

THE FUTURE OF COSMOLOGICAL LIKELIHOOD-BASED INFERENCE: ACCELERATED HIGH-DIMENSIONAL PARAMETER ESTIMATION AND MODEL COMPARISON

DAVIDE PIRAS^{*†1,2}, ALICJA POLANSKA^{†3}, ALESSIO SPURIO MANCINI^{4,3}, MATTHEW A. PRICE³, JASON D. MCEWEN^{3,5}

¹ Centre Universitaire d'Informatique, Université de Genève, 7 route de Drize, 1227 Genève, Switzerland

² Département de Physique Théorique, Université de Genève, 24 quai Ernest Ansermet, 1211 Genève 4, Switzerland

³ Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey, RH5 6NT, UK

⁴ Department of Physics, Royal Holloway, University of London, Egham Hill, Egham, UK and

⁵ Alan Turing Institute, London, NW1 2DB, UK

(Dated: 22nd May 2024)

submitted XXX; accepted YYY

Abstract

We advocate for a new paradigm of cosmological likelihood-based inference, leveraging recent developments in machine learning and its underlying technology, to accelerate Bayesian inference in high-dimensional settings. Specifically, we combine (i) *emulation*, where a machine learning model is trained to mimic cosmological observables, e.g. `CosmoPower-JAX`; (ii) *differentiable and probabilistic programming*, e.g. `JAX` and `NumPyro`, respectively; (iii) *scalable Markov chain Monte Carlo (MCMC) sampling* techniques that exploit gradients, e.g. Hamiltonian Monte Carlo; and (iv) *decoupled and scalable Bayesian model selection* techniques that compute the Bayesian evidence purely from posterior samples, e.g. the learned harmonic mean implemented in `harmonic`. This paradigm allows us to carry out a complete Bayesian analysis, including both parameter estimation and model selection, in a fraction of the time of traditional approaches. First, we demonstrate the application of this paradigm on a simulated cosmic shear analysis for a Stage IV survey in 37- and 39-dimensional parameter spaces, comparing Λ CDM and a dynamical dark energy model (w_0w_a CDM). We recover posterior contours and evidence estimates that are in excellent agreement with those computed by the traditional nested sampling approach while reducing the computational cost from 8 months on 48 CPU cores to 2 days on 12 GPUs. Second, we consider a joint analysis between three simulated next-generation surveys, each performing a 3x2pt analysis, resulting in 157- and 159-dimensional parameter spaces. Standard nested sampling techniques are simply not feasible in this high-dimensional setting, requiring a projected 12 years of compute time on 48 CPU cores; on the other hand, the proposed approach only requires 8 days of compute time on 24 GPUs. All packages used in our analyses are publicly available.

1. INTRODUCTION

The evolution of cosmological likelihood-based data analysis is heading towards a high-dimensional future. Fueled by the acquisition of ever more constraining observational data thanks to ongoing and upcoming surveys like *Euclid* (Laureijs et al. 2011)¹, the Dark Energy Spectroscopic Instrument (DESI, Levi et al. 2019)², the Nancy Grace Roman Space Telescope (Spergel et al. 2015)³, the Simons Observatory (Ade et al. 2019)⁴ and the Vera Rubin Observatory (Ivezić et al. 2019)⁵, this trajectory simultaneously entails more stringent accuracy requirements and a subsequently higher number of parameters to describe various systematic effects. At the same time, the Λ CDM model of cosmology is put under extreme pressure, with ever more complex theoretical models being developed to explain tensions in the values of cosmological parameters, thus growing the size of the parameter space. While these developments promise significant advancements in our understanding of the Universe, they present a formidable challenge for Bayesian

inference, stretching the capabilities of traditional inference methods.

We advocate for a new paradigm of cosmological likelihood-based inference to tackle the challenges of next-generation surveys. This approach leverages recent developments in machine learning (ML) and its underlying technology to accelerate both parameter estimation and model selection — the fundamentals of Bayesian inference — in high-dimensional settings. We suggest combining (i) *emulation*; (ii) *differentiable and probabilistic programming*; (iii) *scalable Markov chain Monte Carlo (MCMC) sampling* techniques that exploit gradient information and (iv) *decoupled and scalable Bayesian model selection* techniques computing the Bayesian evidence purely from posterior samples. At the basis of these developments is the ability to exploit modern hardware accelerators, such as graphics processing units (GPUs) and tensor processing units (TPUs), to provide a high degree of parallelization for significant computational acceleration.

Emulation is based on statistical and ML models trained to replicate cosmological quantities of interest with high accuracy in a fraction of the time required by traditional methods. By replacing slow forward models underlying the likelihood function with, e.g. neural networks, which efficiently run on GPUs, accurate physical simulations can be computed with significantly fewer computational resources, providing massive speed-ups.

* E-mail: davide.piras@unige.ch

† Joint first authors

¹ <https://www.euclid-ec.org/>

² <https://www.desi.lbl.gov/>

³ <https://roman.gsfc.nasa.gov/>

⁴ <https://simonsobservatory.org/>

⁵ <https://www.lsst.org/>

Additionally, one can pair emulators with a likelihood fully written using differentiable and probabilistic programming languages. This unlocks the acceleration provided by ML environments such as `TensorFlow` (Abadi et al. 2015)⁶, `PyTorch` (Paszke et al. 2019)⁷ and `JAX` (Bradbury et al. 2018)⁸ to enable GPU execution and differentiability of the likelihood with respect to the input parameters. Moreover, probabilistic programming languages can be used, e.g. to easily formulate complex hierarchical models.

Differentiability also unlocks gradient-based sampling algorithms, such as Hamiltonian Monte Carlo (HMC, Duane et al. 1987; Neal 1996) and its variants, which represent a promising avenue to address the challenges of high-dimensional inference. These techniques allow one to efficiently explore the complex parameter spaces that characterize cosmological data analyses. Yet, they crucially rely on having access to the derivatives of the likelihood function with respect to input parameters, which are in general expensive and inaccurate, especially when computed by finite differences. This obstacle is being overcome through the development of fully-differentiable frameworks for cosmological data analysis, which also benefit from the acceleration provided by the same dedicated hardware of ML models.

The complexity of cosmological data analysis is not limited to parameter estimation in high-dimensional settings; it also encompasses the task of model comparison. This is particularly important for cosmology, whose main goal is the identification of the most accurate cosmological model of our Universe given observational data. Bayesian model comparison requires the computation of the model evidence (also known as the marginal likelihood) for different models. Consequently, in order to meaningfully and efficiently compare competing cosmological models in light of new data, it becomes imperative to develop novel methodologies that enable evidence estimation for model comparison in the high-dimensional landscapes that will characterize next-generation surveys. Such methodologies should not be necessarily bound to a particular sampling method in order to provide full flexibility. Being able to compute the evidence from posterior chains independently from the sampling algorithm is thus of paramount importance for enabling next-generation cosmological model comparison.

This paper focuses on both parameter estimation and model comparison in the context of Stage IV likelihood-based cosmological analyses, proposing a comprehensive approach towards cosmological inference from upcoming next-generation datasets. We particularly highlight how parameter estimation and model comparison can be performed from posterior samples obtained using the No U-Turn Sampler (NUTS, Hoffman & Gelman 2014), a highly efficient and adaptive variant of HMC. The expensive cosmological Boltzmann solvers are replaced by the `CosmoPower-JAX` emulators (Spurio Mancini et al. 2022; Piras & Spurio Mancini 2023), and evidence estimation is performed using `harmonic`, a software implementation of the learned harmonic mean estimator

(McEwen et al. 2021).

We demonstrate this paradigm on two scenarios. First, we sample the posterior distribution for a simulated Stage IV cosmic shear survey configuration using NUTS and a fully differentiable pipeline. We then compute the evidence from the posterior chains using `harmonic` and compare the estimate with that obtained using a nested sampler. We perform this operation for two competing cosmological models, with the goal of performing Bayesian model comparison, obtaining results in excellent agreement with nested sampling, but taking only a fraction of the time. We also showcase joint inference on three simulated next-generation surveys, each performing a 3x2pt analysis, obtaining values of the Bayes factor in 8 days, as opposed to the 12 years estimated for the same result to be obtained by traditional techniques.

This paper is structured as follows. We begin in Sect. 2 with some background on parameter estimation and model comparison, highlighting the challenges associated with the large parameter spaces characterizing upcoming surveys. We then outline in Sect. 3 each of the four pillars of the proposed new paradigm of cosmological likelihood-based inference for complete Bayesian analysis, including both parameter estimation and model selection, in high-dimensional settings. In Sect. 4 and Sect. 5 we introduce two simulated survey scenarios, to validate and demonstrate our framework for Bayesian inference. We conclude in Sect. 6.

2. BACKGROUND

Bayesian analyses are the backbone of modern cosmology, providing a principled framework to obtain cosmological parameter constraints and compare models. Here we review the basics of Bayesian inference, discussing both parameter estimation and model comparison.

2.1. Parameter estimation

The fundamental goal of Bayesian parameter estimation is to recover an accurate estimate of the posterior distribution $p(\theta|d, \mathcal{M})$, which encapsulates our understanding of the parameters θ given observed data d and a model \mathcal{M} . By employing Bayes' theorem, this distribution is related to the prior distribution $p(\theta|\mathcal{M}) \equiv \pi(\theta)$, the likelihood function $p(d|\theta, \mathcal{M}) \equiv \mathcal{L}(\theta)$, and the model evidence $p(d|\mathcal{M}) \equiv z_{\mathcal{M}}$ through the relation:

$$p(\theta|d, \mathcal{M}) = \frac{p(d|\theta, \mathcal{M})p(\theta|\mathcal{M})}{p(d|\mathcal{M})} = \frac{\mathcal{L}(\theta)\pi(\theta)}{z_{\mathcal{M}}}. \quad (1)$$

Estimating the posterior distribution through its direct evaluation on a grid of parameters is computationally infeasible due to the curse of dimensionality, which becomes a prohibitive factor even in spaces of moderate dimension. Additionally, the likelihood function often involves solving intricate physical models or running detailed simulations, and the posterior landscape may be multi-modal or have complex correlations between parameters. A grid-based evaluation might miss important regions of the parameter space or require an impractically fine grid to capture these features. Sampling algorithms such as MCMC methods, on the other hand, provide a stochastic exploration that can adapt to the shape of the distribution, and allocate computational resources more efficiently by focusing on the regions of parameter space with high posterior density.

⁶ <https://www.tensorflow.org/>

⁷ <https://pytorch.org/>

⁸ <https://github.com/google/jax>

Commonly used MCMC algorithms in cosmology include the Metropolis-Hastings (Lewis & Bridle 2002; Lewis 2013), affine-invariant ensemble (Goodman & Weare 2010; Foreman-Mackey et al. 2013), and nested sampling algorithms (Skilling 2006; Feroz & Hobson 2008; Feroz et al. 2009, 2019; Handley et al. 2015a,b), which have found widespread use thanks to their clear formulations and publicly available implementations. Another sampling technique is HMC (Duane et al. 1987; Neal 1996), which is a more efficient algorithm exploiting Hamiltonian dynamics. By leveraging the information coming from the gradients of the likelihood with respect to the input parameters, HMC concentrates sampling in regions of high posterior mass (the so-called typical set), resulting in a more efficient sampler, particularly in high-dimensional settings. A significantly improved version of HMC is NUTS (Hoffman & Gelman 2014), which overcomes the required tuning of hyperparameters for the numerical integration of Hamilton equations by preventing the sampler from taking U-turns. By exploring the parameter space in a more efficient way, NUTS results in a significantly higher acceptance rate, and has thus found widespread use.

2.2. Model comparison

For Bayesian model comparison it is necessary to compute the Bayesian evidence

$$z_{\mathcal{M}} = \int p(d|\theta, \mathcal{M})p(\theta|\mathcal{M})d\theta, \quad (2)$$

which is a challenging computational problem even in moderate dimensional settings. From the Bayesian evidence it is possible to compute the Bayes factor (BF_{12}), a crucial quantity used to assess which model (\mathcal{M}_1 or \mathcal{M}_2) is favored given current data d and defined as:

$$\text{BF}_{12} \equiv \frac{p(d|\mathcal{M}_1)}{p(d|\mathcal{M}_2)} = \frac{z_{\mathcal{M}_1}}{z_{\mathcal{M}_2}}. \quad (3)$$

In the common case where the models have equal prior probabilities, i.e. $p(\mathcal{M}_1) = p(\mathcal{M}_2)$, the Bayes factor reduces to the posterior model odds:

$$\text{BF}_{12} = \frac{p(\mathcal{M}_1|d)p(\mathcal{M}_2)}{p(\mathcal{M}_2|d)p(\mathcal{M}_1)} = \frac{p(\mathcal{M}_1|d)}{p(\mathcal{M}_2|d)}. \quad (4)$$

Estimating the evidence robustly and efficiently can therefore unlock the full potential of Bayesian analyses, in that it allows one to assess which model is preferred by observational data.

Nested sampling, introduced by Skilling (2006), provides a strategy to estimate the evidence for Bayesian model comparison. By a clever reparameterization of the likelihood in terms of the enclosed prior volume, nested sampling allows the evidence to be computed by a one-dimensional integral. While nested sampling targets the evidence, as a by-product posterior inferences can also be calculated by appropriate importance weighting (Skilling 2006). While highly successful, nested sampling places tight constraints on how sampling is performed (sampling the prior subject to likelihood isocontour constraints). Nested sampling couples the sampling strategy to the evidence calculation, limiting flexibility and scalability to high-dimensional parameter spaces (a notable exception for high-dimensional inference is proximal nested

sampling, although this is only applicable for log-convex likelihoods; Cai et al. 2022).

Recently, methods to compute the evidence that are agnostic to the sampling technique and only require posterior samples have gained popularity (McEwen et al. 2021; Srinivasan et al. 2024; Rinaldi et al. 2024). Among them, the learned harmonic mean estimator (McEwen et al. 2021) has been shown to provide robust estimates of the Bayesian evidence in a variety of scenarios, both in likelihood- and simulation-based inference (Polanska et al. 2023, 2024; Spurio Mancini et al. 2023). In particular, Polanska et al. (2024) integrated normalizing flows into the learned harmonic mean framework to learn the internal importance sampling target distribution, enhancing the robustness and scalability of the estimator. In contrast to the learned harmonic mean, alternative approaches (Srinivasan et al. 2024; Rinaldi et al. 2024) compute the evidence for a surrogate of the posterior, limiting their accuracy (see Polanska et al. 2024 for further discussion).

3. METHODOLOGY

We outline the four pillars of a new paradigm of cosmological likelihood-based inference that will be key to successfully tackling the challenges set by Stage IV cosmological surveys, including emulation (typically based on neural networks), differentiable and probabilistic programming, scalable gradient-based MCMC sampling, and scalable evidence estimation that is agnostic to sampling. All of these pillars provide the ability to exploit modern hardware accelerators, such as GPUs, to provide high degrees of parallelization for significant computational acceleration.

3.1. Emulation: CosmoPower-JAX

The typical bottleneck in a standard Bayesian analysis is the likelihood function, since it incorporates the entire physical knowledge on the observable being probed. Thousands or even millions of likelihood evaluations are needed to explore the posterior distribution in a complete manner, especially in high-dimensional settings. Even in low-dimensional settings, though, the forward model can be computationally expensive since it describes increasingly complex physical models, and extensions to the Λ CDM model might introduce additional correlations among an increasing number of parameters. Each call to the likelihood function thus has to be fast.

Emulation based on ML techniques, and neural networks in particular, have already shown promising results to accelerate the forward model, while providing excellent agreement with traditional techniques in a fraction of the computational time. Neural networks are trained once on a dataset of key quantities that represent a bottleneck in the forward model, and then can be easily integrated into the likelihood function, providing significant acceleration while retaining the accuracy of the predictions. One additional advantage of neural networks is that they effectively exploit the highly parallel computing provided by modern hardware accelerators, such as GPUs. This allows for faster training, batch prediction and scalable architectures with ever more expressive power.

ML-based emulation has found widespread use in cosmology (see e.g. Auld et al. 2007; Auld et al. 2008;

Manrique-Yus & Sellentin 2019; Angulo et al. 2021; Aricò et al. 2022; Nygaard et al. 2022; Zennaro et al. 2023; Bonici et al. 2022, 2024). In this work, we employ the `CosmoPower-JAX` package (Piras & Spurio Mancini 2023)⁹, a JAX version of the `CosmoPower` framework (Spurio Mancini et al. 2022)¹⁰. These dense neural networks were trained to map cosmological parameters to linear and nonlinear matter power spectra, serving as an efficient replacement for the Boltzmann solvers `CAMB` (Lewis & Challinor 2011) or `CLASS` (Blas et al. 2011) to integrate into the likelihood. By replacing Boltzmann solvers with accurate neural networks, we accelerate the likelihood evaluations and thus solve the first significant bottleneck of next-generation surveys.

3.2. Differentiable and probabilistic programming: JAX and NumPyro

Within the current ML paradigm, models are trained using optimizers that leverage gradients computed by the backpropagation of gradient information through automatic differentiation (Baydin et al. 2018). Consequently, each operation upon which a model may be built must necessarily be differentiable; that is, rules by which tangent and cotangent vectors propagate under the operators’ Jacobian must be defined (Wengert 1964; Rumelhart et al. 1986; Baydin et al. 2018).

Modern ML ecosystems, such as JAX and PyTorch, are constructed with this in mind, requiring that each primitive operation within their libraries is coupled with associated gradient rules. In this way, using the chain rule one can construct more complex operations, which are automatically differentiable (Bartholomew-Biggs et al. 2000; Margossian 2019). By automatic differentiation gradients can then be computed efficiently and accurately for complex operations, replacing numerical approximations of the derivatives, which are typically slow and unstable. Primarily due to its relative maturity, PyTorch provides a larger bank of primitive operations. However, JAX provides a more flexible ecosystem that is functional in nature (Wadler 1992) and thus well suited to the implementation of general physical models. In JAX complex operations may more easily be expressed, making it the natural ecosystem for the integration of differentiable programming in, e.g. cosmological models (Campagne et al. 2023), cosmological simulations (Li et al. 2024) or underlying mathematical methods, such as spherical harmonic transforms (Price & McEwen 2024).

Recently, probabilistic programming languages (PPLs) have been constructed leveraging these differentiable programming ecosystems. Through PPLs, statistical operations such as sampling and conditioning can be expressed using a familiar and intuitive Python syntax, whilst also inheriting the differentiability of the underlying ecosystems upon which they are constructed. For instance, PPLs allows one to easily formulate Bayesian hierarchical models, which are widely used in cosmology (see e.g. Mandel et al. 2011; Shariff et al. 2016; Alsing et al. 2016, 2017; Hinton et al. 2019; Porqueres et al. 2021; Mandel et al. 2022; Porqueres et al. 2023; Loureiro et al. 2023; Sellentin et al. 2023).

In this work we use NumPyro (Phan et al. 2019)¹¹, the NumPy backend for Pyro (Bingham et al. 2019). NumPyro is well-documented and easily interfaces with JAX, thus providing a mature framework to build a gradient-based inference pipeline. We implement the likelihood for the problems considered (see Sect. 4 and Sect. 5) in the NumPyro PPL, which requires a separate likelihood implementation. Critically, the likelihood must be differentiable: `CosmoPower-JAX` provides emulation of the matter power spectrum that is differentiable, and we couple it with `jax-cosmo` (Campagne et al. 2023)¹² for differentiable cosmological models to simulate observable spectra.

3.3. Scalable MCMC sampling: NUTS

HMC is a gradient-based sampling algorithm that is scalable to high-dimensional parameter spaces and does not suffer from the low acceptance rate of the standard Metropolis-Hastings algorithm. NUTS further improves on the acceptance rate of HMC by preventing the sampling trajectory from returning to previously-visited regions of the parameter space. Though it requires a higher number of model evaluations, NUTS typically leads to more reliable convergence, and is thus ideal for sampling in the context of Stage IV surveys.

Traditionally, the gradients required for HMC sampling are computed by finite differences, which is computationally costly and inaccurate. Since we have a differentiable pipeline, gradients can be computed by automatic differentiation, providing acceleration and improved accuracy. In the future, new samplers that further improve the convergence rate leading to even more robust inference can be straightforwardly incorporated into this paradigm.

3.4. Decoupled and scalable Bayesian model comparison: learned harmonic mean

The learned harmonic mean (McEwen et al. 2021) provides a robust and scalable estimator of the Bayesian evidence that is decoupled from the sampling strategy. It only requires samples from the posterior, and their corresponding unnormalized probability density. Hence, it can be used with posterior samples obtained through whatever method is best suited for the problem at hand. This property allows the use of the efficient and scalable NUTS sampler, while still being able to estimate the evidence for Bayesian model comparison.

The learned harmonic mean provides an estimator of the reciprocal evidence $\rho = z^{-1}$, defined by

$$\hat{\rho} = \frac{1}{N} \sum_{i=1}^N \frac{\varphi(\theta_i)}{\mathcal{L}(\theta_i)\pi(\theta_i)}, \quad \theta_i \sim p(\theta|d), \quad (5)$$

where $\varphi(\theta)$ is an arbitrary normalized probability density. The estimator can be viewed through the lens of importance sampling, with the posterior acting as the sampling density and $\varphi(\theta)$ as the target density (e.g. McEwen et al. 2021). If the importance sampling target density has fatter tails than the sampling density, the variance will grow large (even catastrophically) and the estimator will fail. The key to preventing this issue

⁹ <https://github.com/dpiras/cosmopower-jax>

¹⁰ <https://github.com/alessiospurioMancini/cosmopower>

¹¹ <https://num.pyro.ai/en/stable/>

¹² https://github.com/DifferentiableUniverseInitiative/jax_cosmo

is to find an appropriate target $\varphi(\theta)$ that is normalized and contained within the posterior. The optimal target density is the normalized posterior itself (McEwen et al. 2021), although this requires knowledge of its normalizing constant, which is precisely the quantity we are attempting to estimate. In the learned harmonic mean the target is learned to approximate the posterior from posterior samples, subject to the constraint that it has narrower tails. Several benchmark examples have been considered where the learned harmonic mean has been demonstrated to provide precise and accurate evidence estimates (McEwen et al. 2021).

More recently, normalizing flows have been integrated within the learned harmonic mean framework for the internal machine learning technique to learn the important target distribution (Polanska et al. 2023, 2024). Normalizing flows (Papamakarios et al. 2021) can be elegantly coupled with the learned harmonic mean to provide an approach that is more robust, flexible and scalable than the machine learning models considered previously.

Normalizing flows work by taking a simple base distribution, often a standard Gaussian, through a series of invertible transformations with learned parameters. Typically, in order to train the model the Kullback-Leibler divergence between the unknown target and the flow is minimized, resulting in training by maximum likelihood. Using normalizing flows it is possible to approximate potentially complex probability distributions, draw samples from them and evaluate their normalized density.

Once the flow is trained, its probability density can then be concentrated by reducing the “temperature” T of its base distribution. Specifically, this is achieved by scaling the base distribution’s variance by a factor $T \in (0, 1)$. This has the effect of concentrating the base distribution’s density, and in turn the trained flow’s density due to the continuity and differentiability of the flow (Polanska et al. 2024). The concentrated flow is then a normalized approximation of the posterior, maintaining a good topological agreement while having thinner tails — a perfect candidate for the target $\varphi(\theta)$. Estimates of the evidence error and other sanity checks can also be computed (McEwen et al. 2021; Polanska et al. 2024). The learned harmonic mean has already been thoroughly validated against nested sampling for numerous cosmological problems (McEwen et al. 2021; Polanska et al. 2023, 2024), and is implemented in the `harmonic` open-source Python package¹³. The code is written in JAX (Bradbury et al. 2018) to provide automatic differentiation and an efficient and scalable implementation that can be run on accelerators such as GPUs.

Critically, this approach decouples sampling from evidence calculation, allowing the use of scalable MCMC sampling techniques (e.g. NUTS). By integrating normalizing flows inside the overarching statistical framework of the learned harmonic mean we recover a robust estimator of the evidence that is also computationally scalable, proving the missing final component of a complete Bayesian analysis.

4. COSMIC SHEAR ANALYSIS WITH 37 AND 39 PARAMETERS

We first demonstrate the application of the proposed paradigm for a simulated next-generation survey performing a power spectrum cosmic shear analysis. Our goal is to estimate the Bayes factor between the Λ CDM and w_0w_a CDM models in high-dimensional scenarios, exposing the challenges of next-generation surveys if approached with traditional methods, while validating and showcasing the scalability of our approach.

4.1. Likelihood

The likelihood details are the same as in Piras & Spurio Mancini (2023), which we summarize here, following the notation of Spurio Mancini et al. (2022). We assume a tomographic survey with $N_{\text{bins}} = 10$ bins, where each galaxy is placed in a bin according to its estimated photometric redshift z . We compute the angular power spectra between pairs of redshift bins $i, j = 1, \dots, N_{\text{bins}}$ and for different probes. The cosmic shear angular power spectrum $C_{ij}^{\epsilon\epsilon}(\ell)$ is defined as:

$$C_{ij}^{\epsilon\epsilon}(\ell) = C_{ij}^{\gamma\gamma}(\ell) + C_{ij}^{\gamma I}(\ell) + C_{ij}^{I\gamma}(\ell) + C_{ij}^{II}(\ell), \quad (6)$$

which includes the contributions from pure shear (γ) and intrinsic alignment (I). Assuming the extended Limber approximation (LoVerde & Afshordi 2008), we can write:

$$C_{ij}^{AB}(\ell) = \int_0^{\chi_H} \frac{W_i^A(\chi)W_j^B(\chi)}{\chi^2} P_{\delta\delta} \left(k = \frac{\ell + 1/2}{\chi}, z \right) d\chi, \quad (7)$$

where $P_{\delta\delta}(k, z)$ is the matter power spectrum, $\chi_H = c/H_0$ (with c the speed of light and H_0 the Hubble constant today) is the Hubble radius, and $W(\chi)$ indicates a window function for each probe $\{A, B\} = \{\gamma, I\}$ as a function of the comoving distance χ .

For the pure cosmic shear γ the window function can be written as:

$$W_i^\gamma(\chi) = \frac{3\Omega_m H_0^2}{2c^2} \frac{\chi}{a} \int_\chi^{\chi_H} n_{i,\text{source}}(\chi') \frac{\chi' - \chi}{\chi'} d\chi', \quad (8)$$

where a is the scale factor, Ω_m is the matter density parameter today and $n_{i,\text{source}}(z)$ represents the tomographic redshift bin distribution of the sources being observed. For the intrinsic alignment field, we consider the non-linear alignment model (NLA, Hirata & Seljak 2004; Joachimi et al. 2011) modified as in Piras & Spurio Mancini (2023):

$$W_i^I(\chi) = -A_{IA,i} \frac{C_1 \rho_{\text{cr}} \Omega_m}{D(\chi)} n_{i,\text{source}}(\chi), \quad (9)$$

where $D(\chi)$ is the linear growth factor, ρ_{cr} is the critical density of the Universe, C_1 is a constant, and with one intrinsic alignment amplitude for each redshift bin $A_{IA,i}$, to allow for more flexibility in the modeling. For each redshift bin we also include a multiplicative bias m_i (Huterer et al. 2006; Amara & Réfrégier 2008; Kitching et al. 2015; Taylor & Kitching 2018; Mandelbaum et al. 2018), which rescales the cosmic shear power spectrum by a factor $(1 + m_i)(1 + m_j)$, and a shift parameter $D_{z_i,\text{source}}$ (Eifler et al. 2021), which shifts the mean of the bin redshift distribution so that we actually consider $n'_{i,\text{source}}(z) = n_{i,\text{source}}(z - D_{z_i,\text{source}})$.

¹³ <https://github.com/astro-informatics/harmonic/>

Table 1: Evidence and Bayes factors for the two analyses considered in this work. The evidence values obtained with nested sampling are the mean and standard deviation across three independent runs with different seeds, while the computation time is the typical one for a single run, and includes sampling and evidence estimation. The likelihood functions sampled with nested sampling and the No U-Turn Sampler (NUTS) are implemented using different packages and include different normalizations; despite these differences, we recover consistent values of the Bayes factor (BF) between Λ CDM and w_0w_a CDM.

(a) Cosmic shear with 37 (Λ CDM) and 39 (w_0w_a CDM) parameters, described in Sect. 4.

Method	$\log(z_{\Lambda\text{CDM}})$	$\log(z_{w_0w_a\text{CDM}})$	$\log \text{BF}$	Total computation time
CAMB + nested sampling	-107.03 ± 0.27	-107.81 ± 0.74	0.78 ± 0.79	~ 8 months (48 CPUs)
CosmoPower-JAX + NUTS + harmonic	40956.55 ± 0.06	40955.03 ± 0.04	1.53 ± 0.07	2 days (sampling, 12 GPUs) + 12 minutes (evidence, 1 GPU + 48 CPUs)
CosmoPower-JAX + NUTS + naïve flow estimator	400958 ± 5	40957 ± 4	1 ± 6	Similar to harmonic

(b) $3x(3x2pt)$ with 157 (Λ CDM) and 159 (w_0w_a CDM) parameters, described in Sect. 5.

Method	$\log(z_{\Lambda\text{CDM}})$	$\log(z_{w_0w_a\text{CDM}})$	$\log \text{BF}$	Total computation time
CAMB + nested sampling	Unfeasible	Unfeasible	Unfeasible	12 years (projected, 48 CPUs)
CosmoPower-JAX + NUTS + harmonic	$406689.6^{+0.5}_{-0.3}$	$406687.7^{+0.5}_{-0.3}$	$1.9^{+0.7}_{-0.5}$	8 days (sampling, 24 GPUs) + 17 minutes (evidence, 1 GPU + 48 CPUs)
CosmoPower-JAX + NUTS + naïve flow estimator	406703 ± 39	406701 ± 62	2 ± 73	Similar to harmonic

Finally, we model the redshift distributions with kernel density estimation (KDE), and consider a Gaussian likelihood with a simulated covariance matrix as in Tutusaus et al. (2020), with surface density of galaxies $n_{\text{source}} = 30$ galaxies/arcmin², observed ellipticity dispersion $\sigma_\epsilon = 0.3$, and sky fraction $f_{\text{sky}} = 0.35$. We compute each $C(\ell)$ spectrum for 30 log-spaced bin values between $\ell_{\text{min}} = 30$ and $\ell_{\text{max}} = 3000$. To compute the theoretical predictions for the cosmological observables we use the Core Cosmology Library (CCL, Chisari et al. 2019).

4.2. Models

We compare two cosmological models, namely the fiducial Λ CDM model (considered also in Piras & Spurio Mancini 2023) and the w_0w_a CDM model, where in the latter we assume that the dark energy equation of state evolves with cosmic time according to:

$$w(a) = w_0 + (1 - a)w_a, \quad (10)$$

which is the most common parametrization (also known as the CPL parametrization) for dynamic dark energy (Chevallier & Polarski 2001; Linder 2003). The w_0w_a CDM model introduces two extra parameters (w_0 and w_a), and Λ CDM is recovered for $w_0 = -1$ and $w_a = 0$. The simulated data vector is generated assuming a Λ CDM model, so we expect that the Bayes factor should favor this hypothesis. The prior distributions on the cosmological and nuisance parameters are the same as those in Piras & Spurio Mancini (2023), with the addition of uniform distributions for w_0 and w_a between $(-1.5, -0.5)$ and $(-0.5, 0.5)$, respectively. The total number of parameters being sampled is thus 37 and 39 for Λ CDM and w_0w_a CDM, respectively.

4.3. Computational approaches

We compare three approaches to perform parameter estimation and model comparison: the “traditional” ap-

proach, the proposed “future” paradigm, and an approach based on a naïve flow estimation of the evidence for comparison (discussed in Spurio Mancini et al. 2023; Polanska et al. 2024).

First, we use the nested sampler PolyChord (Handley et al. 2015a,b) to sample the posterior distribution and run the inference pipeline within Cobaya (Torrado & Lewis 2021), using CAMB to predict the matter power spectrum. We run PolyChord with default settings on 48 CPU cores twice, each time assuming either Λ CDM or w_0w_a CDM. Since we found the values of the evidence to fluctuate significantly between different runs, we repeated the process three times for each model, quoting the mean and standard deviation of the log evidence, and the average time to obtain the Bayes factor.

Second, we compare this “traditional” approach with our “future” paradigm, which replaces CAMB with the CosmoPower-JAX emulator, rewrites the likelihood in the auto-differentiable language JAX, making use of `jax-cosmo`, samples the posterior distribution with NUTS implemented in NumPyro, accelerated by automatic differentiation, and computes a robust estimate of the evidence from posterior samples using the learned harmonic mean implemented in `harmonic`. We run the NUTS chains in parallel on 12 A100 80GB GPUs, collecting 60 chains of 2000 samples each (after 400 samples of warm-up) for both Λ CDM and w_0w_a CDM. For `harmonic`, we train a rational-quadratic spline flow (Durkan et al. 2019), consisting of 4 layers and 64 spline bins, on a single GPU. We separate the NUTS samples into a training set the flow is trained on and an inference set that is used for estimation. We use 30% of chains for training and the remaining 70% for inference to compute the evidence on 48 CPUs with temperature $T = 0.8$.

Third, the evidence can also be estimated directly from the normalizing flow in an approach introduced by Spurio Mancini et al. (2023), and further discussed in Polanska et al. (2024), called the “naïve flow estimator”.




Figure 1. Corner plot for the 37-dimensional Λ CDM model, showing the posterior contours obtained with CosmoPower-JAX in red and the concentrated flow with temperature $T = 0.8$ in blue.

Since the flow is a normalized approximation of the posterior, the expectation of the ratio of the unnormalized posterior density and the flow density should be the evidence itself. This can be estimated via a Monte Carlo expectation as the mean of the ratios evaluated across samples from the posterior. In contrast to the learned harmonic mean, this approach is highly sensitive to the flow being a close approximation of the posterior and consequently can suffer from a large bias and variance. We also consider this method of estimating the evidence to demonstrate the need for an alternative principled estimator of the evidence, as provided by the learned harmonic mean. We compare the results obtained repla-

cing harmonic with the naïve flow estimator on the same NUTS samples and hardware.

4.4. Results

The results are reported in Table 1(a). The values of the evidence when using CAMB + nested sampling and CosmoPower-JAX + NUTS + harmonic use different implementations of the likelihood (the latter implemented in a probabilistic programming framework), including different normalizations, so their values cannot be compared directly. Nevertheless, the log BF's are in good agreement between the different approaches, with 0.78 ± 0.79 for nested sampling and 1.53 ± 0.07 for




Figure 2. Same as Fig. 1 for the 39-dimensional w_0w_a CDM model.

harmonic. Note that the evidence values correctly favour the ground truth model considered for the simulated data. While the evidence computed by the naïve flow estimator is also in agreement with a value of 1 ± 6 , notice that its error is considerably larger than the other estimators, as anticipated. While the evidence values computed by the “traditional” and “future” paradigms are in close agreement, computational times are dramatically different, with the “future” paradigm being two orders of magnitude faster, taking 2 days (with only 12 additional minutes required to compute the evidence) instead of roughly 8 months.

The contours for the posterior samples (red) alongside the concentrated flow at $T = 0.8$ (blue) used for inference

for this analysis are shown in Fig. 1 and Fig. 2 for the Λ CDM and w_0w_a CDM models respectively. Note that the flow is contained within the posterior.

5. 3X2PT ANALYSIS WITH 157 AND 159 PARAMETERS

We further showcase the scalability and robustness of our proposed paradigm by considering three different next-generation simulated surveys, each performing a 3x2pt analysis, where information on cosmic shear, galaxy clustering and their cross-correlation is combined to obtain more stringent constraints on the cosmological parameters (Joachimi & Bridle 2010). We refer to this application as a 3x(3x2pt) analysis. We present the likelihood and results below; the cosmological models and

computational approaches are the same as Sect. 4, although as we discuss in Sect. 5.2 we are not able to run the “traditional” approach in this high-dimensional setting.

5.1. Likelihood

The galaxy clustering field (n) power spectrum $C_{ij}^{nn}(\ell)$ can be expressed using Eq. (7) assuming that galaxies are a linearly-biased tracer of dark matter, with a free parameter b_i for each redshift bin. The corresponding window function can then be written as:

$$W_i^n(\chi) = b_i n_{i,\text{lens}}(\chi), \quad (11)$$

where $n_{i,\text{lens}}$ is a different sample of redshift distributions, also including a shift for each bin $D_{z_i,\text{lens}}$. Finally, the cross power spectrum between the shear field and the galaxy clustering field, also called the galaxy-galaxy lensing power spectrum, is written as:

$$C_{ij}^{m\epsilon}(\ell) = C_{ij}^{n\gamma}(\ell) + C_{ij}^{n\ell}(\ell). \quad (12)$$

The total number of parameters for this cross-survey analysis is 157 for ΛCDM and 159 for $w_0w_a\text{CDM}$, since each survey comes with its own 50 nuisance parameters, and all surveys share 7 (or 9) cosmological parameters. The specific details of every survey are reported in Piras & Spurio Mancini (2023).

5.2. Results

All numerical results for this higher dimensional analysis are reported in Table 1(b). The high number of parameters and the need to consider multiple probes pose a significant challenge for “traditional” methods, despite being quite realistic for next-generation surveys. We estimated that the computation of the Bayes factor in such a scenario would require about 12 years on 48 CPU cores, making it effectively impossible to perform model selection in a reasonable time.

On the other hand, the proposed “future” paradigm provides an estimate of the Bayes factor based on 40 (45) converged chains for the ΛCDM ($w_0w_a\text{CDM}$) model, totalling 8 days on 24 GPUs, with only 17 additional minutes required to compute the evidence. In this case, the flow model we consider is constructed from 2 layers and 128 spline bins. The log Bayes factor computed by `harmonic` is $1.9^{+0.7}_{-0.5}$, which correctly favours the ground truth model considered for the simulated data. The evidence computed by the naïve flow estimator is 2 ± 73 : its error is very large, as anticipated, rendering this approach unusable in practice.

The contours plots for the posterior samples (red) alongside the concentrated flow at $T = 0.8$ (blue) used for inference for this analysis are shown in Fig. 3 and Fig. 4 for the ΛCDM and $w_0w_a\text{CDM}$ models respectively. Fig. 5 and Fig. 6 show the one-dimensional marginal plots of the posterior samples (red) alongside the concentrated flow at $T = 0.8$ (blue), for all the model parameters for ΛCDM and $w_0w_a\text{CDM}$ respectively. It can be seen that the flow is indeed contained within the posterior as required by the learned harmonic mean.

6. CONCLUSIONS

We propose a combination of state-of-the-art techniques from machine learning (ML), differentiable and

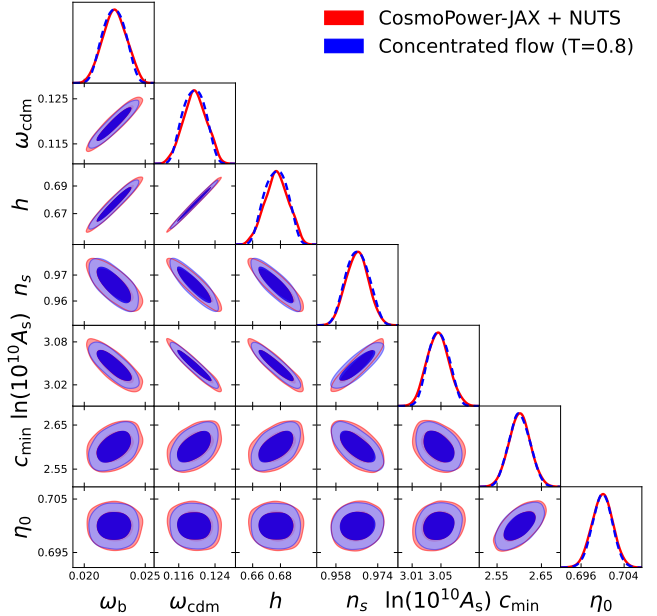


Figure 3. Corner plot of the subset of cosmological parameters for the 157-dimensional ΛCDM model, showing the posterior contours in red and the concentrated flow with temperature $T = 0.8$ in blue.




Figure 4. Same as Fig. 3 for the 159-dimensional $w_0w_a\text{CDM}$ model.

probabilistic programming, high-dimensional sampling and robust statistical estimation to tackle the Bayesian likelihood-based cosmological analyses of the future. Instead of relying on nested sampling with standard Boltzmann solvers, we leverage ML-based emulation (`CosmoPower-JAX`), differentiable programming (`JAX`), probabilistic programming (`NumPyro`), more efficient samplers (`NUTS`), and robust evidence estimation (`harmonic`) to perform end-to-end cosmological analyses including parameter estimation and model comparison in





Figure 5. Marginal distributions of all parameters for the 157-dimensional Λ CDM model, with the posterior distribution obtained with `CosmoPower-JAX` in red, and the concentrated flow with temperature $T = 0.8$ in blue.

Figure 6. Same as Fig. 5 for the 159-dimensional $w_0 w_a$ CDM model.

two challenging cosmological scenarios scaling up to 159 parameters.

We first computed the Bayes factor for a simulated Stage IV cosmic shear analysis comparing the Λ CDM model with 37 parameters and a dynamical dark energy model w_0w_a CDM with 39 parameters. We demonstrated excellent agreement with nested sampling despite different implementations of the likelihood, while requiring two orders of magnitude less computational time. We additionally demonstrated the scalability of our approach by considering the joint analysis of three Stage IV surveys, each performing a 3x2pt analysis, totalling either 157 (Λ CDM) or 159 (w_0w_a CDM) parameters. We computed the Bayes factor obtaining results requiring only 8 days on 24 GPUs, while projecting that a “traditional” analysis on 48 CPU cores would require 12 years, assuming nested sampling is capable of scaling to such a high-dimensional setting.

We advocate for a combination of ML emulation, differentiable and probabilistic programming, scalable sampling and robust evidence estimation to tackle forthcoming cosmological likelihood-based analyses. This “future” paradigm unlocks parameter estimation and model comparison for Stage IV cosmological surveys with an unprecedented number of parameters. Given that all packages used in our analysis are already publicly available, we envision this approach could become the standard in future likelihood-based analyses, allowing the community to extract the maximal amount of information from upcoming observations in a practical timescale.

AUTHOR CONTRIBUTIONS

DP: conceptualization; formal analysis; investigation; methodology; validation; software; visualization; resources; writing – original draft. **AP:** formal analysis; methodology; validation; software; visualization; writing – original draft. **ASM:** conceptualization; methodology; validation; software; supervision; writing – review & editing. **MAP:** software; validation; writing – review & editing. **JDM:** conceptualization; methodology; software; validation; project administration; funding acquisition; supervision; writing – review & editing.

ACKNOWLEDGEMENTS

DP was supported by a Swiss National Science Foundation (SNSF) Professorship grant (No. 202671), and by the SNF Sinergia grant CRSII5-193826 “AstroSignals: A New Window on the Universe, with the New Generation of Large Radio-Astronomy Facilities”. AP is supported by the UCL Centre for Doctoral Training in Data Intensive Science (STFC grant number ST/W00674X/1). ASM acknowledges support from the MSSL STFC Consolidated Grant ST/W001136/1. MAP and JDM are supported in part by EPSRC (grant number EP/W007673/1). The authors are pleased to acknowledge that part of the work reported on in this paper was performed using the Princeton Research Computing resources at Princeton University which is a consortium of groups led by the Princeton Institute for Computational Science and Engineering (PICSciE) and Office of Information Technology’s Research Computing.

REFERENCES

- Abadi M., et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <https://www.tensorflow.org/>
- Ade P., et al., 2019, *Journal of Cosmology and Astroparticle Physics*, 2019, 056–056
- Alsing J., Heavens A., Jaffe A. H., Kiessling A., Wandelt B., Hoffmann T., 2016, *MNRAS*, 455, 4452
- Alsing J., Heavens A., Jaffe A. H., 2017, *MNRAS*, 466, 3272
- Amara A., Réfrégier A., 2008, *Monthly Notices of the Royal Astronomical Society*, 391, 228
- Angulo R. E., Zennaro M., Contreras S., Aricò G., Pellejero-Ibañez M., Stücker J., 2021, *MNRAS*, 507, 5869
- Aricò G., Angulo R., Zennaro M., 2022, *Open Research Europe*, 1
- Auld T., Bridges M., Hobson M. P., Gull S. F., 2007, *MNRAS*, 376, L11
- Auld T., Bridges M., Hobson M. P., 2008, *Monthly Notices of the Royal Astronomical Society*, 387, 1575
- Bartholomew-Biggs M., Brown S., Christianson B., Dixon L., 2000, *Journal of Computational and Applied Mathematics*, 124, 171
- Baydin A., Pearlmutter B., Radul A., Siskind J., 2018, *Journal of Machine Learning Research*, 18, 1
- Bingham E., et al., 2019, *J. Mach. Learn. Res.*, 20, 28:1
- Blas D., Lesgourgues J., Tram T., 2011, *Journal of Cosmology and Astroparticle Physics*, 2011, 034–034
- Bonici M., Biggio L., Carbone C., Guzzo L., 2022, *arXiv e-prints*, p. [arXiv:2206.14208](https://arxiv.org/abs/2206.14208)
- Bonici M., Baxter E., Bianchini F., Ruiz-Zapatero J., 2024, *The Open Journal of Astrophysics*, 7, 10
- Bradbury J., et al., 2018, JAX: composable transformations of Python+NumPy programs, <http://github.com/google/jax>
- Cai X., McEwen J. D., Pereyra M., 2022, *Statistics and Computing*, 32
- Campagne J.-E., et al., 2023, *Open J. Astrophys.*, 6, 1
- Chevallier M., Polarski D., 2001, *International Journal of Modern Physics D*, 10, 213
- Chisari N. E., et al., 2019, *The Astrophysical Journal Supplement Series*, 242, 2
- Duane S., Kennedy A., Pendleton B. J., Roweth D., 1987, *Physics Letters B*, 195, 216
- Durkan C., Bekasov A., Murray I., Papamakarios G., 2019, *Advances in neural information processing systems*, 32
- Eifler T., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 1514
- Feroz F., Hobson M. P., 2008, *Monthly Notices of the Royal Astronomical Society*, 384, 449
- Feroz F., Hobson M. P., Bridges M., 2009, *Monthly Notices of the Royal Astronomical Society*, 398, 1601
- Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2019, *The Open Journal of Astrophysics*, 2
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Goodman J., Weare J., 2010, *Communications in Applied Mathematics and Computational Science*, 5, 65
- Handley W. J., Hobson M. P., Lasenby A. N., 2015a, *Monthly Notices of the Royal Astronomical Society: Letters*, 450, L61
- Handley W. J., Hobson M. P., Lasenby A. N., 2015b, *Monthly Notices of the Royal Astronomical Society*, 453, 4384
- Hinton S. R., et al., 2019, *ApJ*, 876, 15
- Hirata C. M., Seljak U., 2004, *Phys. Rev. D*, 70, 063526
- Hoffman M. D., Gelman A., 2014, *J. Mach. Learn. Res.*, 15, 1593–1623
- Huterer D., Takada M., Bernstein G., Jain B., 2006, *Monthly Notices of the Royal Astronomical Society*, 366, 101
- Ivezić Ž., et al., 2019, *ApJ*, 873, 111
- Joachimi B., Bridle S. L., 2010, *A&A*, 523, A1
- Joachimi B., Mandelbaum R., Abdalla F. B., Bridle S. L., 2011, *A&A*, 527, A26
- Kitching T. D., Taylor A. N., Cropper M., Hoekstra H., Hood R. K. E., Massey R., Niemi S., 2015, *Monthly Notices of the Royal Astronomical Society*, 455, 3319
- Laureijs R., et al., 2011, *arXiv e-prints*, p. [arXiv:1110.3193](https://arxiv.org/abs/1110.3193)
- Levi M., et al., 2019, in *Bulletin of the American Astronomical Society*. p. 57 ([arXiv:1907.10688](https://arxiv.org/abs/1907.10688)), [doi:10.48550/arXiv.1907.10688](https://doi.org/10.48550/arXiv.1907.10688)
- Lewis A., 2013, *Phys. Rev. D*, 87, 103529
- Lewis A., Bridle S., 2002, *Phys. Rev. D*, 66, 103511

- Lewis A., Challinor A., 2011, CAMB: Code for Anisotropies in the Microwave Background, Astrophysics Source Code Library, record ascl:1102.026 (ascl:1102.026)
- Li Y., Modi C., Jamieson D., Zhang Y., Lu L., Feng Y., Lanasse F., Greengard L., 2024, *Astrophysical Journal, Supplement Series*, 270
- Linder E. V., 2003, *Phys. Rev. Lett.*, **90**, 091301
- Mandel K. S., Afshordi N., 2008, *Physical Review D*, 78
- Loureiro A., Whiteway L., Sellentin E., Silva Lafaurie J., Jaffe A. H., Heavens A. F., 2023, *The Open Journal of Astrophysics*, **6**, 6
- Mandel K. S., Narayan G., Kirshner R. P., 2011, *ApJ*, **731**, 120
- Mandel K. S., Thorp S., Narayan G., Friedman A. S., Avelino A., 2022, *MNRAS*, **510**, 3939
- Mandelbaum R., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 3170
- Manrique-Yus A., Sellentin E., 2019, *Monthly Notices of the Royal Astronomical Society*
- Margossian C. C., 2019, *WIREs Data Mining and Knowledge Discovery*, **9**, e1305
- McEwen J. D., Wallis C. G. R., Price M. A., Mancini A. S., 2021, Machine learning assisted Bayesian model comparison: learnt harmonic mean estimator ([arXiv:2111.12720](https://arxiv.org/abs/2111.12720))
- Neal R. M., 1996, Bayesian Learning for Neural Networks, Vol. 118 of Lecture Notes in Statistics. Springer-Verlag
- Nygaard A., Brinch Holm E., Hannestad S., Tram T., 2022, *arXiv e-prints*, p. [arXiv:2205.15726](https://arxiv.org/abs/2205.15726)
- Papamakarios G., Nalisnick E., Rezende D. J., Mohamed S., Lakshminarayanan B., 2021, *The Journal of Machine Learning Research*, **22**, 2617
- Paszke A., et al., 2019, in , Advances in Neural Information Processing Systems 32. Curran Associates, Inc., pp 8024–8035, <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- Phan D., Pradhan N., Jankowiak M., 2019, *arXiv preprint arXiv:1912.11554*
- Piras D., Spurio Mancini A., 2023, *The Open Journal of Astrophysics*, 6
- Polanska A., Price M. A., Spurio Mancini A., McEwen J. D., 2023, *Physical Sciences Forum*, 9
- Polanska A., Price M. A., Piras D., Mancini A. S., McEwen J. D., 2024, Learned harmonic mean estimation of the Bayesian evidence with normalizing flows ([arXiv:2405.05969](https://arxiv.org/abs/2405.05969))
- Porqueres N., Heavens A., Mortlock D., Lavaux G., 2021, *MNRAS*, **502**, 3035
- Porqueres N., Heavens A., Mortlock D., Lavaux G., Mäkinen T. L., 2023, *arXiv e-prints*, p. [arXiv:2304.04785](https://arxiv.org/abs/2304.04785)
- Price M. A., McEwen J. D., 2024, *Journal of Computational Physics*, p. 113109
- Rinaldi S., Demasi G., Del Pozzo W., Hannuksela O. A., 2024, *arXiv e-prints*, p. [arXiv:2405.07504](https://arxiv.org/abs/2405.07504)
- Rumelhart D. E., Hinton G. E., Williams R. J., 1986, *Nature*, **323**, 533
- Sellentin E., Loureiro A., Whiteway L., Lafaurie J. S., Balan S. T., Olamaie M., Jaffe A. H., Heavens A. F., 2023, *The Open Journal of Astrophysics*, **6**, 31
- Shariff H., Jiao X., Trotta R., van Dyk D. A., 2016, *ApJ*, **827**, 1
- Skilling J., 2006, *Bayesian Analysis*, **1**, 833
- Spergel D., et al., 2015, *arXiv e-prints*, p. [arXiv:1503.03757](https://arxiv.org/abs/1503.03757)
- Spurio Mancini A., Piras D., Alsing J., Joachimi B., Hobson M. P., 2022, *Monthly Notices of the Royal Astronomical Society*, **511**, 1771–1788
- Spurio Mancini A., Docherty M. M., Price M. A., McEwen J. D., 2023, *RAS Techniques and Instruments*, **2**, 710
- Srinivasan R., Crisostomi M., Trotta R., Barausse E., Breschi M., 2024, *arXiv e-prints*, p. [arXiv:2404.12294](https://arxiv.org/abs/2404.12294)
- Taylor A. N., Kitching T. D., 2018, *Monthly Notices of the Royal Astronomical Society*, **477**, 3397
- Torrado J., Lewis A., 2021, *Journal of Cosmology and Astroparticle Physics*, 2021, 057
- Tutusaus I., et al., 2020, *A&A*, **643**, A70
- Wadler P., 1992, in Proceedings of the 19th ACM SIGPLAN-SIGACT symposium on Principles of programming languages. pp 1–14
- Wengert R. E., 1964, *Commun. ACM*, **7**, 463–464
- Zennaro M., Angulo R. E., Pellejero-Ibáñez M., Stücker J., Contreras S., Aricò G., 2023, *MNRAS*, **524**, 2407

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which

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