



HAL
open science

Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology

Justin Alsing, Benjamin Wandelt, Stephen Feeney

► **To cite this version:**

Justin Alsing, Benjamin Wandelt, Stephen Feeney. Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology. *Monthly Notices of the Royal Astronomical Society*, 2018, 477 (3), pp.2874-2885. 10.1093/mnras/sty819 . hal-01703719

HAL Id: hal-01703719

<https://hal.science/hal-01703719v1>

Submitted on 11 Nov 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology

Justin Alsing,^{1,2★} Benjamin Wandelt^{1,3,4,5} and Stephen Feeney¹

¹Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York City, NY 10010, USA

²Imperial Centre for Inference and Cosmology, Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK

³Institut d'Astrophysique de Paris (IAP), UMR 7095, CNRS UPMC Université Paris 6, Sorbonne Université, 98bis boulevard Arago, F-75014 Paris, France

⁴Institut Lagrange de Paris (ILP), Sorbonne Université, 98bis boulevard Arago, F-75014 Paris, France

⁵Department of Physics and Astronomy, University of Illinois at Urbana-Champaign, 1002 W Green St, Urbana, IL 61801, USA

Accepted 2018 March 23. Received 2018 March 7; in original form 2017 December 29

ABSTRACT

Many statistical models in cosmology can be simulated forwards but have intractable likelihood functions. Likelihood-free inference methods allow us to perform Bayesian inference from these models using only forward simulations, free from any likelihood assumptions or approximations. Likelihood-free inference generically involves simulating mock data and comparing to the observed data; this comparison in data space suffers from the curse of dimensionality and requires compression of the data to a small number of summary statistics to be tractable. In this paper, we use massive asymptotically optimal data compression to reduce the dimensionality of the data space to just one number per parameter, providing a natural and optimal framework for summary statistic choice for likelihood-free inference. Secondly, we present the first cosmological application of Density Estimation Likelihood-Free Inference (DELFI), which learns a parametrized model for joint distribution of data and parameters, yielding both the parameter posterior and the model evidence. This approach is conceptually simple, requires less tuning than traditional Approximate Bayesian Computation approaches to likelihood-free inference and can give high-fidelity posteriors from orders of magnitude fewer forward simulations. As an additional bonus, it enables parameter inference and Bayesian model comparison simultaneously. We demonstrate DELFI with massive data compression on an analysis of the joint light-curve analysis supernova data, as a simple validation case study. We show that high-fidelity posterior inference is possible for full-scale cosmological data analyses with as few as $\sim 10^4$ simulations, with substantial scope for further improvement, demonstrating the scalability of likelihood-free inference to large and complex cosmological data sets.

Key words: methods: data analysis.

1 INTRODUCTION

In cosmological data analysis, we are often faced with scenarios where we can generate mock data with sophisticated forward simulations, but are unable to write down a tractable likelihood function. For example, physics associated with non-linear structure formation on small scales (Springel 2005; Klypin, Trujillo-Gomez & Primack 2011), baryonic feedback (Hellwing et al. 2016; Springel et al. 2018; Chisari et al. 2018), gravitational and hydrodynamical evolutions of the intergalactic medium (Arinyo-i Prats et al. 2015; Bolton et al. 2016), epoch of reionization (Mesinger, Greig & Sobacchi

2016; Kern et al. 2017) etc., may be captured (to varying degrees) by simulations, whilst compact and accurate models for the statistical properties of these processes are often elusive. Similarly on the measurement side, complicated noise models, subtle measurement, and selection biases etc., can often be simulated but are challenging to incorporate exactly into a likelihood function.

The standard approach is to build an approximate likelihood that tries to capture as much of the known physics and measurement processes underlying the data as possible, in the hope that the adopted approximations do not lead to biased posterior inferences. Even if the means and variances of inferences are not appreciably biased, assessing tensions between data sets (Marshall, Rajguru & Slosar 2006), combining inferences, and comparing models can be strongly affected by posterior tail probabilities that are unlikely to

* E-mail: jalsing@flatironinstitute.org

be accurate when using popular likelihood approximations. With widely reported tensions between key state-of-the-art cosmological data sets, most notably weak-lensing and cosmic microwave background (CMB) measurements of the amplitude of matter clustering (Ade et al. 2016; Joudaki et al. 2017; Alsing, Heavens & Jaffe 2016; Hildebrandt et al. 2017) and local versus CMB measurements of the Hubble constant (Riess et al. 2011; Ade et al. 2016; Feeney, Mortlock & Dalmaso 2017), it is worthwhile seeking methods that might eliminate likelihood approximations from the chain of scientific reasoning.

Likelihood-free inference methods allow us to perform Bayesian inference using forward simulations only, free from any likelihood assumptions or approximations (see Lintusaari et al. 2017 for a review). This approach has great appeal for cosmological data analysis, since encoding complex physical processes, instrumental effects, selection biases, etc., into a forward simulation is typically much easier than incorporating these effects into a complicated likelihood function and solving the inverse problem.

Likelihood-free methods are emerging as a viable way forward for analysing complex data sets in cosmology, with recent applications to inference of the quasar luminosity function (Schafer & Freeman 2012), galaxy merger rate evolution at early times (Cameron & Pettitt 2012), cosmological parameters from supernova observations (Weyant, Schafer & Wood-Vasey 2013), galaxy formation (Robin et al. 2014), weak-lensing peak statistics (Lin & Kilbinger 2015), the galaxy–halo connection (Hahn et al. 2017), cosmological redshift distributions (Kacprzak et al. 2017), photometric evolution of galaxies (Carassou et al. 2017), and Lyman- α and $-\beta$ forests (Davies et al. 2017), with public likelihood-free inference codes implementing Approximate Bayesian Computation (ABC) facilitating the rise in popularity of these methods (Ishida et al. 2015; Akeret et al. 2015; Jennings, Wolf & Sako 2016).

In its simplest form, likelihood-free inference with ABC involves forward simulating mock data given a set of input parameters drawn from the prior, and then comparing the simulated data to the observed data, accepting parameters when the simulated data are close (by some distance metric) to the observed data. This comparison in data space suffers from the curse of dimensionality, scaling exponentially with the size of the data set; for large data sets such a comparison is completely impractical and it is essential to compress the data down to a small number of summary statistics. The need for data compression is common amongst even sophisticated likelihood-free inference methods. Data compression schemes should be carefully designed to reduce the data to the smallest set of summaries possible, whilst retaining as much information about the parameters of interest as possible (see Blum et al. 2013, for a review).

Once a data compression scheme has been prescribed, the second hurdle for achieving scalable likelihood-free inference is choosing how to propose parameters and run forward simulations in the most efficient way, minimizing the number of simulations required to obtain high-fidelity posterior inferences. This is of particular importance for applications in cosmology, where forward simulations are often extremely computationally expensive; even with very aggressive data compression, ABC methods typically require an unfeasibly large number of simulations for many cosmological applications.

This paper tackles the two key hurdles for scalable likelihood-free inference: (1) how do we compress large cosmological data sets down to a small number of summaries, whilst retaining as much information about the cosmological parameters as possible, and (2) how do we perform likelihood-free inference using a feasible number of forward simulations. We propose a general two-step data compression scheme, first compressing the full data set $\mathbf{D} \in \mathbb{R}^N$

down to $\mathbf{d} \in \mathbb{R}^M$ well-chosen heuristic summary statistics as is standard practice in cosmological data analysis (e.g. compressing maps down to power spectra, supernova light curves, and spectra down to estimated distance moduli and redshifts), and secondly asymptotically optimally¹ compressing the M summaries down to just n numbers $\mathbf{t} \in \mathbb{R}^n$ – one for each parameter of interest – while preserving the Fisher information content of the data, following Alsing & Wandelt (2018), Heavens, Jimenez & Lahav (2000), and Tegmark, Taylor & Heavens (1997).

With the two-step compression scheme defined, we then introduce Density Estimation Likelihood-Free Inference (DELFI, Bonassi, You & West 2011; Fan, Nott & Sisson 2013; Papamakarios & Murray 2016), which learns a parametrized model for the joint density of the parameters and compressed statistics $P(\boldsymbol{\theta}, \mathbf{t})$, from which we can extract the posterior density by simply evaluating the joint density at the observed data \mathbf{t}_o , i.e. $P(\boldsymbol{\theta}|\mathbf{t}_o) \propto P(\boldsymbol{\theta}, \mathbf{t} = \mathbf{t}_o)$. We will show that high-fidelity posterior inference can be achieved with orders of magnitude fewer forward simulations than, for example, available implementations of Population Monte Carlo ABC (PMC-ABC), making likelihood-free inference feasible for full-scale cosmological data analyses where simulations are expensive. As a case study, we will demonstrate DELFI with massive data compression on an analysis of the joint light-curve analysis (JLA) supernova data set (Betoule et al. 2014).

The structure of this paper is as follows: in Section 2, we discuss massive asymptotically optimal data compression for application to likelihood-free inference methods. In Section 3, we introduce likelihood-free inference methods, discussing ABC and introducing DELFI (Bonassi et al. 2011; Fan et al. 2013; Papamakarios & Murray 2016) as a scalable alternative. In Section 4, we validate the DELFI method with asymptotically optimal data compression on a simple analysis of the JLA supernova data set, and compare to PMC-ABC. We conclude in Section 5.

2 MASSIVE ASYMPTOTICALLY OPTIMAL DATA COMPRESSION

In this section, we describe the two-step data compression, first from N data to a set M of well-chosen summary statistics, and then compressing the M summaries down to just n numbers, where n is equal to the number of parameters to be inferred.

The first step is already common practice in cosmological data analysis; for example, data from cosmological surveys are typically compressed to maps and then estimated n -point statistics (power spectra or correlation functions, bispectra, etc.) or other summary statistics. The second step, compressing down to just one number per parameter whilst retaining as much (Fisher) information about the parameters as possible, has been considered by Tegmark et al. (1997), Heavens et al. (2000), and Alsing & Wandelt (2018); we follow the most general of these studies here (Alsing & Wandelt 2018). These Fisher-information preserving data compression ideas are already widely used in astronomy and cosmology, with applications spanning determining galaxy star formation histories (Reichardt, Jimenez & Heavens 2001; Heavens et al. 2004; Panter et al. 2007), CMB data analysis (Gupta & Heavens 2002; Zablacki & Dodelson 2016), gravitational waves (Graff, Hobson & Lasenby 2011), transient detection (Protopapas, Jimenez & Alcock 2005), fast covariance-matrix estimation (Heavens et al. 2017), galaxy

¹ For a discussion of precisely what is meant by optimal in this context, see Section 2.3.

power spectrum and bispectrum analyses (Gualdi et al. 2017), and optimal power spectrum estimation (Tegmark, Taylor & Heavens 1997; Bond, Jaffe & Knox 1998, 2000).

2.1 Two-step data compression

The two-step compression proceeds as follows:

(I) Compress the full data set $\mathbf{d} \in \mathbb{R}^N$ into a set of summary statistics $\mathbf{d} \in \mathbb{R}^M$, with the aim of retaining as much information as possible about the parameters of interest:

$$\mathbf{D} \rightarrow \mathbf{d}(\mathbf{D}) = \{\text{list of summary statistics}\} \quad (1)$$

(II) Compress the vector of M summary statistics \mathbf{d} into a vector of n numbers $\mathbf{t} \in \mathbb{R}^n$, as follows: assume an approximate form for the log-likelihood function \mathcal{L} , and define the compressed statistics \mathbf{t} to be the score function – the gradient of the log-likelihood – evaluated at some fiducial parameter set θ_* :

$$\mathbf{d} \rightarrow \mathbf{t} = \nabla_{\theta} \mathcal{L}_*. \quad (2)$$

In the case where we assume a Gaussian likelihood for the summary statistics, $2\mathcal{L} = -(\mathbf{d} - \boldsymbol{\mu})^T \mathbf{C}^{-1} (\mathbf{d} - \boldsymbol{\mu}) - \ln|\mathbf{C}|$ with the mean and covariance depending on the parameters, the compressed statistics \mathbf{t} are given by (Alsing & Wandelt 2018):

$$\mathbf{t} = \nabla_{\theta} \boldsymbol{\mu}_*^T \mathbf{C}_*^{-1} (\mathbf{d} - \boldsymbol{\mu}_*) + \frac{1}{2} (\mathbf{d} - \boldsymbol{\mu}_*)^T \mathbf{C}_*^{-1} \nabla_{\theta} \mathbf{C}_* \mathbf{C}_*^{-1} (\mathbf{d} - \boldsymbol{\mu}_*), \quad (3)$$

where $\boldsymbol{\mu}_* \equiv E_{\theta_*}[\mathbf{d}]$ and $\mathbf{C}_* = E_{\theta_*}[(\mathbf{d} - \boldsymbol{\mu})(\mathbf{d} - \boldsymbol{\mu})^T]$ are the mean and covariance of the summary statistics \mathbf{d} (evaluated at the fiducial parameter values), which can be estimated from forward simulations.

Let us consider the practical considerations and limitations of steps (I) and (II) in turn.

2.2 Step I: initial compression to summary statistics

Compression step (I) retains as much information as can be captured by the carefully chosen summary statistics \mathbf{d} . In special cases, where *sufficient* summary statistics can be found, this step will be lossless, but in general step (1) will be lossy. Nevertheless, it is standard practice to reduce large cosmological data sets to sets of summary statistics in the hope of retaining as much information as possible, whilst making the subsequent inference task tractable. A substantial body of literature exists on summary statistic choice for cosmological data analysis. In the context of likelihood-free inference, one may include in \mathbf{d} a list of as many relevant summary statistics for the problem as is feasible in the hope of capturing as much information as possible, without needing to be able to write down a joint likelihood function for the summaries.

Importantly, the only requirement for likelihood-free inference is that realizations of the summary statistics can be generated by forward simulation given a set of input parameters; in contrast to likelihood-based analyses, we do not require a predictive model for the expected summary statistics $\boldsymbol{\mu}(\boldsymbol{\theta}) = E_{\theta}[\mathbf{d}]$. For cosmological data analysis, this has great appeal, since many problems have summary statistics that are expected to contain a wealth of information but that we may not have a reliable predictive model for. Examples in the context of large-scale structure analyses include the galaxy power spectrum and bispectrum in redshift space and

on small scales, the weak-lensing power spectrum on small angular scales, cosmic void statistics, the flux power spectrum of the Lyman- α forest, and many others.

Whilst step (I) typically results in an enormous reduction in the size of the data space, for cosmological applications the number of summary statistics M is still often $\sim 10^2$ or larger. For example, we may have compressed a vast number of time-ordered data points from a CMB survey down to a few hundred or thousand estimated power spectrum modes, or measured supernova light curves and spectra down to an estimated apparent magnitude, redshift, colour, and stretch parameter for each of the sources. Hence, the space of M summaries is still typically much too large for practical data-space comparisons and likelihood-free inference; further compression of these summaries down to a small number of compressed statistics is still required.

2.3 Step II: asymptotically optimal compression to the score function of an approximate likelihood

Once an appropriate set of summary statistics has been chosen, the massive compression in step (II) proceeds by assuming an approximate form for the likelihood and compressing to the score function – the gradient of the log-likelihood evaluated at some fiducial parameter set θ_* (equation 2). This compression results in just n numbers – one per parameter of interest – that are optimal in the sense that they preserve the Fisher information, to the extent that the assumed form for the (unknown or intractable) likelihood is a good approximation to the true likelihood function and the fiducial expansion point is close to the maximum likelihood (this can be iterated if necessary; see Alsing & Wandelt 2018 for details). A detailed discussion of optimality in this context is given at the end of this section.

Crucially, the assumed approximate likelihood function is used for the sole purpose of performing the data compression; once the compression is done, all likelihood assumptions are dropped and the subsequent inference is genuinely likelihood-free. Better likelihood approximations for the compression step will lead to more optimally compressed statistics, but there is no sense in which these choices will bias the final parameter inferences.

For many applications, a Gaussian likelihood may be a reasonable first approximation, but not accurate enough in detail to be used for likelihood-based inferences. In these situations, a Gaussian likelihood may be assumed for the data compression, leading to compressed statistics given by equation (3). Computing the compression in equation (3) requires an estimate of the mean, covariance matrix, and their derivatives at some fiducial parameter set θ_* . These can all be obtained using forward simulations only, as is already common practice for mean and covariance estimation for conventional likelihood-based analyses. Derivatives can be estimated quickly by performing simulations with matched random seeds but perturbed parameter values.

For likelihood-based analyses assuming Gaussian likelihoods, poorly estimated covariance matrices will in general lead to biased parameter inferences and great care needs to be taken to ensure these are determined precisely, including any parameter dependencies, and any covariance-matrix uncertainties should be formally marginalized over (Sellentin & Heavens 2015). In contrast, for likelihood-free analyses, if the covariance matrix used for the compression in equation (3) is approximate, the worst outcome is that the resulting compression will be sub-optimal; crucially, there is no sense in which poorly estimated covariance matrices can bias the final parameter inferences and parameter-dependent covariance

matrices are not required. The same principle applies to the mean and derivatives appearing in equation (3); approximations are always safe, only leading to sub-optimality. This means that when fast approximate models for the covariances and means are available, they can be used safely for rapid compression, reducing the total number of forward simulations required (at the cost of some optimality).

2.4 Asymptotic optimality

Compression to the score function of a given likelihood promises to be optimal in the sense that it preserves the Fisher information content of the data, to the extent that the likelihood assumed for the compression is a good approximation to the true likelihood and the gradient (score) can be evaluated close to the maximum likelihood.

In typical likelihood-free inference applications, the form of the likelihood function under the model is either not known or not computationally tractable, and using an approximate likelihood function for the compression will be lossy. In these cases, if the compression is performed under a ‘best guess’ for the true likelihood then there is still a sense in which the compression is ‘optimal’; it preserves as much Fisher information as possible under the level of knowledge and resources available for making a good likelihood approximation for the compression. In other words, you can only do as well as your likelihood ignorance allows.

Even when the compression is performed under the exact likelihood, compression to the score only promises to preserve the Fisher information content of the data. Whilst this is a clearly stated definition of optimality, in cases where the likelihood is a highly non-Gaussian or multimodal function of the parameters the Fisher information is not guaranteed to be a good measure of the information content of the data and there may be more effective compression schemes. This is rarely the case for cosmological applications. Nevertheless, in the asymptotic limit where the likelihood becomes Gaussian with (expected) curvature specified by the Fisher information matrix, compression to the score exactly preserves the (expected) uncertainties on the inferred parameters. In this sense, compression to the score can be said to be asymptotically optimal.

We note that a new approach to data compression is emerging that does not require the compression to be performed under an approximate-likelihood function; Charnock, Lavaux & Wandelt (2018) develop information maximizing neural networks trained on forward simulations that can learn optimal compression schemes without specifying a likelihood function.

3 LIKELIHOOD-FREE INFERENCE

In this section, we discuss likelihood-free inference. In Section 3.1, we discuss ABC methods, highlighting some of their limitations in the context of cosmological data analysis. In Section 3.2, we present DELFI, which overcomes many of the key shortcomings of ABC methods.

3.1 Approximate Bayesian Computation

In its simplest incarnation, rejection ABC works as follows (Rubin et al. 1984):

- (i) Draw parameters from the prior $\theta \leftarrow P(\theta)$;
- (ii) Simulate mock data $\mathbf{d} \leftarrow P(\mathbf{d}|\theta)$;
- (iii) If distance between observed and mock data is smaller than some threshold, $\rho(\mathbf{d}, \mathbf{d}_o) < \epsilon$, accept, else reject;

- (iv) Repeat until desired number of samples are obtained.

In the limit where $\epsilon \rightarrow 0$, the accepted samples are drawn from the true posterior, whilst for any non-zero ϵ , the samples drawn are from an approximate posterior that is by construction broader than the true posterior density. The distance metric ρ for comparing simulated and observed data needs to be specified (with many options existing, McKinley, Cook & Deardon 2009), as does the distance threshold ϵ . Since the acceptance rate becomes vanishingly small as $\epsilon \rightarrow 0$, ABC posteriors are always broader than the true posterior, but are unbiased; provided one can make ϵ small enough, good posterior approximations can be recovered.

Proposing parameters from the prior in rejection ABC is typically inefficient when the posterior density occupies a small portion of the total prior volume (see e.g. Toni et al. 2009; Toni & Stumpf 2009). In this case, drawing parameters from a proposal distribution that preferentially samples the relevant portion of parameter space (followed by importance reweighting) leads to more efficient ABC sampling. PMC and Sequential Monte Carlo (SMC) ABC methods (Del Moral, Doucet & Jasra 2006; Sisson, Fan & Tanaka 2007; Beaumont et al. 2009; Toni et al. 2009; Bonassi et al. 2015) are popular advancements on rejection ABC that adaptively learn a more intelligent proposal distribution, whilst at the same time implementing a ‘cooling’ scheme for ϵ , gradually lowering the distance threshold as the proposal distribution becomes more optimized (see Ishida et al. 2015; Akeret et al. 2015; Jennings et al. 2016 for applications in the astronomy literature).

ABC methods have been applied successfully to a number of problems in cosmological data analysis (Schafer & Freeman 2012; Cameron & Pettitt 2012; Weyant et al. 2013; Robin et al. 2014; Lin & Kilbinger 2015; Hahn et al. 2017; Kacprzak et al. 2017; Davies et al. 2017). However, even sophisticated ABC algorithms suffer from vanishingly small acceptance rates as $\epsilon \rightarrow 0$ by construction, scaling poorly with the number of parameters of interest, so for high-fidelity posterior inference this usually means running very large numbers of forward simulations. For many applications in cosmology where simulations are expensive, this is impractical.

In the next section, we describe a totally different approach to likelihood-free inference that is ‘ ϵ -free’, circumventing the need to do direct comparisons in data space and ultimately making much more efficient use of forward simulations.

3.2 Density Estimation Likelihood-Free Inference

DELFI works by learning a parametrized model for the joint density $P(\theta, \mathbf{d})$, from a set of samples drawn from that density (Bonassi et al. 2011; Fan et al. 2013; Papamakarios & Murray 2016). In its simplest form, we start by generating a set of samples $\{\theta, \mathbf{d}\}$ from $P(\theta, \mathbf{d})$ by drawing parameters from the prior and forward simulating mock data:

$$\begin{aligned} \theta &\leftarrow P(\theta) \\ \mathbf{d} &\leftarrow P(\mathbf{d}|\theta). \end{aligned} \quad (4)$$

We then write down a model for the joint density $P(\theta, \mathbf{d}; \eta)$, parametrized by η , and fit this model to the samples $\{\theta, \mathbf{d}\}$. The estimated² posterior density and Bayesian evidence can then be

² Recall that widely used MCMC methods also produce *estimates* of the posterior density (and its properties) and/or the model evidence, from a set of posterior samples.

easily extracted from the fit to the joint density as follows:

$$\begin{aligned} \hat{P}(\theta | \mathbf{d}_o) &\propto P(\theta, \mathbf{d} = \mathbf{d}_o; \eta) \\ \hat{P}(\mathbf{d}_o) &= \int P(\theta, \mathbf{d} = \mathbf{d}_o; \eta) d\theta, \end{aligned} \quad (5)$$

i.e. taking a slice through the joint distribution evaluated at the observed data $\mathbf{d} = \mathbf{d}_o$, and subsequently integrating over the parameters for the Bayesian evidence. For many practical choices of parametrized models for the joint density, e.g. Gaussian mixture models (GMM, see below), the evidence integral in equation (5) is analytically tractable. This means that the evidence comes for free, and if the parametrized model for the joint density is fit to the samples in a principled way, the uncertainties on the fit parameters can be propagated through to a principled uncertainty on the estimated Bayesian evidence.

In contrast to ABC, DELFI uses all of the available forward simulations $\{\theta, \mathbf{d}\}$ to inform the inference of the joint density $P(\theta, \mathbf{d})$, and hence the posterior density and evidence estimation. In practice, this means that far fewer forward simulations may be needed to obtain high-fidelity posterior inferences (compared to ABC that has a vanishingly small acceptance rate as $\epsilon \rightarrow 0$), as demonstrated by Papamakarios & Murray (2016).

3.2.1 Gaussian mixture density estimation

In this work, we parametrize the joint density with a GMM,

$$P(\theta, \mathbf{d}; \eta) = \sum_{i=1}^K w_i \mathcal{N}(\mu_i, \mathbf{C}_i), \quad (6)$$

where $\mathcal{N}(\cdot)$ denotes the Gaussian density, and the mixture model is parametrized by the weights, means, and covariances $\eta = \{\{w\}, \{\mu\}, \{\mathbf{C}\}\}$ of each of the K components, with $\sum w = 1$. GMMs are capable of representing any probability density arbitrarily accurately, provided the number of components K is sufficiently large, and are straightforward to fit to data using expectation-maximization or other methods (see e.g. Bishop 2006). They also have the appeal that the evidence integrals appearing in equation (5) are analytically tractable, so the Bayesian evidence comes for free:

$$\hat{P}(\mathbf{d}_o) = \int P(\theta, \mathbf{d} = \mathbf{d}_o; \eta) d\theta = \sum_{i=1}^K w_i \mathcal{N}(\mu_{di}, \mathbf{C}_{ddi}), \quad (7)$$

where μ_{di} and \mathbf{C}_{ddi} are the component means and covariances corresponding to the data dimensions in the joint density.

3.2.2 Data compression

The need for data compression for DELFI is still clear: the joint density of the data and parameters has dimensionality $N + n$, which presents an intimidating density estimation task for even modestly large data sets. However, implementing the two-step compression scheme described in Section 2 means we only have to estimate the joint density of the parameters and compressed statistics $P(\mathbf{t}, \theta)$, whose dimensionality is just $2n$. For many cosmological applications, the number of parameters of interest is typically $n \lesssim 10$.

When using DELFI with a data compression scheme, samples $\{\theta, \mathbf{t}\}$ are generated from $P(\theta, \mathbf{t})$ as before by drawing parameters from the prior and forward simulating mock data, with the addition of

the subsequent compression step:

$$\begin{aligned} \theta &\leftarrow P(\theta), \\ \mathbf{d} &\leftarrow P(\mathbf{d} | \theta), \\ \mathbf{t} &= \mathbf{t}(\mathbf{d}). \end{aligned} \quad (8)$$

These samples $\{\theta, \mathbf{t}\}$ are then fitted with a mixture density model in the usual way as described above. Note that when the Bayesian evidence is estimated from the joint density $P(\theta, \mathbf{t})$, this will be the evidence for the compressed statistics $P(\mathbf{t}_o)$ and not for the original data vector $P(\mathbf{d}_o)$; whilst these are not numerically equivalent, the evidence under compressed statistics can still be readily used for model comparison purposes provided alternative models are compared under the same set of compressed summaries.

Importantly, the complexity of the inference problem stays as a $2n$ -dimensional density estimation task irrespective of the size of the data set (or the number of first-level summaries used), once the compression scheme has been prescribed. Therefore, the inference step scales easily to large data sets.

3.2.3 Implementation

We run forward simulations to generate a set of samples $\{\theta, \mathbf{t}\}$ and fit a GMM using PYGMMIS (Melchior & Goulding 2016), which uses expectation-maximization whilst properly taking into account any hard prior boundaries.

Gaussian mixture density estimation with a large number of components can fall foul to overfitting. One simple way to mitigate overfitting is to set a minimum threshold for the diagonals of the mixture component covariances (we adopt this approach). For a more sophisticated implementation that avoids overfitting without having to specify thresholds by hand, see Papamakarios & Murray (2016).

Note that when there are hard prior boundaries, the evidence integral in equation (7) is no longer analytically tractable. In these cases, one can estimate the evidence as follows: fit a GMM to the samples of $\{\mathbf{t}\}$ alone, ignoring $\{\theta\}$ (this effectively pre-marginalizes over θ). Then the evidence can be estimated by simply evaluating at the observed data, i.e. $\hat{P}(\mathbf{t} = \mathbf{t}_o)$.

3.2.4 Sophistications

Papamakarios & Murray (2016) developed a sophisticated implementation of DELFI with two key advancements on the vanilla set-up described above. First, they parametrize the joint distribution with a mixture density network (MDN) – a neural network parametrization of a GMM – which is fit to the samples using stochastic variational inference (SVI; see Bishop 2006 for a review of MDN and Hoffman et al. 2013 for SVI methods). Secondly, rather than drawing samples from the prior, they adaptively learn a proposal distribution that preferentially samples regions of high posterior density, and subsequently importance reweight the samples (in the same spirit as PMC-ABC methods). They find that this set-up is highly resistant to overfitting even for small numbers of samples, enabling the number of forward simulations to be reduced further. We leave implementation of these sophistications to future work.

3.2.5 Scaling with number of parameters and dealing with nuisances

With the compression scheme employed, the inference task is reduced to learning a $2n$ -dimensional density from a set of forward

simulations, irrespective of the size of the data set. The complexity of the inference problem will increase with the number of parameters n .

For typical cosmological applications, the number of parameters of interest θ (i.e. the cosmological parameters) will be $\lesssim 10$. However, in many situations, there will be additionally a number of nuisance parameters ξ , capturing observational and astrophysical systematics, selection effects, etc., which need inferring simultaneously and subsequently marginalizing over. If there are n parameters of interest and m nuisance parameters, DELFI involves learning a $2(n + m)$ -dimensional probability density over the parameters, nuisances, and their respective compressed summaries, $P(\theta, \xi, \mathbf{t}_\theta, \mathbf{t}_\xi)$. However, if the goal is the posterior marginalized over the nuisance parameters, it may be possible to keep the complexity of the inference task as a $2n$ -dimensional density estimation problem. This can be achieved by choosing compressed summaries $\mathbf{t}_\theta \in \mathbb{R}^n$ that contain as much information as possible about the parameters of interest, whilst being as insensitive as possible to the nuisance parameters. Then, draw samples $\{\theta, \xi, \mathbf{t}_\theta\}$ by drawing from the prior and forward simulating as usual, but only attempt to fit the density $P(\theta, \mathbf{t}_\theta)$, i.e. a $2n$ -dimensional density pre-marginalized over the nuisances.

Data compression for marginalized parameter subsets (under Gaussian likelihoods) is treated in Zablocki & Dodelson (2016). We leave a more general treatment of optimal compression in the presence of nuisances and learning the nuisance-marginalized posterior density to future work.

4 VALIDATION CASE STUDY: JLA SUPERNOVA DATA ANALYSIS

To demonstrate the use of DELFI (Section 3) with massive optimal data compression (Section 2), in this section, we perform an analysis of the JLA supernova data set (Betoule et al. 2014). For the purposes of validating the likelihood-free approach, we perform a simple analysis assuming a Gaussian likelihood for the JLA data so that we can compare to an exact likelihood-based analysis, allowing us to demonstrate the fidelity of the likelihood-free posterior inference against a ground truth.³ We will compare DELFI and PMC-ABC against a long-run Markov Chain Monte Carlo (MCMC) analysis of the exact (assumed) posterior distribution.

In Sections 4.1–4.3, we describe the JLA data, model, and Gaussian likelihood assumptions under which we validate the likelihood-free approach. In Section 4.6, we discuss the implementation of the likelihood-free inference, and in Section 4.7, we show the results.

4.1 JLA supernova data

We use the JLA sample comprised of observations of 740 type Ia supernovae, as analysed in Betoule et al. (2014). The sample is a compilation of supernova observations from a number of surveys – see Betoule et al. (2014) and references therein for details.

The full data set comprises multicolour light curves and spectroscopic (or sometimes photometric) observations of each supernova. These light curves and spectra are then used to estimate apparent magnitudes m_B and redshifts z , as well as colour at maximum-brightness C and stretch X_1 parameters characterizing the light

curves (see e.g. Tripp 1998). In the data analysis (see Section 4.3, also Betoule et al. 2014), the data vector will be assumed to be the vector of estimated apparent magnitudes $\mathbf{d} = (\hat{m}_B^1, \hat{m}_B^2, \dots, \hat{m}_B^M)$, where uncertainties in the redshift, colour, and stretch parameters are propagated through to the covariance matrix of the observed apparent magnitudes. This compression of the multicolour light curves and spectra down to a set of estimated apparent magnitudes can be thought of as step (I) of the data compression described in Section 2.

4.2 Λ CDM and light-curve calibration model

As standardizable candles, we assume that the apparent magnitudes of type Ia supernovae depend on the luminosity distance to the source at a given redshift $D_L^*(z)$ (which is a function of the cosmological model and parameters), a reference absolute magnitude for type Ia supernovae (as a function of host-galaxy mass), and calibration corrections for the light-curve stretch X_1 and colour at maximum-brightness C ,

$$m_B = 5 \log_{10} \left[\frac{D_L^*(z)}{10 \text{pc}} \right] + \tilde{M}_B(M_{\text{stellar}}; M_B, \delta M) - \alpha X_1 + \beta C \quad (9)$$

where α and β are calibration parameters for the stretch and colour, respectively. The absolute magnitude \tilde{M}_B is assumed to be dependent on the properties of the host galaxy; following Betoule et al. (2014), we model the dependence of the reference absolute magnitude on the stellar mass of the host as $\tilde{M}_B = M_B + \delta M \Theta(M_{\text{stellar}} - 10^{10} M_\odot)$, where Θ is the Heaviside function.

The cosmological model enters in the luminosity distance–redshift relation. We will assume a flat universe with cold dark matter and dark energy characterized by equation-of-state $p/\rho = w_0$ (hereafter, Λ CDM). In a Λ CDM universe, the luminosity distance is given by,

$$D_L^*(z) = \frac{(1+z)c}{100} \int_0^z \frac{dz'}{\sqrt{\Omega_m(1+z')^3 + (1-\Omega_m)(1+z')^{3(w_0+1)}}, \quad (10)$$

where Ω_m is the matter density parameter, c is the speed of light (in vacuum), and w_0 is the equation-of-state of dark energy.

The resulting Λ CDM model with colour and stretch calibration and host-mass dependent absolute magnitude has six free parameters of interest: $\theta = (\Omega_m, w_0, \alpha, \beta, M_B, \delta M)$.

4.3 Likelihood

Following Betoule et al. (2014), for this validation case, we will assume the data $\mathbf{d} = (\hat{m}_B^1, \hat{m}_B^2, \dots, \hat{m}_B^M)$ are Gaussian distributed,

$$\mathcal{L} = -\frac{1}{2}(\mathbf{d} - \boldsymbol{\mu}(\theta))^T \mathbf{C}^{-1}(\mathbf{d} - \boldsymbol{\mu}(\theta)) - \frac{1}{2} \ln |\mathbf{C}|, \quad (11)$$

where the mean depends on the parameters and is given by equation (9), and we will assume a fixed covariance matrix,⁴ shown in Fig. 1, that is assumed to have already accounted for the uncertainties in

⁴ Betoule et al. (2014) constructed a covariance matrix that depends on α and β , and also dropped the $|\mathbf{C}|$ term from the likelihood (equation 11) following March et al. (2011). However, since α and β are very well constrained by the data, the covariance dependence has a small impact on the final parameter inference. For this study, we compute the covariance described in Betoule et al. (2014), but with α and β fixed to their maximum likelihood values: $\alpha = 0.1257$ and $\beta = 2.644$. This also avoids issues arising from dropping the $|\mathbf{C}|$ term from the Gaussian likelihood.

³ Note that this simple validation case study is for method-validation purposes only. It is not intended to incorporate new physics or systematics over-and-above the standard JLA analysis of Betoule et al. (2014).

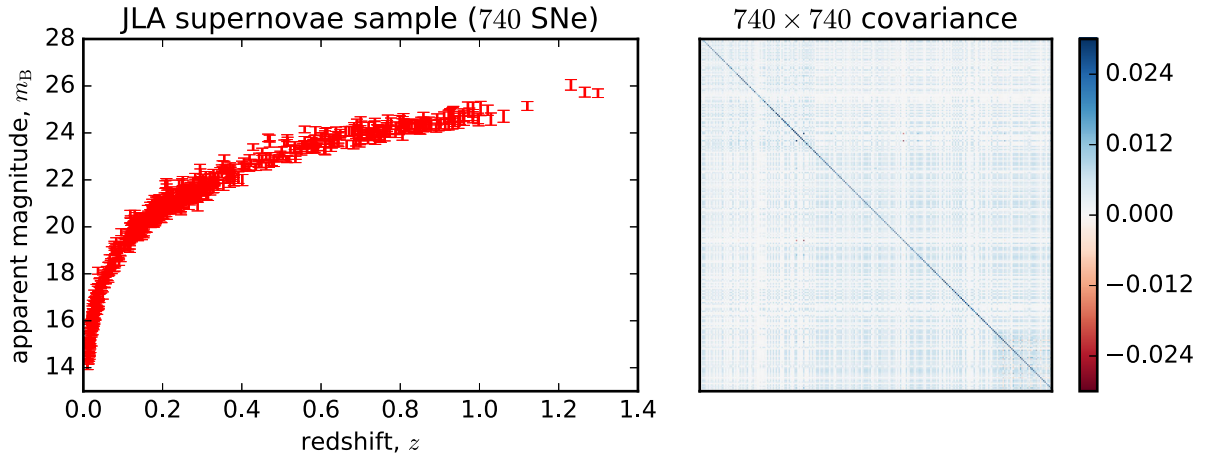


Figure 1. Left: measured apparent magnitudes with their associated uncertainties (from the diagonal of the covariance matrix) for the sample of 740 supernovae in the JLA sample. Right: covariance matrix of the measured apparent magnitudes, having taken into account redshift and light-curve calibration uncertainties – see Betoule et al. (2014) for details of the covariance-matrix construction.

the colour, stretch, and redshift of each measured supernova (see Betoule et al. 2014 for details of the covariance matrix construction).

4.4 Priors

We assume broad Gaussian priors on the parameters $\theta = (\Omega_m, w_0, \alpha, \beta, M_B, \delta M)$ with the following mean and covariance:

$$\mu_p = (0.3, -0.75, -19.05, 0.125, 2.6, -0.05),$$

$$\mathbf{C}_p = \begin{pmatrix} 0.4^2 & -0.24 & 0 & 0 & 0 & 0 \\ -0.24 & 0.75^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.1^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.025^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.25^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.05^2 \end{pmatrix}. \quad (12)$$

In addition to the Gaussian prior, we impose hard prior boundaries on $\Omega_m \in [0, 0.6]$ and $w_0 \in [-1.5, 0]$. The (truncated) Gaussian prior is much broader than the resulting posterior, having a negligible impact on the posterior inference relative to (infinite) uniform priors.

The correlations in the Gaussian prior are chosen to roughly follow the correlation structure of the inverse Fisher matrix for the parameters; this allows us to form a broad, weakly informative prior whilst improving the volume ratio of the posterior and prior (i.e. giving low prior weight to regions of parameter space that are anticipated to be strongly disfavoured by the likelihood, based on the Fisher matrix). We find this has negligible impact on the parameter inferences whilst improving the performance of the likelihood-free inference methods.⁵

4.5 Massive asymptotically optimal compression

For step (II) of the data compression, from $N = 740$ estimated apparent magnitudes down to $n = 6$ numbers (one per parameter,

following Section 2), we assume a Gaussian likelihood as in equation (11) where only the mean depends on the parameters. In this case, following equation (3) (Alsing & Wandelt 2018; Heavens et al. 2000), the optimally compressed statistics are given by:⁶

$$\mathbf{t} = \nabla_{\theta} \mu_*^T \mathbf{C}_*^{-1} (\mathbf{d} - \mu_*) + \mathbf{C}_p^{-1} (\mu_p - \theta_*), \quad (13)$$

where the mean, its derivative, and the covariance matrix are evaluated at some fiducial point θ_* and the second term includes the impact of the Gaussian prior.⁷ To choose an optimal fiducial parameter set for the compression, we iterate the parameters using the Fisher scoring method (Alsing & Wandelt 2018):

$$\theta_{k+1} = \theta_k + \mathbf{F}_k^{-1} \mathbf{t}_k, \quad (14)$$

where $\mathbf{F} = \nabla \mu^T \mathbf{C}^{-1} \nabla \mu$ is the Fisher information matrix (Tegmark et al. 1997), and \mathbf{t}_k are the compressed statistics computed about the fiducial point θ_k . We find that equation (14) converges very quickly and gives the following expansion point: $\theta_* = (0.202, -0.748, -19.04, 0.126, 2.644, -0.0525)$.

The derivatives of the mean with respect to the parameters can be written down analytically for α , β , M_B , and δM (see equation 9), whilst we use a simple leap-frog approximation for the derivatives with respect to Ω_m and w_0 .

4.6 Likelihood-free inference implementation

For this validation case, we are assuming that the data are Gaussian distributed, with mean given by equation (9) and fixed covariance as shown in Fig. 1. Forward simulating data realizations given parameters is hence as simple as drawing Gaussian random variates, and samples from the joint data-parameter density $\{\theta, \mathbf{t}\}$ are generated as follows:

$$\theta \leftarrow \mathcal{N}(\mu_p, \mathbf{C}_p),$$

$$\mathbf{d} = \mu(\theta) + \mathbf{n}, \text{ where } \mathbf{n} \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{C}_*)$$

⁶ Note that under the assumptions of Gaussian data where the only parameter dependence is in the mean, the compressed statistics are equivalent to the MOPED linear data compression of Heavens et al. (2000).

⁷ This is a minor extension of the derivation in Alsing & Wandelt (2018), replacing the log-likelihood in $\mathbf{t} = \nabla_{\theta} \mathcal{L}$ with the sum of the log-likelihood and log prior, to incorporate the impact of the prior into the optimal compression.

⁵ Note that PMC ABC and similar PMC approaches to DELFI (Papamakarios & Murray 2016) will be less sensitive to the posterior-prior volume ratio, since they adaptively learn a proposal density rather than blindly proposing parameters from the prior.

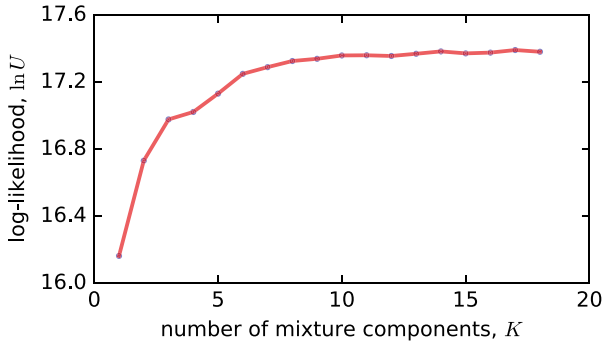


Figure 2. Log-likelihood of the samples $\{t, \theta\}$ under the GMM fits to the joint density $P(t, \theta)$, as a function of the number of Gaussian mixture components K . The log-likelihood converges with respect to the number of mixture components, where regularization of the mixture component covariances protects against overfitting (see Section 4.6).

$$t = \nabla_{\theta} \mu_*^T C_*^{-1} (d - \mu_*) + C_p^{-1} (\mu_p - \theta_*), \quad (15)$$

We generated a set of 20000 samples $\{\theta, t\}$ from the joint distribution and fit them with GMM. The GMM fits are performed using expectation-maximization (implemented using `PYGMMIS`; Melchior & Goulding 2016) as described in Section 3.2, with $K = 1$ through to 18 mixture components. The mixture component covariances are regularized by setting a minimum threshold value of 10^{-6} for the diagonal values, to avoid overfitting. We assess convergence with respect to the number of mixture components by looking at the total log-likelihood of the samples under the GMM, as a function of the number of mixture components (see Fig. 2, discussion in Section 4.7). Convergence with respect to the number of samples fed to the GMM is assessed by looking for convergence in the recovered posterior means and covariances (see Fig. 3, discussion in Section 4.7).

Since the prior in this case has hard boundaries, the Bayesian evidence integral from the GMM (equation 7) is no longer analytically tractable. We estimate the evidence by fitting a GMM directly to the samples $\{t\}$ to obtain an estimate of the density $P(t)$, and then evaluating at the observed data to estimate the evidence $P(t_o)$. We use the same GMM set-up (i.e. number of components) for the evidence as was used for the final posterior inference (discussion in Section 4.7).

4.7 Results

To assess convergence with respect to the number of mixture components, Fig. 2 shows the log-likelihood of the samples under the GMM model fits to the joint density, as a function of the number of mixture components K (using the full 20000 samples; see below). The log-likelihood of the samples clearly converges with the number of components, reaching a point where adding more components does not improve the fit further. The regularization of the mixture-component covariances has ostensibly protected against overfitting successfully (cf., Section 4.6). We use the $K = 12$ component GMM model moving forward, which appears to be well into the regime where the log-likelihood of the fit has converged.

Fig. 3 shows convergence of the posterior means and standard deviations (for each parameter) as a function of the number of samples N that are fed to the GMM model fit to the joint density. The posterior means and standard deviations converge after $N \approx 8000$, with some residual scatter of $\lesssim 0.05\sigma$ for each parameter in the means, and a few per cent in the standard deviations. This gives confidence that a reasonable posterior approximation may be obtained from just a few thousand forward simulations. However, if we are interested in high-fidelity posterior inference capturing the detailed shape of the posterior, a most sophisticated convergence diagnostic may be required and we leave this to future work. We proceed using the GMM fit to 20000 forward simulations to ensure excellent convergence, but note that in practice perfectly adequate posterior inferences may be made with far fewer simulations.

Fig. 4 shows the GMM fit (blue) to the samples (red). The mixture model with 12 components is a remarkably good representation of the samples, successfully capturing the substantial and varied non-Gaussianity, degeneracies between parameters and hard prior boundaries.

Fig. 5 shows the posterior reconstruction from a long MCMC run on the exact posterior⁸ (red), and the posterior recovered from DELFI (blue), following equation (5) applied to the GMM fit shown in Fig. 4. The posterior recovered from DELFI, using only 20000 forward simulations, is an excellent representation of the true posterior. This is not surprising given the fidelity of the GMM fit to the joint data-parameter density (Fig. 4). Crucially, this also demonstrates that

⁸ 10^6 posterior samples drawn using the affine-invariant MCMC code `EMCEE` (Foreman-Mackey et al. 2013).

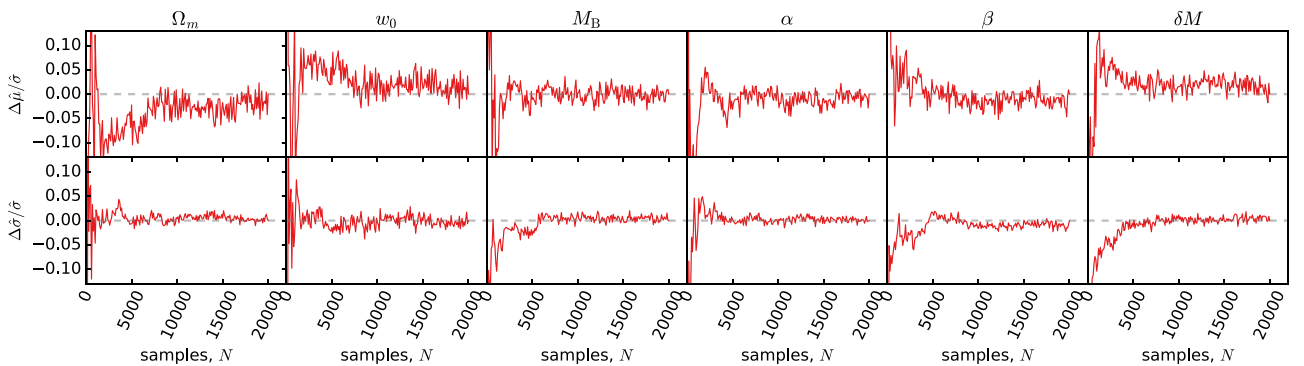


Figure 3. Top row: convergence in the estimated posterior mean as a function of the number of forward simulations fed to the GMM fit to the joint density $\hat{P}(\theta, t)$. The panels show $\Delta \hat{\mu} / \hat{\sigma} = (\hat{\mu}_N - \hat{\mu}_{N=20000}) / \hat{\sigma}_{N=20000}$, where $\hat{\mu}_N$ and $\hat{\sigma}_N$ are the estimated posterior mean and standard deviation from a Gaussian mixture fit to N forward simulated samples $\{t, \theta\}$. The posterior mean converges after a few thousand simulations, with some residual scatter of $\lesssim 0.05\sigma$ for each parameter. Bottom row: similarly, convergence of the estimated posterior standard deviation for each parameter as a function of the number of simulations fed to the GMM. The standard deviations also converge after a few thousand forward simulations, with some residual scatter at the level of a few per cent. Much of the residual scatter in the posterior means and standard deviations is due to small residual GMM fitting uncertainties.




Figure 4. Samples from the joint density $P(t, \theta)$ (red) against the GMM fit to the samples (blue) for a $K = 12$ component mixture model. The GMM model fit provides an excellent representation of the joint density $P(t, \theta)$.

the massive optimal compression step – compressing the 740 supernovae apparent magnitudes down to just six numbers – is for all intents and purposes, lossless; we find no perceptible increase in width of the DELFI posteriors computed from the compressed statistics compared to the exact (MCMC sampled) posterior computed on the full data set. In this validation case, since we are comparing to a ground truth with an assumed likelihood function, the compression is indeed optimal (and in this case, effectively lossless).

For comparison against the state of the art in ABC, in Fig. 6 we show the recovered posterior from a long PMC-ABC run.⁹ The

⁹ The PMC-ABC implementation used here follows the algorithm described in Ishida et al. (2015). We also tested modified PMC-ABC algorithms following Jennings et al. (2016), Akeret et al. (2015), and Bonassi et al. (2015); our conclusions are unchanged by these modifications to the PMC-ABC set-up.

PMC-ABC was run through 14 population iterations, generating 20 000 accepted samples in the final iteration. This required $> 2 \times 10^6$ forward simulations in total, since the vast majority of samples are rejected in the PMC-ABC approach. The posterior approximation obtained from the final set of samples is shown (blue) against the long run MCMC chain (red). The PMC-ABC run yields a reasonable approximation to the true posterior, which is unbiased but broader than the exact posterior, as expected for ABC methods. The massive optimal data compression has enabled us to successfully perform ABC, which would have been unfeasible in the full data space. However, in comparison to DELFI (Fig. 5), PMC ABC required $\sim 10^6$ forward simulations compared to $\sim 10^4$ for DELFI, for a poorer approximation to the true posterior in the end. We conclude that whilst current implementations of ABC have limited applicability for scalable cosmological data analyses where forward simulations are expensive,




Figure 5. The DELFI posterior (blue) obtained from only 20 000 forward simulations matches the exact posterior (red) obtained from a long-run MCMC chain.

DELFI allows us to perform scalable likelihood-free Bayesian inference with a reasonable number of forward simulations (with scope for further improvement).

Finally, we estimate the Bayesian evidence using DELFI. We separately fit a 12 component GMM to the 20 000 samples $\{t\}$ (neglecting the θ samples), and evaluate this estimated density at the observed data t_0 to obtain the evidence. This gives an evidence estimate of $\ln P(t_0) = 7.38$. In this validation case, we can compare to the evidence estimated directly from the known likelihood using nested sampling (MULTINEST; Feroz & Hobson 2008). Using MULTINEST, we find $\ln P(t_0) = 7.4(1)$, so the evidence estimates from DELFI and MULTINEST are in excellent agreement.

5 CONCLUSIONS

Likelihood-free inference methods allow us to perform Bayesian inference using forward simulations only, with no reference to a likelihood function. This is of particular appeal for cosmological data analysis problems where complex physical processes and in-

strumental effects can often be simulated, but incorporating them into a likelihood function and solving the inverse inference problem is much harder.

Likelihood-free methods generically require large data sets to be compressed down to a small number of summary statistics in order to be scalable. We have developed a two-step compression scheme that has widespread applicability for cosmological data analysis problems. First, we compress the data down to a list of summary statistics that are carefully chosen to contain as much information about the parameters as possible, e.g. compressing galaxy survey data to power spectra or other summary statistics. This type of compression is already standard practice in the analysis of cosmological surveys. Secondly, we use the optimal data compression scheme of Alsing & Wandelt (2018) (following earlier work by Tegmark et al. 1997; Heavens et al. 2000) to compress the list of summary statistics down to just n numbers – one per parameter – whilst preserving the Fisher information with respect to the parameters of interest. This second compression step requires the assumption of an approximate likelihood function, and will be optimal to the extent that this is a

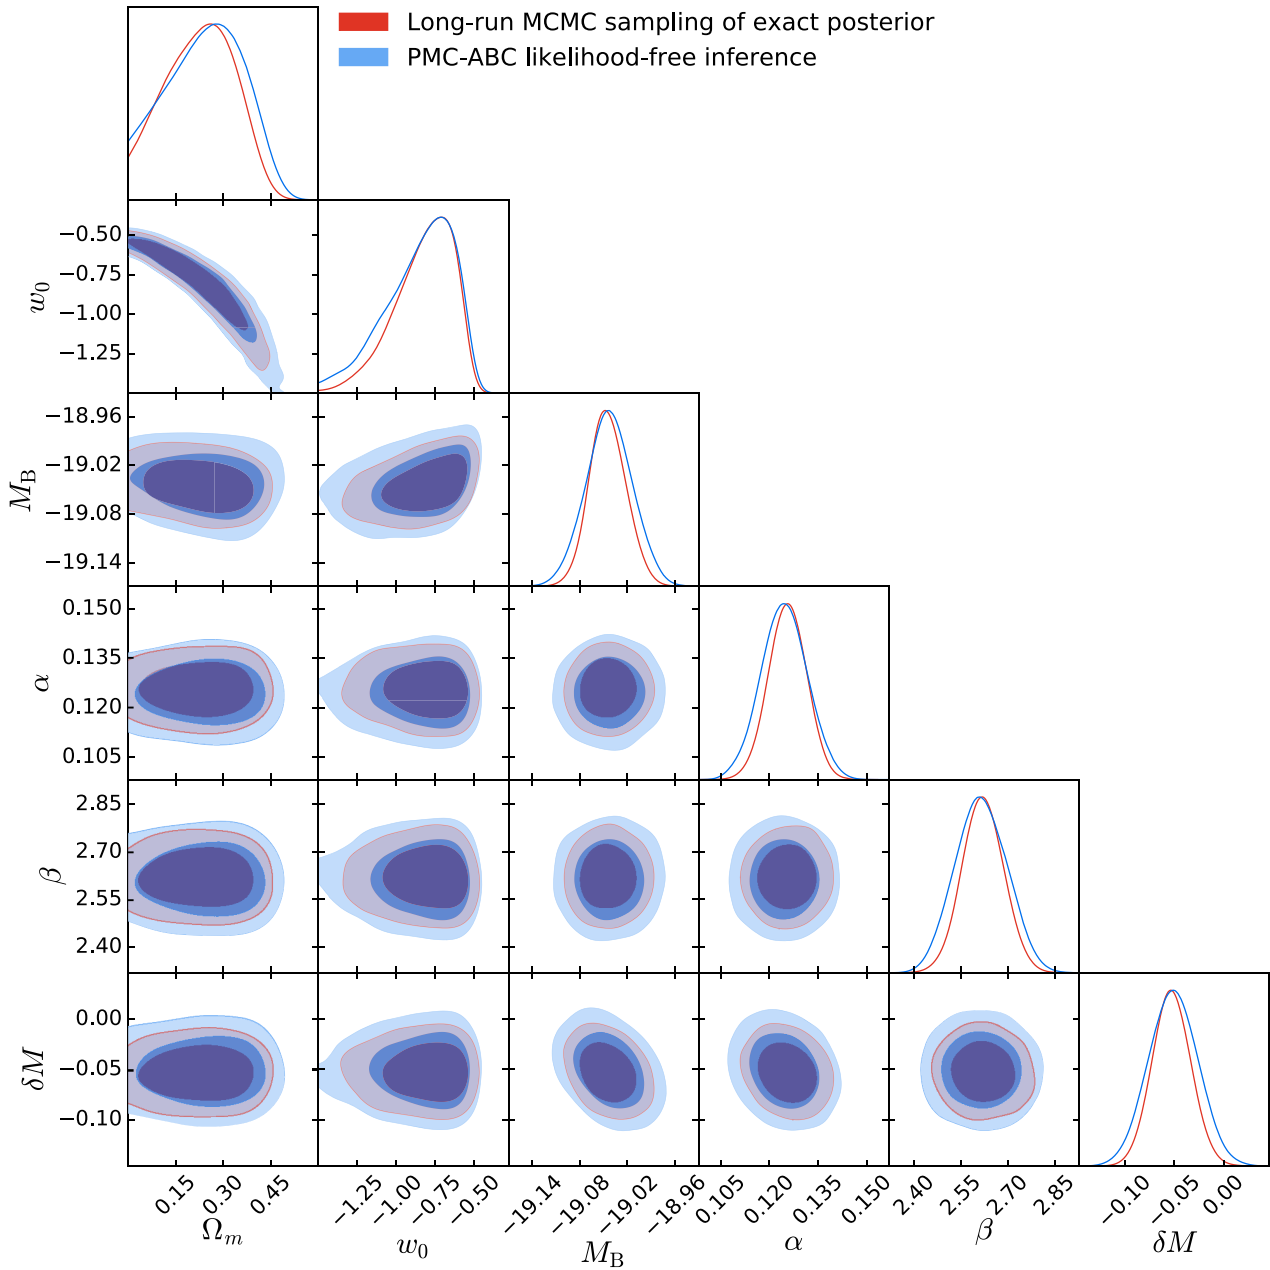


Figure 6. The PMC-ABC posterior (blue) based on 20 000 accepted samples after $>2 \times 10^6$ forward simulations bounds but does not tightly approximate the exact posterior (red) obtained from a long-run MCMC chain.

reasonable approximation to the true (unknown) likelihood. Once the data have been compressed, all subsequent likelihood-free inference based on the massively compressed statistics is genuinely likelihood-free.

ABC approaches to likelihood-free inference draw parameters from the prior and forward simulate mock data, accepting points when the simulated data fall inside some ϵ -ball around the observed data. This generates samples from an approximate posterior density that becomes exact in the limit $\epsilon \rightarrow 0$. However, ABC methods suffer from vanishingly small acceptance rates as $\epsilon \rightarrow 0$, leading to either the need for unfeasibly large numbers of forward simulations, or poor approximations to the posterior (with undesirably large values of ϵ), or both.

We have presented a new approach to likelihood-free inference for cosmology – DELFI (Bonassi et al. 2011; Fan et al. 2013; Papamakarios & Murray 2016) – that involves learning a parametrized model for the joint data-parameter probability density, from which (analytical models for) the posterior density and Bayesian evidence can be straightforwardly extracted. We have shown that when combined with the massive two-step data compression scheme, DELFI is able to recover high-fidelity posterior inferences for full-scale cosmological data analyses from $\sim 10^4$ forward simulations (for a six-parameter inference task), with scope for further improvement. In contrast, current implementations of ABC methods require orders of magnitude more forward simulations for approximate posterior inferences.

Together, massive data compression and DELFI provide a framework for performing scalable likelihood-free inference from large cosmological data sets, even when forward simulations are computationally expensive. This opens the door to a new paradigm for principled, simulation-based Bayesian inference in cosmology.

ACKNOWLEDGEMENTS

We thank Dan Foreman-Mackey and Boris Leistedt for useful discussions. This work is supported by the Simons Foundation. Benjamin Wandelt acknowledges support by the Labex Institut Lagrange de Paris (reference ANR-10-LABX-63) part of the Idex SUPER, and received financial state aid managed by the Agence Nationale de la Recherche, as part of the programme Investissements d'avenir under the reference ANR-11-IDEX-0004-02.

REFERENCES

- Ade P. A. et al., 2016, *A&A*, 594, A13
- Akeret J., Refregier A., Amara A., Seehars S., Hasner C., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 043
- Alsing J., Wandelt B., 2018, *MNRAS*, 476, L60
- Alsing J., Heavens A., Jaffe A. H., 2016, *MNRAS*, 466, 3272
- Arinyo-i Prats A., Miralda-Escudé J., Viel M., Cen R., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 017
- Beaumont M. A., Cornuet J.-M., Marin J.-M., Robert C. P., 2009, *Biometrika*, 96, 983
- Betoule M. et al., 2014, *A&A*, 568, A22
- Bishop C. M., 2006, *Pattern Recognition and Machine Learning*. Springer, New York
- Blum M. G. et al., 2013, *Stat. Sci.*, 28, 189
- Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2016, *MNRAS*, 464, 897
- Bonassi F. V., You L., West M., 2011, *Stat. App. Genetics Mol. Biol.*, 10, 49
- Bonassi F. V. et al., 2015, *Bayesian Anal.*, 10, 171
- Bond J., Jaffe A. H., Knox L., 1998, *Phys. Rev. D*, 57, 2117
- Bond J., Jaffe A. H., Knox L., 2000, *ApJ*, 533, 19
- Cameron E., Pettitt A., 2012, *MNRAS*, 425, 44
- Carassou S., de Lapparent V., Bertin E., Borgne D. L., 2017, *A&A*, 605, A9
- Charnock T., Lavaux G., Wandelt B. D., 2018, preprint ([arXiv:1802.03537](https://arxiv.org/abs/1802.03537))
- Chisari N. E. et al., 2018, preprint ([arXiv:1801.08559](https://arxiv.org/abs/1801.08559))
- Davies F. B., Hennawi J. F., Eilers A.-C., Lukić Z., 2017, *ApJ*, 855, 106
- Del Moral P., Doucet A., Jasra A., 2006, *J. R. Stat. Soc.: B (Stat. Methodol.)*, 68, 411
- Fan Y., Nott D. J., Sisson S. A., 2013, *Statistics*, 2, 34
- Feeney S. M., Mortlock D. J., Dalmasso N., 2017, *MNRAS*, 442, 3861
- Feroz F., Hobson M., 2008, *MNRAS*, 384, 449
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Graff P., Hobson M. P., Lasenby A., 2011, *MNRAS*, 413, L66
- Gualdi D., Manera M., Joachimi B., Lahav O., 2017, *MNRAS*, 476, 4045
- Gupta S., Heavens A. F., 2002, *MNRAS*, 334, 167
- Hahn C., Vakili M., Walsh K., Hearin A. P., Hogg D. W., Campbell D., 2017, *MNRAS*, 469, 2791
- Heavens A. F., Jimenez R., Lahav O., 2000, *MNRAS*, 317, 965
- Heavens A., Panter B., Jimenez R., Dunlop J., 2004, *Nature*, 428, 625
- Heavens A. F., Sellentin E., de Mijolla D., Vianello A., 2017, *MNRAS*, 472, 4244
- Hellwing W. A., Schaller M., Frenk C. S., Theuns T., Schaye J., Bower R. G., Crain R. A., 2016, *MNRAS*, 461, L11
- Hildebrandt H. et al., 2017, *MNRAS*, 465, 1454
- Hoffman M. D., Blei D. M., Wang C., Paisley J., 2013, *J. Mach. Learn. Res.*, 14, 1303
- Ishida E. et al., 2015, *Astron. Comput.*, 13, 1
- Jennings E., Wolf R., Sako M., 2016, preprint ([arXiv:1611.03087](https://arxiv.org/abs/1611.03087))
- Joudaki S. et al., 2017, *MNRAS*, 465, 2033
- Kacprzak T., Herbel J., Amara A., Réfrégier A., 2017, *J. Cosmol. Astropart. Phys.*, 2, 042
- Kern N. S., Liu A., Parsons A. R., Mesinger A., Greig B., 2017, *ApJ*, 848, 23
- Klypin A. A., Trujillo-Gomez S., Primack J., 2011, *ApJ*, 740, 102
- Lin C.-A., Kilbinger M., 2015, *A&A*, 583, A70
- Lintusaari J., Gutmann M. U., Dutta R., Kaski S., Corander J., 2017, *Syst. Biol.*, 66, e66
- March M., Trotta R., Berkes P., Starkman G., Vaudrevange P., 2011, *MNRAS*, 418, 2308
- Marshall P., Rajguru N., Slosar A., 2006, *Phys. Rev. D*, 73, 067302
- McKinley T., Cook A. R., Deardon R., 2009, *Int. J. Biostat.*, 5
- Melchior P., Goulding A. D., 2016, preprint ([arXiv:1611.05806](https://arxiv.org/abs/1611.05806))
- Mesinger A., Greig B., Sobacchi E., 2016, *MNRAS*, 459, 2342
- Panther B., Jimenez R., Heavens A. F., Charlot S., 2007, *MNRAS*, 378, 1550
- Papamakarios G., Murray I., 2016, *Advances in Neural Information Processing Systems*. p. 1028
- Protopapas P., Jimenez R., Alcock C., 2005, *MNRAS*, 362, 460
- Reichardt C., Jimenez R., Heavens A. F., 2001, *MNRAS*, 327, 849
- Riess A. G. et al., 2011, *ApJ*, 730, 119
- Robin A., Reylé C., Fliri J., Czekaj M., Robert C., Martins A., 2014, *A&A*, 569, A13
- Rubin D. B. et al., 1984, *Ann. Stat.*, 12, 1151
- Schafer C. M., Freeman P. E., 2012, in *Statistical Challenges in Modern Astronomy V*. Springer, p. 3
- Sellentin E., Heavens A. F., 2015, *MNRAS*, 456, L132
- Sisson S. A., Fan Y., Tanaka M. M., 2007, *Proc. Natl. Acad. Sci.*, 104, 1760
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V. et al., 2018, *MNRAS*, 475, 676
- Tegmark M., Taylor A. N., Heavens A. F., 1997, *ApJ*, 480, 22
- Toni T., Stumpf M. P., 2009, *Bioinformatics*, 26, 104
- Toni T., Welch D., Strelkowa N., Ipsen A., Stumpf M. P., 2009, *J. R. Soc. Interf.*, 6, 187
- Tripp R., 1998, *A&A*, 331, 815
- Weyant A., Schafer C., Wood-Vasey W. M., 2013, *ApJ*, 764, 116
- Zablocki A., Dodelson S., 2016, *Phys. Rev. D*, 93, 083525

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