining structural param-

eters from a galaxy image and explore compressing the extended modes to a subset of those that are most discrepant with the theory. We show that retaining even a few percent of modes may be sufficient to challenge the baseline model and constrain its competitors. In total, this implies public reanalyses of Euclid- or LSST-sized data do not need to succeed in analyzing the full unwieldy data sets, but may study new physics from much smaller, carefully prepared, compressed data sets. Finally, in section 6 we present our conclusions.

2 MOPED WITH UNCHALLENGED THEORIES

We first review MOPED data compression. Introduced by [Heavens, Jimenez & Lahav \(2000\)](#), it is a linear compression method that is optimised for Gaussian-distributed data of fixed covariance matrix. The N -dimensional data vector $\mathbf{x}$ is drawn as

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}), \quad (1)$$

where $\mathcal{N}$ is the multivariate Gaussian distribution of mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{C}$. Denoting expectation values of the data with angled brackets, it is now conjectured that there exists a parametric prediction for the expectation value, e.g. a model from theoretical physics which might explain the data. Assuming this model has M free parameters, we then have

$$\langle \mathbf{x} \rangle \equiv \boldsymbol{\mu}(\boldsymbol{\Theta}), \quad (2)$$

where $\boldsymbol{\Theta} = (\Theta_1, \dots, \Theta_M)$ are the M free parameters of the theory. In a usual inference problem, the values of these parameters are to be inferred.

We assume that the expectation value $\boldsymbol{\mu}(\boldsymbol{\Theta})$ is differentiable at least twice with respect to its parameters, and write $\partial \boldsymbol{\mu} / \partial \Theta_\alpha = \boldsymbol{\mu}_{,\alpha}$, where the minor Greek index α here runs over the M free parameters of the theory, i.e. $\alpha \in [1, M]$. Then, the MOPED algorithm defines one MOPED weighting vector $\mathbf{b}_\alpha$ per free parameter of the theory. These MOPED vectors act as filters when applied to the data. The dot products of the MOPED filter vectors with the original data set $\mathbf{x}$ thus produce M compressed data points

$$y_\alpha = \mathbf{b}_\alpha^T \mathbf{x}, \quad \alpha \in [1, M]. \quad (3)$$

The MOPED filters $\mathbf{b}_\alpha$ have been derived in [Heavens, Jimenez & Lahav \(2000\)](#), and follow from an optimization procedure, such that the set of compressed data points $\{y_\alpha\}$ contains the same Fisher information on the model parameters $\boldsymbol{\Theta}$ as the full data set $\mathbf{x}$. Under this optimization criterion, the weighting vectors were shown by [\(Heavens, Jimenez & Lahav 2000\)](#) to be

$$\mathbf{b}_1 = \frac{\mathbf{C}^{-1} \boldsymbol{\mu}_{,1}}{\sqrt{\boldsymbol{\mu}_{,1}^T \mathbf{C}^{-1} \boldsymbol{\mu}_{,1}}}, \quad \mathbf{b}_{\alpha>1} = \frac{\mathbf{C}^{-1} \boldsymbol{\mu}_{,\alpha} - \sum_{\beta=1}^{\alpha-1} (\boldsymbol{\mu}_{,\alpha}^T \mathbf{b}_\beta) \mathbf{b}_\beta}{\sqrt{\boldsymbol{\mu}_{,\alpha}^T \mathbf{C}^{-1} \boldsymbol{\mu}_{,\alpha} - \sum_{\beta=1}^{\alpha-1} (\boldsymbol{\mu}_{,\alpha}^T \mathbf{b}_\beta)^2}}. \quad (4)$$

For $M > 1$ the modes are Gram-Schmidt orthogonalized with respect to the covariance matrix as metric tensor. This is an optional step which does not interfere with optimality, and one can also use simpler non-orthonormalised

weights such as $\mathbf{C}^{-1} \boldsymbol{\mu}_\alpha$, as used by [Gualdi et al. \(2018\)](#) and [Zablocki & Dodelson \(2016\)](#), and the compression is equivalent to MOPED. To set the weights, a fiducial parameter set needs to be defined, at which the derivatives are computed.

The Gram-Schmidt procedure ensures orthonormality of the modes, which is defined by the relation

$$\mathbf{b}_\alpha^T \mathbf{C} \mathbf{b}_\beta = \delta_{\alpha\beta}. \quad (5)$$

With Gram-Schmidt orthogonalisation, the ultimate outcome is that the compressed data points of Eq. (3) are uncorrelated, and each has unit variance. i.e. the covariance matrix of the data points $\{y_\alpha\}$ is thus by construction the $M \times M$ identity matrix. Note that this only holds approximately if the fiducial model is not correct, and this is explored further in [Heavens et al. \(2017\)](#), which also treats parameter-dependent covariance matrices: data compression facilitates covariance estimation, and in particular inversion, if the covariance matrices are estimated from simulations ([Sellentin & Heavens 2016, 2017](#)). The MOPED compression has been extended to nonlinear compression by [Alsing & Wandelt \(2018\)](#), who provided further insight into why it works, connecting it to the M components of the gradient of the likelihood function. In the same spirit, neural networks can be used for non-Gaussian data distributions to find the optimal weights ([Charnock, Lavaux & Wandelt 2018](#)). Finally, [Zablocki & Dodelson \(2016\)](#) formed linear combinations of the MOPED coefficients to remove some parameter dependence, at least locally.

In practice, the compression can be massive, since it reduces the original dataset of dimension N to the length of the parameter vector, which is M . The compressed data set then contains one data point per parameter to be measured. MOPED is exactly preserving the Fisher information if the derivatives of the weighting vectors are evaluated at the correct value of the model parameters. For most applications the Fisher matrix will however vary little throughout parameter space, such that MOPED compression is almost optimal also in a realistic setting where the correct parameter values are not known a priori. In total, MOPED thus can lose no information in the Fisher sense, despite its very radical compression. However, a limitation is that the compression is tied to a particular theory, and does not lend itself naturally to the comparison of different theories.

3 DATA COMPRESSION WITH MODEL COMPARISON

3.1 Adding additional filters

With N data points and M theoretical parameters used by the baseline physical model, there are M MOPED vectors, and $N - M$ additional degrees of freedom that could be explored. We thus use these additional degrees of freedom to introduce auxiliary parameters, for the purpose of testing whether the baseline physical model is sufficient to explain the data, or whether a theory with extended parameter space is required.

Contrasting two theories with each other, and computing which of the two is preferred by the data, is answered by Bayesian model comparison. However, introducing new

parameters in an arbitrary way will compromise the precision with which Θ , the parameters of the baseline physical model, are measured. If more parameters are to be jointly inferred, then the uncertainty on each of them will generally increase. If we do not wish to compromise the precision of the baseline physical parameters, then the only possibility is to introduce new parameters under the additional constraint that these be statistically *independent*.

Enforcing independence between parameters can, at the level of a Gaussian approximation, be ensured if the extended model is linear in the new parameters, and if the Fisher matrix between old and new parameters is a block-diagonal matrix, as the correlation between the old and new parameters is then zero. To meet these goals, we thus introduce the additional parameters in the following manner.

The baseline physical model will restrict the expected data vector to an M -dimensional hypersurface in the N -dimensional data space. The actual data vector will not lie in this hypersurface, either due to noise, and/or because the physical model is not correct. Our task is to determine which.

Let us consider the part of the data vector $\mathbf{x}$ that is not in the baseline model hypersurface. This part results from removing the projections of the original data vector onto the MOPED vectors, locally at $\mu_F \equiv \mu(\Theta_F)$, where Θ_F is the fiducial set of baseline parameters. For convenience, we subtract this from $\mathbf{x}$:

$$\mathbf{X} \equiv (\mathbf{x} - \mu_F) - \sum_{\alpha=1}^M [\mathbf{b}_\alpha^T (\mathbf{x} - \mu_F)] \mathbf{C} \mathbf{b}_\alpha. \quad (6)$$

The resultant vector $\mathbf{X}$ effectively points orthogonally out of the hypersurface at the point μ_F . With this definition, $\mathbf{X}$ is orthogonal to all the MOPED vectors, defined at Θ_F , in the normal sense, i.e. $\mathbf{b}_\alpha^T \mathbf{X} = 0$ for all $\alpha \in [1, M]$. It then follows that the covariance matrix of $\mathbf{X}$ is given by

$$\mathbf{C}_{ij}^X = \mathbf{C}_{ij} - \sum_{\alpha=1}^M (\mathbf{C} \mathbf{b}_\alpha)_i (\mathbf{C} \mathbf{b}_\alpha)_j. \quad (7)$$

The Roman indices i, j label the different vector or matrix elements, whereas the Greek index α runs over the parameters.

3.2 Agnostic extension of the baseline theory

The procedure that we follow is to define a set of basis vectors $\mathbf{u}_\gamma; \gamma \in [M+1, N]$, that span the $N - M$ -dimensional space orthogonal to the baseline model hypersurface at the fiducial model μ_F . Deviations from the baseline model will be reflected in non-zero amplitudes of these modes, but $\mathbf{X}$ will also be non-zero because of noise. Our aim is to define modes whose amplitudes can be computed analytically – this frees us from otherwise having to perform an N -dimensional MCMC analysis, which could be prohibitively expensive. We can avoid the MCMC by constructing a linear data model that is effectively a Taylor expansion of the expected signal, adding extra degrees of freedom corresponding to directions in which the expected data vector can move, so the expected signal is modified to

$$\mu(\Theta, \Gamma) = \mu(\Theta) + \sum_{\gamma=M+1}^N \Gamma_\gamma \mathbf{C} \mathbf{u}_\gamma, \quad (8)$$

where we have expressed as $\mathbf{C} \mathbf{u}_\gamma$ the partial derivative of the signal with respect to the new parameter Γ_γ . The first term on the right hand side is $\mu(\Theta, \Gamma = \mathbf{0})$, which is just the expected signal $\mu(\Theta)$ for the baseline physics model. This first-order Taylor expansion in the directions off the hypersurface yields linear parameters in the model that can be treated analytically.

To find the new modes $\mathbf{u}_\gamma$, it is convenient if they have the same orthonormality properties as the MOPED vectors, which span the hyperplane locally. i.e. they are normal to each other,

$$\mathbf{u}_\alpha^T \mathbf{C} \mathbf{u}_\beta = \delta_{\alpha\beta}, \quad \alpha, \beta \in [M+1, N] \quad (9)$$

and also to the MOPED vectors,

$$\mathbf{b}_\alpha^T \mathbf{C} \mathbf{u}_\gamma = 0; \quad \alpha \in [1, M], \gamma \in [M+1, N]. \quad (10)$$

We can satisfy these constraints by solving the generalised eigenvalue problem

$$\mathbf{C}^X \mathbf{u}_\gamma = \Lambda_\gamma \mathbf{C} \mathbf{u}_\gamma. \quad (11)$$

The resulting $N - M$ generalized eigenvectors $\mathbf{u}_\gamma$ with non-zero eigenvalues will form our additional generalised PCA (gPCA) weighting vectors. As many of these as desired are then added to the set of MOPED vectors.

To solve the generalised eigenvalue problem of Eq. (11), we Cholesky decompose the positive definite covariance matrix $\mathbf{C} = \mathbf{L} \mathbf{L}^T$. Rearranging, and introducing the shorthand $\mathbf{L}^{-T} = (\mathbf{L}^T)^{-1}$, this leads to

$$(\mathbf{L}^{-1} \mathbf{C}^X \mathbf{L}^{-T}) \mathbf{L}^T \mathbf{u}_\gamma = \Lambda_\gamma \mathbf{L}^T \mathbf{u}_\gamma, \quad (12)$$

which is now an ordinary eigenvalue problem

$$\mathbf{A} \mathbf{v}_\gamma = \Lambda_\gamma \mathbf{v}_\gamma, \quad (13)$$

where the matrix $\mathbf{A} \equiv \mathbf{L}^{-1} \mathbf{C}^X \mathbf{L}^{-T}$, and

$$\mathbf{v}_\gamma = \mathbf{L}^T \mathbf{u}_\gamma, \quad (14)$$

are the eigenvectors of $\mathbf{A}$ with eigenvalues Λ_γ . We can find our generalized eigenvectors $\mathbf{u}_\gamma$ by solving Eq. (14).

The eigenvectors $\mathbf{v}_\gamma$ can be made orthogonal in the usual sense, and indeed have to be, since it is straightforward from the expression for $\mathbf{C}^X$ (equation 7) to show that the eigenvalues of the problem are either 0 (with multiplicity M), or 1 (with multiplicity $N - M$). The off-hypersurface eigenvectors (with eigenvalue 1) are arbitrary provided they are orthogonal, but their precise directions are not important – any mutually orthogonal set will do. The M eigenvectors $\mathbf{u}_\gamma = \mathbf{L}^{-T} \mathbf{v}_\gamma$ with eigenvalues equal to 0 are discarded, as they lie in the local hyperplane already spanned by the MOPED vectors, and the data dependence there is already captured by the MOPED vectors.

From the (usual) orthogonality of $\mathbf{v}_\gamma$, the orthogonality condition for the $\mathbf{u}_\gamma$ is

$$\delta_{\alpha\beta} = \mathbf{v}_\alpha^T \mathbf{v}_\beta = (\mathbf{L}^T \mathbf{u}_\alpha)^T (\mathbf{L}^T \mathbf{u}_\beta) = \mathbf{u}_\alpha^T \mathbf{C} \mathbf{u}_\beta, \quad (15)$$

as desired. We then add the gPCA $\mathbf{u}_\gamma$ eigenvectors for which $\Lambda_\gamma = 1$ (i.e. $\gamma \in [M+1, N]$) to the set of MOPED filters $\mathbf{b}_\alpha$.

The filters with which we now sift through the data are then

$$\mathbf{f}_\alpha = \begin{cases} \mathbf{b}_\alpha, & \text{for } 1 \leq \alpha \leq M, \\ \mathbf{u}_\alpha, & \text{for } M+1 \leq \alpha \leq N. \end{cases} \quad (16)$$

and the data points are

$$Y_\alpha = \mathbf{f}_\alpha^T (\mathbf{x} - \boldsymbol{\mu}_F), \quad \alpha \in [1, N], \quad (17)$$

where we have for convenience modified the MOPED algorithm to subtract the expected signal at the fiducial point.

The first M data points, corresponding to the MOPED modes, give the maximum useful compression if the theory is correct. The remaining $N - M$ data points serve the purpose of being able to falsify the baseline physical theory. If the baseline theory is correct, then these data points contain no information, only noise. If the dataset is to be published for reanalysis, then the M MOPED datapoints would be provided, and an ‘interesting’ subset (to be defined in Sect. 3.4) of the additional points should be added, to enable model comparison and studies of non-standard theories with the public dataset.

This establishes our complete set of weight vectors. Of course, if we choose to keep all $N - M$ gPCAs, then there is no compression at all, but we can combine the advantages of massive data compression with an investigation of whether the physical model requires extension by including a subset of the extra modes. As we will see in the next sections, this construction has a number of very useful properties.

3.3 Extended model parameters

In this section, we derive the posterior distribution of the extended model parameters, Γ_γ , introduced in Eq. (8). We will show that, subject to certain benign conditions, their marginal posteriors are independent Gaussian distributions of unit variance. The mode of their posterior can be computed analytically, and they fully decouple from the baseline model parameters; introducing $N - M$ additional parameters Γ_γ for the sake of model comparison thus does not deteriorate our knowledge of the primary physical parameters.

The analytic solution is an important practical point: We may infer, if desired, all $N - M$ extra parameters without undertaking an expensive N -dimensional MCMC sampling.

3.3.1 Gaussianity

If the original data $\mathbf{x}$ are Gaussian-distributed, then so are the compressed data, since they are linear combinations of $\mathbf{x}$. If we assume a Gaussian prior on Γ (a uniform distribution also works), then the posterior probability as a function of Γ_γ will be Gaussian, since

$$\mathcal{P}(\Gamma, \Theta | \mathbf{Y}) \propto \mathcal{P}(\mathbf{Y} | \Theta, \Gamma) \pi(\Theta, \Gamma) \quad (18)$$

where π is the prior, Gaussian in Γ , and the likelihood $\mathcal{P}(\mathbf{Y} | \Theta, \Gamma)$ is Gaussian in both $\mathbf{Y}$ and Γ :

$$\mathcal{P}(\mathbf{Y} | \Theta, \Gamma) = (2\pi)^{-N/2} \exp \left(-\frac{1}{2} \sum_{\beta=1}^N \left\{ Y_\beta - \mathbf{f}_\beta^T [\boldsymbol{\mu}(\Theta, \Gamma) - \boldsymbol{\mu}_F] \right\}^2 \right), \quad (19)$$

where $\boldsymbol{\mu}(\Theta, \Gamma)$ is linear in Γ_γ and given by equation (8).

The maximum of the likelihood (ML) is at $(\Theta_{\text{ML}}, \Gamma_{\text{ML}})$, where Γ_{ML} can be computed analytically for given Θ_{ML} , since $\mathbf{f}$ and $\mathbf{u}$ are orthogonal:

$$\Gamma_{\text{ML}, \gamma} = Y_\gamma - \mathbf{f}_\gamma^T [\boldsymbol{\mu}(\Theta_{\text{ML}}) - \boldsymbol{\mu}_F]; \quad \gamma \in [M + 1, N]. \quad (20)$$

This is an important result. If the prior on Γ_γ is uniform, then the peak of the posterior is given analytically by equation (20), and this is easily generalised to the peak of the posterior when a Gaussian prior is assumed. So we find that once a MOPED analysis has been done to find Θ_{ML} , only trivial computations are required to infer Γ_γ .

3.4 Hessian matrix and data compression

The likelihood of the Γ_γ is Gaussian, with the mean given by equation (20). The covariance is obtained from the Hessian matrix, which is composed of four blocks, whose expressions simplify because of the orthonormality of the $\mathbf{f}_\beta$. The Θ - Θ block is the same as in the MOPED analysis, given at the peak by

$$-\frac{\partial^2 \ln \mathcal{L}}{\partial \Theta_\alpha \partial \Theta_\beta} = \sum_{\rho=1}^N (\mathbf{f}_\rho^T \boldsymbol{\mu}_{\alpha, \rho}) (\mathbf{f}_\rho^T \boldsymbol{\mu}_{\beta, \rho}). \quad (21)$$

The gPCA-gPCA block of additional modes is the identity matrix in all circumstances:

$$-\frac{\partial^2 \ln \mathcal{L}}{\partial \Gamma_\gamma \partial \Gamma_\rho} = \delta_{\gamma \rho}. \quad (22)$$

The MOPED-gPCA off-diagonal block is

$$-\frac{\partial^2 \ln \mathcal{L}}{\partial \Theta_\alpha \partial \Gamma_\gamma} = \mathbf{f}_\gamma^T \boldsymbol{\mu}_{\alpha, \gamma}. \quad (23)$$

Here we see that the off-diagonal blocks are zero, provided that the ML point is the fiducial model, or the baseline physics model is linear in the parameters. Otherwise the tangent hyperplane spanned by the $\mathbf{C}\mathbf{b}_\alpha$, which is defined at the fiducial parameters, is not tangent at the ML point.

As a result, if the ML point is not known from a previous uncompressed analysis, then the MOPED analysis may need to be repeated, first to find the ML point, and then compressed during the second time with the found ML point as the fiducial model parameters. In this way, the baseline physical model parameters and the extended parameters Γ_γ are made independent, and the joint distribution of (Θ, Γ) is

$$\mathcal{P}(\Theta, \Gamma | \mathbf{x}) = \mathcal{P}(\Theta | \mathbf{y}) \prod_{\gamma=M+1}^N \mathcal{N}(\Gamma_\gamma | \Gamma_{\text{ML}, \gamma}, 1), \quad (24)$$

where the first term is the MOPED posterior.

This construction has a number of useful consequences. Firstly, the introduction of the extra parameters has no effect on the inference of the baseline model parameters, whose credible regions are not degraded by their inclusion. Secondly, any number of the $N - M$ possible additional modes may be included. Leaving any out is equivalent to marginalising over them. Thirdly, in the baseline model, the expected values of Γ_γ are zero and their variance unity (for uniform priors), so the means of the posterior (in the more general case) can be used to assess by analytic model comparison techniques whether extension to the baseline model should be considered. Finally, if we wish to retain a compressed dataset for comparing future models with the baseline model, then a sensible ordering is to keep those that are most discrepant with baseline model (those with the largest $|Y_\beta|$).

In conclusion for this section, we find that with the generalised PCA approach, the likelihood is separable, with the

gPCA parameters being independent of each other and of the baseline physical model parameters used in the MOPED part. If the derivatives are evaluated at the ML point, the expected likelihood of the gPCA parameters is the product of zero mean, unit variance Gaussians.

3.5 Bayes factor: the need for an extended model?

In this section we describe how to quantify whether the compressing baseline theory is insufficient and whether a new model should accordingly be sought. If the baseline is found to be insufficient, then Sect. 4. describes how a subset of the auxiliary parameters Γ_γ can be used to quickly, and analytically, predict which other theory is likely to alleviate the defects of the baseline theory.

We quantify the insufficiency of the baseline theory by the Bayesian evidence: the evidence is the probability of the observed data given a model, independent of the values its free parameters might take. We denote the extended MOPED+gPCA model by M_E and the original baseline physical model by M_B . We then compute the Bayesian evidence Z_E for the extended MOPED+gPCA model M_E (with as many extra gPCA parameters as desired) and the baseline physical model, M_B . These are combined to the Bayes factor, $B_{EB} = Z_E/Z_B$, which gives the probabilistic ratio of whether the agnostically extended theory is preferred ($B_{EB} > 1$), or whether the physical baseline theory is favoured ($B_{EB} < 1$), assuming equal prior model probabilities. For concreteness, we here assume all gPCAs are included, so there is no data compression, but fewer may be considered if desired. This would simply correspond to truncating the following sums early.

The evidence for the extended MOPED+gPCA model is

$$Z_E = \int \mathcal{P}(\mathbf{Y}|\boldsymbol{\Theta}, \boldsymbol{\Gamma}, M_E) \pi(\boldsymbol{\Theta}, \boldsymbol{\Gamma}|M_E) d^M \boldsymbol{\Theta} d^{N-M} \boldsymbol{\Gamma}, \quad (25)$$

where π is the prior. For the baseline model,

$$Z_B = \int \mathcal{P}(\mathbf{Y}|\boldsymbol{\Theta}, M_B) \pi(\boldsymbol{\Theta}|M_B) d^M \boldsymbol{\Theta}. \quad (26)$$

Since the models are nested (the baseline model corresponds to the extended model with $\Gamma_\gamma = 0$), the Bayes factor is the Savage-Dickey density ratio (see Trotta (2007) for details):

$$B_{BE} = \frac{\mathcal{P}(\boldsymbol{\Gamma} = \mathbf{0}|\mathbf{Y}, M_E)}{\pi(\boldsymbol{\Gamma} = \mathbf{0}|M_E)}, \quad (27)$$

where we have assumed that the prior is separable, and the prior for the common parameters is the same for both models, $\pi(\boldsymbol{\Theta}|M_E) = \pi(\boldsymbol{\Theta}|M_B)$:

$$\pi(\boldsymbol{\Theta}, \boldsymbol{\Gamma}|M_E) = \pi(\boldsymbol{\Theta}|M_B) \pi(\boldsymbol{\Gamma}|M_E). \quad (28)$$

Now we assume an uncorrelated Gaussian prior centred on zero for the extended model parameters

$$\pi(\boldsymbol{\Gamma}|M_E) = \frac{1}{\prod_\gamma \sqrt{2\pi}\Sigma_\gamma} \exp\left(-\sum_{\gamma=M+1}^N \frac{\Gamma_\gamma^2}{2\Sigma_\gamma^2}\right). \quad (29)$$

Performing the Gaussian integrals to normalise properly the

posterior in the numerator, we find an analytic expression for the Bayes factor

$$\ln B_{BE} = \frac{1}{2} \sum_{\gamma=M+1}^N \left[\ln(1 + \Sigma_\gamma^2) - \frac{\Gamma_{ML,\gamma}^2}{1 + \Sigma_\gamma^2} \right]. \quad (30)$$

There is no obvious way to set Σ_γ , which we will assume to be the same for all γ , but the problem does have a natural scale, since by construction the expected width of the likelihood for these components in the baseline model is unity. It makes sensible to choose a prior that is greater than unity, and for this paper, we have typically chosen $\Sigma_\gamma = 5$, but quote some results with a prior width of 10.

However, it enters only logarithmically if $\Sigma_\gamma \gg 1$. The sum is over those extra gPCA parameters that are included. The greater the values of $\Gamma_{ML,\gamma}^2$, the more the gPCA model is favoured, i.e. the less adequate is the original baseline theory.

For completeness, we note that the χ^2 of the best-fit base model may be computed in this basis as

$$\chi^2 = \sum_\gamma \Gamma_{ML,\gamma}^2. \quad (31)$$

In summary, our setup allows to compute the Bayes factors for agnostically extended models analytically, which avoids the otherwise highly non-trivial numerical integration to yield an evidence. This therefore removes an important computational bottleneck. Since inferring the amplitudes of additional modes is almost certainly much faster than performing the baseline MCMC analysis, there is little reason not to infer all of the gPCA mode amplitudes.

Finally, if we assume the evidence calculation indicates that the baseline model is indeed insufficient, then we now require a fast route to establishing which other physically motivated theory will fix the flaws of the baseline theory. This will be solved in the following Sect. 4.

4 PARAMETER INFERENCE WITH SPECIFIC EXTENDED MODELS

One advantage of constructing our basis of agnostic extended models is that we can easily investigate *specific extended models* by projecting them on to the gPCA eigenmodes. This again enables us to constrain the parameters of extended theories analytically, alongside testing whether these extended theories solve shortcomings of the baseline theory. Both these issues would otherwise need to be investigated numerically.

The caveat here is that we will use a first-order Taylor expansion, so a fiducial set of extended parameters needs to be assumed. If the derivatives $\mu_{,\phi}$ with respect to the extended model parameters ϕ , are strongly parameter-dependent, it may be necessary to iterate to a solution in the extended model space, and perform the linear expansion there. Fortunately this can all be done analytically without higher-dimensional MCMC chains in the extended parameter space.

We thus assume there exists one, or multiple, extended theories which use the physically motivated parameters Φ_ϕ with $\phi = 1 \dots s$. We assume here that the baseline theory

is ‘nested’ into the extended theory, i.e. for $\Phi = 0$ the extended model² falls back onto the baseline theory. Although this is not strictly necessary, it simplifies the mathematical formalism below.

Constraining the extended theory implies the posterior distribution of the new parameters Φ must be calculated. Additionally, the introduction of new parameters may also shift the ML values of the former parameters Θ . We thus also have to compute the shifts $\Delta\Theta_\alpha; \alpha = 1 \dots m$ of the baseline model parameters arising from the shift in Φ_ϕ away from zero.

With Θ representing the baseline parameters, Γ the agnostic extension parameters, Φ the specific extended model parameters, and $\bar{Y}_\beta$ being the mean predicted by the model, the log-likelihood reads

$$-2 \ln \mathcal{L} = \sum_{\beta=1}^n [Y_\beta - \bar{Y}_\beta(\Theta, \Gamma \text{ or } \Phi)]^2. \quad (32)$$

Treating $\mathcal{L}$ as a function of Θ, Φ , the ML is at

$$\begin{aligned} \frac{\partial \ln \mathcal{L}}{\partial \Phi_\phi} &= \sum_{\beta} [Y_\beta - \bar{Y}_\beta] \frac{\partial \bar{Y}_\beta}{\partial \Phi_\phi} = 0 \\ \frac{\partial \ln \mathcal{L}}{\partial \Theta_\alpha} &= \sum_{\beta} [Y_\beta - \bar{Y}_\beta] \frac{\partial \bar{Y}_\beta}{\partial \Theta_\alpha} = 0. \end{aligned} \quad (33)$$

Here we make an approximation that $\bar{Y}_\beta$ is linearly dependent on Θ and Φ . This is accurate locally, but whether it will be a sufficiently accurate approximation globally depends on an interplay of parameter non-linearity and constraining power of the data. As shown in [Sellentin, Quartin & Amendola \(2014\)](#) and [Sellentin \(2015\)](#), the approximation that $\bar{Y}_\beta$ is linearly dependent on Θ and Φ could be extended to higher-order terms, but if the data are sufficiently constraining, then those higher-order terms are rendered irrelevant as they only dominate in parameter regions that are assigned negligible posterior probability. Thus, our linear approximation is expected to also handle a wide range of non-linear parameters accurately, as long as the data are sufficiently constraining.

Expanding around the baseline ML point, where $\bar{Y} = \bar{Y}^0$ and $\Theta_\alpha = \Theta_\alpha^0; \Phi_\phi = 0$, we introduce the shift of the old parameters

$$\Delta\Theta_\alpha \equiv \Theta_\alpha - \Theta_\alpha^0. \quad (34)$$

This shift describes the displacement in the posterior of the baseline model due to having introduced new parameters. It gives the distance (and direction) between the old best-fit and the location of the new best-fit to be expected when introducing the parameters Φ . We then have

$$\bar{Y}_\beta = \bar{Y}_\beta^0 + \sum_{\phi=1}^s \mathbf{f}_\beta^T \cdot \left(\Phi_\phi \frac{\partial \mu}{\partial \Phi_\phi} \right) + \sum_{\alpha=1}^m \mathbf{f}_\beta^T \cdot \left(\Delta\Theta_\alpha \frac{\partial \mu}{\partial \Theta_\alpha} \right). \quad (35)$$

Defining non-square matrices V and W with elements

$$V_{\phi\beta} \equiv \mathbf{f}_\beta^T \cdot \frac{\partial \mu}{\partial \Phi_\phi}; \quad W_{\alpha\beta} \equiv \mathbf{f}_\beta^T \cdot \frac{\partial \mu}{\partial \Theta_\alpha}, \quad (36)$$

we have (now with summation convention over ranges that

are set by context)

$$\begin{aligned} 0 &= \frac{\partial \ln \mathcal{L}}{\partial \Phi_\phi} = [Y_\beta - \bar{Y}_\beta^0 - \Phi_\rho V_{\rho\beta} - \Delta\Theta_\alpha W_{\alpha\beta}] V_{\phi\beta}, \\ 0 &= \frac{\partial \ln \mathcal{L}}{\partial \Theta_\alpha} = [Y_\beta - \bar{Y}_\beta^0 - \Phi_\rho V_{\rho\beta} - \Delta\Theta_\eta W_{\eta\beta}] W_{\alpha\beta}. \end{aligned} \quad (37)$$

These equations have vector solutions

$$\begin{aligned} \Delta\Theta &= X^{-1} [W - (WV^T)(VV^T)^{-1}V] \Delta Y, \\ \Phi &= Z^{-1} [V - (VW^T)(WW^T)^{-1}W] \Delta Y, \end{aligned} \quad (38)$$

where $\Delta Y_\beta \equiv Y_\beta - \bar{Y}_\beta^0$ and

$$\begin{aligned} X &\equiv WW^T - WV^T(VV^T)^{-1}VW^T, \\ Z &\equiv VV^T - VW^T(WW^T)^{-1}WV^T. \end{aligned} \quad (39)$$

These calculations require only $O(N^2)$ operations per parameter. Note that we can equally well do this linear model analysis with the uncompressed data, at the expense of computing C^{-1} , and similar analytic expressions are reached.

The central result of this section is Eq. (38). It describes by how much the best-fit of the old parameters Θ will shift, and in which direction. It also yields a direct analytical estimate of the values that the new parameters Φ will take. Eq. (38) therefore provides an analytical study of whether the extended model will improve upon the baseline theory.

The variances of the baseline model parameters, the specific extended model parameters, and their covariances, are obtained as a block matrix constructed from V and W :

$$-\frac{\partial^2 \mathcal{L}}{\partial \Theta_\alpha \partial \Theta_\beta} = W_{\beta\eta} W_{\alpha\eta}; \quad (40)$$

$$-\frac{\partial^2 \mathcal{L}}{\partial \Phi_\phi \partial \Phi_\psi} = V_{\phi\eta} V_{\psi\eta}; \quad (41)$$

$$-\frac{\partial^2 \mathcal{L}}{\partial \Theta_\alpha \partial \Phi_\phi} = W_{\alpha\eta} V_{\phi\eta}. \quad (42)$$

One advantage of this basis is that the data can be compressed, and the agnostic gPCA coefficients with the smallest Γ values discarded. We will see in the example below that this may be an effective way to retain an expanded set of compressed data (including the MOPED coefficients and a subset of the most discrepant gPCA modes), so that alternative models can be confronted very rapidly (albeit approximately) with the compressed data, without undertaking a full MCMC analysis. In the test case considered later in Sect. 5, retaining 5% of the gPCA modes will prove to be enough to show that that baseline model is inadequate.

5 EXAMPLE: GALAXY IMAGES

The method is quite general, but for illustration, we consider the same 4-parameter model as in [Heavens et al. \(2017\)](#), for the surface brightness of a thin exponential disk galaxy seen at an angle. Parameters to be inferred are the scalelength a , central amplitude A , eccentricity ϵ and position angle ϕ . The variables r and φ are polar coordinates of pixels in the image, meaning they are required to define the pixelized data, and are consequently not to be inferred. We generate artificial data of such pixelized galaxy images and add noise which is

² We assume w.l.o.g. that the baseline model has $\Phi_\phi \equiv 0$.




Figure 1. Examples of the galaxy image model of Section 5. Left: Noise-free model of a pixelized galaxy. Right: the same galaxy with additive shot noise.

assumed to be uncorrelated, with a white noise background uncertainty plus a source Poisson term. The physical model of the brightness is then

$$\mu(r, \varphi; a, A, \epsilon, \phi) = A \exp \left[-\frac{r\sqrt{1+\epsilon^2} - 2\epsilon \cos 2(\varphi - \phi)}{a\sqrt{1-\epsilon^2}} \right]. \quad (43)$$

The observed surface brightness is

$$\hat{\mu}_{\text{obs}}(r, \varphi) = \mu(r, \varphi; a, A, \epsilon, \phi) + n(r, \varphi). \quad (44)$$

The Poisson shot noise $n(r, \varphi)$ is assumed to include enough photons to be accurately treated as Gaussian.

We generate an $n \times n$ galaxy image with parameters $a = n/4$, $A = 100$, $\epsilon = 0.25$, $\phi = 1$. We add a constant background level of 100. A visualization with $n = 20$ is shown in Fig. 1.

The posterior using the full, uncompressed, dataset is shown in Fig. 2, for an image with $n = 10$. The posterior constraints derived from the MOPED compressed data set are shown in Fig. 3, revealing that MOPED indeed does not lose precision, despite outputting only 4 compressed data points, and all parameters being non-linear except for A . The inference is done with Hamiltonian Monte Carlo sampling of 2 chains of 30000 points using Stan (Carpenter et al. 2012), with a burnin of 1000 points. The corner plots show posterior densities assuming a uniform prior, so likelihood and posterior will be used interchangeably in the parameter inference discussion. For model comparison purposes later we will assume a Gaussian prior.

The MOPED vectors are shown in Fig. 4, and four gPCA vectors are shown in 5. The inference with the MOPED coefficients plus the projections on to the first 4 gPCA modes is shown in Fig. 6. Note that the extra parameters are consistent with zero, implying the baseline model assumed by MOPED is sufficient to describe the data. The posteriors are, within the sampling errors, unit width gaussians centred on the analytically-calculated means, as expected for a correctly-specified model expanded around the posterior mode. With all additional components, the Bayesian evidence favours the baseline model with $\ln B_{\text{BE}} = 53.7$, which is highly decisive in favour of the simple baseline model. For information only, $\chi^2 = 107$ for 96 extra components.




Figure 2. Posterior constraints of the galaxy's four parameters from the full dataset of 100 data points.




Figure 3. Posterior from MOPED-compressed dataset, which comprises 4 data points only. The comparison with respect to Fig. 2 reveals that the data compression indeed does not lose constraining power on the parameters.

5.1 Conservation of precision despite extended parameter set

The means and standard deviations from MCMC analysis of the full likelihood and the compressed data are shown in Table 1. The MOPED coefficients are derived assuming the correct model parameters. The maximum 96 extra gPCA components are added. MOPED on its own incurs no information loss at Fisher level, and we see that this is still the case while jointly inferring the extra gPCA components. This is expected, since the gPCA components are designed to have probability distributions that are independent of the MOPED coefficients.




Figure 4. The four MOPED filters, which filter the data for the information on the parameters a , A , ϵ , ϕ . The filters and the data are plotted in the pixel domain, using Cartesian coordinates. Applying the filters to the data implies taking the pixelwise dot product between these filters and the pixelized data.




Figure 5. Four gPCA vectors as used in the analysis. The gPCA vectors act as additional filters to sweep through the dataset, thereby outputting agnostic compressed data points from which non-standard physics can first be discovered and in a subsequent step constrained.

	True value	Mean (full data)	Mean (MOPED +gPCA)	s.d. (full)	s.d. (MOPED +gPCA)
a	2.5	2.54	2.54	0.10	0.10
A	100	100.7	100.7	4.3	4.3
ϵ	0.25	0.29	0.29	0.03	0.03
ϕ	1.0	0.94	0.95	0.05	0.05

Table 1. Mean and standard deviation from chains run with the full dataset (100 data points), and with the MOPED+gPCA data. By construction, even with the MCMC runs being different, MOPED+gPCA maintains the optimal constraints of MOPED despite having introduced 96 additional parameters to fit for. The peak positions are maintained as well.

5.2 Wrong physical model: presence of a bulge

To test the ability of the method to detect failures of the baseline model, we generate data from another model. In this case, we analyse the data with the full set of baseline modes, and we might expect to see large values of the gPCA coefficients Γ_γ , and a Bayes factor that prefers their inclusion in the agnostic model.

Specifically, we generate an image which has an additional circularly symmetric bulge with a Gaussian projected radial profile, parameterized as an additive term $A_b \exp(-\alpha_b r^2)$. An example image is shown in Fig. 7 where we show a noisy disk+bulge image, the noise-free model, and the maximum likelihood solution that assumes the baseline model (disk, with no bulge). The latter solution is an adequate fit by eye, but the baseline-model parameters are not correctly recovered.

Fig. 8 shows the posterior for the physical parameters and 4 gPCA coefficients assuming a uniform prior, for a galaxy with $A_b = 0.3A$, $\alpha_b = 1/(2a^2)$ and $n = 10$. Note that the posterior probability density that $\Gamma_8 = 0$ in particular is very low. Bayesian Evidence disfavors the physical model, with $\ln B_{BE} = -6.8$, for a prior width $\Sigma_\gamma = 5$ for all additional gPCA components, showing strong evidence for the need to modify the data model. As is always the case, the Bayes factor depends on the prior: with $\Sigma_\gamma = 10$, $\ln B_{BE} = -9.7$, with the same conclusion. For information only, $\chi^2 = 13.8$ for the 4 gPCA modes. With more components, the evidence would be stronger.

5.3 Finding a better suited model

The agnostic extended (gPCA) modes can be used to confront any physical model. Here we consider a nested model, where the baseline model is a special case of the extended model. By projecting the derivatives in the extended model onto the gPCA modes, we can infer the extended parameters, and correct the baseline parameters, using equation (38). We consider the additive Gaussian bulge described in Sec. 5.2. We infer the central amplitude $\Phi_0 = A_b$ and the scale parameter $\Phi_1 = \alpha_b$, which is highly nonlinear and more challenging to infer. For this reason, we use larger images with $n = 30$. In this example, $A_b = 0.2A$ and $\alpha_b = 1/18$, corresponding to a Gaussian width of 3 pixels.

Fig. 9 shows the difference between the true extended model and the maximum-likelihood baseline fit, showing that the baseline model fails to reproduce the central structure with discrepancies of the order of 1.2%. In the middle panel, we see the results of using the gPCA modes to constrain the extended model, calculating shifts to the baseline parameters and estimating the extended model parameter (the bulge amplitude). These are computed according to equation (38). The residuals are now at the 0.2% level. Fig. 10 shows in dark blue the distribution of recovered maximum likelihood baseline and extended model parameters, with the truth shown as a red circle. For this application, the linear assumption recovers the parameters reasonably well. The pale blue histogram shows that the recovery is still reasonably good when only the most discrepant 50% of the gPCA modes is used. The standard deviations of all the dark blue histograms agree to better than 4% with the posterior widths predicted in eqs. (40-42), consistent with the




Figure 6. Posterior from MOPED + 4 gPCA modes. Note that the extra parameters are consistent with zero. The precision of the baseline model parameters is unchanged from analyzing the full data set, however many new parameters are analysed.




Figure 7. Image generated from the extended Disk+bulge full data model with (left) and without noise (centre). On the right is the maximum likelihood baseline (disk without bulge) model, which visually matches the data quite closely but yields incorrect baseline parameters.

variation expected from the finite number of samples, but note that they are not quite the same quantity.

If we reduce the proportion of gPCA modes further, the recovery of parameters becomes poor, but a radical compression can still be useful. Fig. 11 shows the recovered parameters from only 5% of the gPCA modes, plus the MOPED coefficients. The recovery is poor, but still shows clear evidence (especially in A_b) that the baseline model is inadequate, and

extended models should be considered more carefully. The moral of this part is that a massive data compression to the MOPED coefficients plus a few gPCA modes may be used to test at a basic level the baseline model against various extended models, without further MCMC calculation, and using only fast linear algebra. Finally, for completeness, we show in Fig. 12 that if the baseline model is correct, and there is no bulge, then these 5% most discrepant modes will support the baseline model, and do not erroneously show evidence that an extended model is preferred.

6 CONCLUSIONS

In this paper we have extended the MOPED extreme data compression algorithm to address a specific issue, that the MOPED algorithm is inextricably tied to a particular data model, which we describe as the baseline theoretical model. In the original MOPED compression setup, this has the disadvantage that if a different theory is considered, the analysis has to be redone from scratch.




Figure 8. MOPED + 4 gPCA inferred parameters for data that are drawn from a disk+bulge model, whereas the baseline physical model has disk only (see text for details). Note for example that the distribution of Γ_8 is noticeably offset from zero.




Figure 9. Difference between maximum likelihood baseline model fit and the image generated from the extended model. Centre panel shows the differences after the linear correction of the baseline model parameters and the bulge amplitude (without noise). The right panel shows the difference between the data and the extended model inferred using the MOPED + gPCA modes.

In this paper, we augment the MOPED weighting vectors with a set of orthogonal modes that span the subspace that is not in the local hyperplane that is spanned by the MOPED vectors. The construction of these modes leads to a generalised principle component analysis (gPCA) of the subspace, and the combined set of MOPED-PC filters have a number of desirable features. Between them, they corre-

spond to all possible directions in data space, and each mode has connected with it an amplitude whose expected distribution is a unit variance Gaussian centred on zero, if the baseline model is correct. These modes are independent of each other and of the MOPED coefficients, if the fiducial parameters used for the compression are correct.

By including all the additional modes, we can investigate a ‘generic extended model’ with The advantages of these properties are several: firstly, the maximum likelihood amplitude of the modes can be calculated analytically once the baseline model has been investigated with MCMC - further (higher-dimension) MCMC analyses are not necessary; secondly, because of the independence, the investigation of these extra modes has no bearing on the precision of the baseline parameters, whose inference is not degraded; thirdly, by assigning priors to the extended mode amplitudes, the Bayes factor between the baseline model and the generic extended model can be computed analytically, so the baseline model can be compared with a generic model to see if the data indicate that further models should be investigated.




Figure 10. Maximum likelihood estimates of the bulge amplitude A_b from 500 noise realisations, and of the baseline model parameters, using all the gPCA modes (dark blue), and using half of the modes with the highest values of $|\Delta Y|$, which are most discrepant with the baseline model expectation (light blue). True values are marked with the red dot, and the distributions of the baseline parameters from the baseline MOPED analysis are shown in green.




Figure 11. As in Fig. 10, but with the pale blue histograms showing the ML estimates from retaining only 5% of the gPCA modes. This radical compression does not recover the parameters well, but still provides evidence (especially from $A_b \neq 0$) that the baseline model is not preferred over the extended model, which should be investigated further.

The extended gPCA modes also form a convenient basis for confronting the data with specific extended models, whose data derivatives can be projected onto the modes, and the extra parameters of the extended theory determined under the approximation that the changes in the mean data vector are linear in the parameters. If the extended model holds, then the posterior distributions of the baseline parameters from the MOPED analysis will be systematically wrong, but they can be corrected, and the extended model parameters determined, using linear algebra alone. This is not a surprise for linear models, and the effectiveness of the procedure will depend on how good the linear approximation is.

Finally, there is the option to keep the radical data compression of MOPED, which is optimal for the baseline model, and increase the size of the compressed dataset by retaining the gPCA modes that are most discrepant with the baseline theory. This is a sensible way to maximise sensitivity to new




Figure 12. As in Fig. 11, but here there is no bulge ($A_b = 0$), and the baseline model is correct. The highly compressed dataset of 4 MOPED coefficients plus 5% of the gPCA modes is consistent with the baseline model, as desired.

physics that might come along later. We find in the example case considered that a reduction of the gPCA modes to only 5% of the original size can give a very strong indication that the baseline model needs extension, even if these modes alone are insufficient to correct fully the baseline parameters and to recover the true extended parameters.

Acknowledgements

We thank the anonymous referee for helpful comments.

Data Availability

No new data were generated or analysed in support of this research.

REFERENCES

- Abbott T. M. C., et al., 2018, *PhRvD*, 98, 043526
- Alsing J., Wandelt B., 2018, *MNRAS*, 476, L60
- Alsing J., Wandelt B., 2019, *MNRAS*, 488, 5093
- Alsing J., Charnock T., Feeney S., Wandelt B., 2019, *MNRAS*, 488, 4440
- Amendola L., et al., 2018, *LRR*, 21, 2
- Carpenter B., Gelman A., Hoffman M., Lee D., Goodrich B., Betancourt M., Brubaker M., Guo J., Li P., Riddell A., 2017, *Journal of Statistical Software*, 76, 1
- Charnock T., Lavaux G., Wandelt B. D., 2018, *PhRvD*, 97, 083004
- Gualdi D., Manera M., Joachimi B., Lahav O., 2018, *MNRAS*, 476, 4045
- Heavens A. F., Jimenez R., Lahav O., 2000, *MNRAS*, 317, 965.
- Heavens A. F., Sellentin E., de Mijolla D., Vianello A., 2017, *MNRAS*, 472, 4244
- Hildebrandt H., et al., 2020, *A&A*, 633, A69
- Leauthaud A., et al., 2017, *MNRAS*, 467, 3024
- Leclercq F., 2018, *PhRvD*, 98, 063511
- LSST Science Collaboration, et al., 2009, *arXiv*, arXiv:0912.0201
- Planck Collaboration, et al., 2018, *arXiv*, arXiv:1807.06209
- Reichardt C., Jimenez R., Heavens A. F., 2001, *MNRAS*, 327, 849
- Sellentin E., Quartin M., Amendola L., 2014, *MNRAS*, 441, 1831
- Sellentin E., 2015, *MNRAS*, 453, 893
- Sellentin E., Heavens A.F., 2016, *MNRAS*, 456, 132

- Sellentin E., Heavens A.F., 2017, MNRAS, 464, 4658
Tegmark M., Taylor A.N., Heavens A.F., 1997, ApJ, 480, 22
Trotta R., 2007, MNRAS, 378, 72
Zablocki A., Dodelson S., 2016, PRD, 93, 083525

ORIGINAL UNEDITED MANUSCRIPT