

Eryn: A multi-purpose sampler for Bayesian inference*

Nikolaos Karnesis,¹ Michael L. Katz,² Natalia Korsakova,³ Jonathan R. Gair,² Nikolaos Stergioulas¹

¹ Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece

²Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut,
Am Mühlenberg 1, 14476 Potsdam-Golm, Germany

³Astroparticule et Cosmologie, CNRS, Université Paris Cité, F-75013 Paris, France
(Dated: March 7, 2023)

In recent years, methods for Bayesian inference have been widely used in many different problems in physics where detection and characterization are necessary. Data analysis in gravitational-wave astronomy is a prime example of such a case. Bayesian inference has been very successful because this technique provides a representation of the parameters as a posterior probability distribution, with uncertainties informed by the precision of the experimental measurements. During the last couple of decades, many specific advances have been proposed and employed in order to solve a large variety of different problems. In this work, we present a Markov Chain Monte Carlo (MCMC) algorithm that integrates many of those concepts into a single MCMC package. For this purpose, we have built *Eryn*, a user-friendly and multipurpose toolbox for Bayesian inference, which can be utilized for solving parameter estimation and model selection problems, ranging from simple inference questions, to those with large-scale model variation requiring trans-dimensional MCMC methods, like the LISA global fit problem. In this paper, we describe this sampler package and illustrate its capabilities on a variety of use cases.

I. INTRODUCTION

In physics, and in science in general, one of the most encountered problems is the one of model calibration and comparison. We test our models of the physical world against the measured data, to estimate their parameters and to robustly determine the most suitable model that describes our observations. A crucial first step in this direction, is to efficiently explore the posterior distribution of the parameters given the measured data. Markov Chain Monte Carlo (MCMC) algorithms have proven to be very successful in this regard [1–4], being one of the few methods which can efficiently perform Bayesian inference when the posterior is not analytically tractable and without solving exactly for the marginal likelihood. This is compared to, for example, grid methods, which are often computationally unfeasible. This is especially true in the field of Gravitational-Wave (GW) astronomy, where MCMC methods have been extensively used in order to find physical parameters for signals buried in the data (see e.g. [5–10]), as well as to hierarchically infer the properties of the underlying astrophysical populations (e.g. [11–13]). MCMC approaches can also compute the *marginal likelihood* or *evidence* (see section II), by using techniques such as *Thermodynamic Integration* (see section IID). In a Bayesian framework, the evidence difference between two models can be used to compute the Bayes Factor, which is used to select between different models that could describe the observations.

Thermodynamic integration (or other approximations, see section IIE or [14, 15]), is ideal for cases where the number of competing models is small. However, in situa-

tions where the number of potential models becomes too large, the task of iteratively and hierarchically computing the marginal likelihood can become computationally inefficient, or even practically unachievable. Such is the case for future signal-dominated GW observatories, such as the Laser Interferometer Space Antenna (LISA) [16] or other proposed space-borne GW observatories [17–19]. LISA will observe different types of GW sources, the most numerous of them being the Ultra Compact Binaries (UCBs) within the Milky Way [16, 20–27]. Those are mostly comprised of a population of Double White Dwarfs (DWD), although a small fraction of neutron star - white dwarf (NS-WD) or double neutron stars (NS-NS) binaries are expected [28, 29]. In fact, LISA is going to detect GWs from the complete population of $\mathcal{O}(10^7)$ sources simultaneously, but only a small fraction of them are going to be individually resolvable ($\mathcal{O}(10^4)$). The large majority of signals will generate an anisotropic and non-stationary “confusion” type of signal, which will dominate the LISA band between 0.05 and ~ 0.2 mHz.

In the above context, computing the marginal likelihood for such a large parameter space and for all possible numbers of events that could be in the population becomes computationally prohibitive. Instead, we can employ *dynamical trans-dimensional* MCMC methods [30]. This family of methods can be quite challenging to tune, but it has proven to yield satisfactory results, even for such demanding problems as the LISA data [21, 23]. There are also implementation challenges, which arise from technical aspects of the algorithm; one example being the *dimension matching* requirement when proposing moves between models with different dimensionality. In terms of algorithm efficiency, it is also crucial to choose proposal distributions that allow smooth transitions on a dynamical parameter space, a task which in many cases requires substantial effort. For these technical reasons,

* <https://github.com/mikekatz04/Eryn>

all the available software tools of this kind have been specifically developed for the particular problem they intend to solve.

In this work, we present **Eryn**, a reversible jump MCMC algorithm, capable of *efficiently sampling dynamical parameter spaces, while remaining generic and usable by a large community*. We build upon various ideas from statistics, astronomy, etc., in order to develop an efficient statistical toolbox that can be applied to the majority of problems involving detection and characterization of signals. Our primary goal, however, is to utilize **Eryn** as a basic ingredient for a data analysis pipeline [31] to perform the LISA *Global Fit* [21, 23]. The Global Fit is a data analysis strategy required to tackle the problem of multiple source detection, separation, and characterization in LISA data. For demonstration purposes in this work, we use **Eryn** to analyze a “reduced” scenario of the LISA data in section IV.

This paper is organized as follows: in section II, we begin with explaining the foundations of the MCMC algorithm, as well as some of the relevant methods that we have adopted for our implementation. In section III, we describe how the methods introduced in section II are combined into the actual toolbox implemented in **Eryn**. In section III C, we demonstrate the capabilities of **Eryn** through some toy examples, while in section IV we apply our machinery to more demanding applications in Gravitational-Wave astronomy. In particular, we demonstrate **Eryn** by performing model selection on a simulated population of Ultra Compact Galactic Binaries (UCBs) as measured by the future LISA observatory. Finally, in section V, we summarize our work and discuss future applications. We should state again here, that **Eryn** is available as open source software in <https://github.com/mikekatz04/Eryn>.

II. MARKOV CHAIN MONTE CARLO ALGORITHMS

Nowadays, MCMC methods are considered to be a cornerstone of Bayesian Inference, being very effective in finding solutions to problems encountered across wide-ranging disciplines [e.g. 5, 6, 32–35]. These include the sampling of the posterior densities of parameters of interest, the numerical marginalisation over nuisance parameters, and providing a framework to compute the marginal posterior distributions (or evidences) that can be used for model selection. The Bayesian framework is based around Bayes’ Theorem:

$$p(\vec{\theta}|y, \mathcal{M}) = \frac{p(y|\vec{\theta}, \mathcal{M})p(\vec{\theta}|\mathcal{M})}{p(y|\mathcal{M})}, \quad (1)$$

where y is the measured data and \mathcal{M} our chosen model of analysis. The $p(\vec{\theta}|y)$ term is the posterior distribution of the parameter set $\vec{\theta}$, which is related to the likelihood function of the data $p(y|\vec{\theta}, \mathcal{M})$ and the prior densities

of the parameters $p(\vec{\theta}|\mathcal{M})$. The evidence $p(y|\mathcal{M})$ is the marginal posterior over the parameter space $\vec{\theta} \in \Theta$:

$$\begin{aligned} \mathcal{Z} \equiv p(y|\mathcal{M}) &= \int_{\Theta} p(\vec{\theta}, y|\mathcal{M})d\vec{\theta} \\ &= \int_{\Theta} p(y|\vec{\theta}, \mathcal{M})p(\vec{\theta}|\mathcal{M})d\vec{\theta}. \end{aligned} \quad (2)$$

For parameter estimation purposes, the evidence acts as a normalization constant and can be ignored. However, it is really important if one wants to perform model selection over the measured data. We shall describe in detail how one can numerically approximate the integral of eq. (2) in section II E.

MCMC algorithms work by constructing a Markov Chain sequence, whose elements, $\vec{\theta}(t_i)$, for $i = 0, 1, \dots$, are independent samples from the target distribution, $f(\vec{\theta})$. Under fairly general assumptions, the distribution of samples in the chain will converge to the target distribution provided the algorithm satisfies *detailed balance*:

$$f(\vec{\theta})p(\vec{\theta} \rightarrow \vec{\theta}') = f(\vec{\theta}')p(\vec{\theta}' \rightarrow \vec{\theta}). \quad (3)$$

Here $p(\vec{\theta} \rightarrow \vec{\theta}')$ is the probability that the Markov chain moves from point $\vec{\theta}$ to point $\vec{\theta}'$. The most widely-used MCMC algorithm is Metropolis-Hastings [1, 2], which is explained in algorithm box 1. The first step of the algorithm is to define an initial state, $\vec{\theta}(t_0)$. Then, at each subsequent step i , a new state is proposed, by randomly drawing from a given proposal distribution $q(\vec{\theta}'|\vec{\theta}(t_i))$. The newly proposed state is then accepted with a certain probability, given by eq. (4). If the move is accepted we set $\vec{\theta}(t_{i+1}) = \vec{\theta}'$, otherwise we set $\vec{\theta}(t_{i+1}) = \vec{\theta}(t_i)$. Any reasonable choice of the proposal density will generate a Markov chain with the correct stationary distribution. However, a good choice of q is critical for its efficiency, i.e. achieving the convergence of the MCMC chains within a reasonable computational time. For the special case of a symmetric proposal distribution, such as the widely used multivariate Gaussian distribution, the ratio of eq. (4) in algorithm box 1 becomes simply the ratio of the target densities at the current $\vec{\theta}(t_n)$ and proposed $\vec{\theta}'$ points. For high-dimensional problems, the multivariate Gaussian proposal can be tuned during the burn-in period of sampling to improve efficiency [15, 36–38], or even scaled according to the Cramer-Rao bound, estimated from the Information matrix [39].

Although the MH algorithm has been quite successful in tackling inference problems, there are practical implementation issues to overcome. Improving acceptance rate is crucial for convergence, and sometimes improvements in the proposal distribution are not sufficient to efficiently sample the parameter space. To tackle these issues, various MCMC enhancements have been proposed. A prime example is the Hamiltonian Monte Carlo (HMC) algorithms that utilize local gradients in order to generate proposal points [40, 41]. One variant of HMC is the No-U-Turn sampler which automates part of the required

Algorithm 1: The Metropolis-Hastings algorithm [1, 2]. f is the target density to be sampled.

- 1: Set the initial state of chain: $\vec{\theta}(t_0) \equiv \vec{\theta}$.
- 2: At each subsequent step n , draw a new state $\vec{\theta}'$ from a *proposal distribution* $q(\vec{\theta}'|\vec{\theta}(t_{n-1}))$.
- 3: Take

$$\vec{\theta}(t_n) = \begin{cases} \vec{\theta}' & \text{with probability } \alpha(\vec{\theta}(t_{n-1}), \vec{\theta}') \\ \vec{\theta} & \text{with probability } 1 - \alpha(\vec{\theta}(t_{n-1}), \vec{\theta}') \end{cases},$$

where

$$\alpha(x, y) = 1 \wedge \left\{ \frac{f(y)q(x|y)}{f(x)q(y|x)} \right\}. \quad (4)$$

- 4: Go to 2 and repeat until equilibrium is reached, and enough independent samples have been drawn from the target distribution.
-

tuning of the HMC [42] sampler. Another alternative to the “standard” MH is the Gibbs sampling algorithm, which is particularly useful if the conditional distributions of the parameters of the model are known [43–45]. All of the above developments have been shown to be useful in various disciplines [33, 46–55]. Finally, there have recently been numerous proposals that aim to enhance sampling with machine learning techniques. At their core, many of these methods optimize the exploration of the likelihood surface, either by learning it directly (see for example [56]) or by sampling it in a simpler latent space (for example [57]).

In this work, we introduce **Eryn**, which is built around the `emcee` package [58], enhanced with a variety of sampling mechanisms that allow us to perform inference on dynamical parameter spaces with minimal tuning. We expand on the most important features in the sections below.

A. Affine-invariant samplers

An affine transformation is one of the form $\vec{\theta} \rightarrow \vec{\zeta} = A\vec{\theta} + b$, where A and b are a constant matrix and vector respectively. Under an affine transformation a probability density $p(\vec{\theta}|y)$ transforms to

$$p_{A,b}(\vec{\zeta}|y) = p(A^{-1}(\vec{\zeta} - b)|y)/\det(A). \quad (5)$$

Such transformations can help to transform difficult-to-sample distributions into easier-to-sample ones. A simple example is a multi-variate Normal distribution. If the dynamical range of the eigenvalues of the covariance matrix is very large, then sampling can be difficult, but any multi-variate Normal distribution can be transformed into a spherical distribution via an affine transformation.

Affine-invariant MCMC is a class of samplers that are designed to have equal sampling efficiency for all distributions that are related by an affine transformation [58, 59]. The sequence of samples in a Markov chain, $\{X(t)\}$, can be written deterministically as a function of a sequence of random variables, $\xi(t)$, which represent the random draws used to propose new points and evaluate the accept/reject decision. Specifically we can always write

$$X(t+1) = R(X(t), \xi(t), p) \quad (6)$$

where p denotes the target density. An affine-invariant sampler has the property

$$R(AX(t) + b, \xi(t), p_{A,b}) = AR(X(t), \xi(t), p) + b, \quad (7)$$

i.e., the sequence of points visited when sampling an affine-transformed density are the affine transformations of the states visited when sampling the original density. If an affine transformation exists that maps the given target density to one which is more straightforward to sample from, an affine-invariant sampler should sample it as efficiently as it could the simpler distribution, so the convergence of affine-invariant samplers is less affected by correlations between the parameters [58].

In practice, this goal is achieved by following an ensemble of points, called walkers, and basing proposed moves on the distribution of other points in the ensemble. In [58], the primary update move is the so-called *stretch-move* proposal. Each walker at state $X_i(t)$ is updated by randomly selecting another walker j and proposing a new value $Y = X_j(t) + Z[X_i(t) - X_j(t)]$, where Z is a random variable drawn from the distribution [59]

$$g(z) \propto \begin{cases} \frac{1}{\sqrt{z}} & z \in [\frac{1}{a}, a] \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

The parameter a can be tuned to improve convergence, but $a = 2$ works well in the majority of applications [58]. The proposed point is accepted with probability

$$\alpha(X_i(t), Y) = 1 \wedge \left\{ z^{d-1} \frac{p(Y)}{p(X_i(t))} \right\} \quad (9)$$

where p is the target density and d is the dimension of the parameter space. This acceptance probability is specific to the stretch proposal distribution given by Eq. (8). For other stretch proposals, the term z^{d-1} must be replaced by $z^{d-2}g(1/z)/g(z)$. Following this scheme, detailed balance is maintained, and it can be proven that affine-invariant samplers converge faster to their target distribution [58]. Below in section III we discuss the extension of the stretch-move proposal to Reversible-Jump MCMC methods. The benefits of running MCMC chains in parallel, combined with a proposal distribution that requires almost no tuning, have contributed to an increasing popularity of affine-invariant samplers. In particular, the `emcee` package [58], has been used widely in Astrophysics and Cosmology [60–63, e.g.].

B. Delayed Rejection

The Delayed Rejection (DR) scheme of sampling was devised in order to improve two aspects of MCMC algorithms. First, it allows for improvements in the acceptance rate of the proposals, yielding “healthier” parameter chains, with better mixing. Secondly, it is more robust against becoming trapped in local maxima of the posterior surface [64–67]. The strategy, as the name suggests, is, at each iteration, instead of immediately rejecting the newly proposed point based on algorithm 1, we keep proposing new points while maintaining detailed balance by computing both the forward and backward transition probabilities. Suppose we are at a point $\vec{\theta}_0$ and use a proposal $q(\vec{\theta}_1|\vec{\theta}_0)$ to propose a new point $\vec{\theta}_1$. The usual acceptance probability, following the notation of eq. (1), is

$$\alpha_1(\vec{\theta}_0, \vec{\theta}_1) = 1 \wedge \left\{ \frac{p(\vec{\theta}_1|y)q(\vec{\theta}_0|\vec{\theta}_1)}{p(\vec{\theta}_0|y)q(\vec{\theta}_1|\vec{\theta}_0)} \right\}, \quad (10)$$

as per eq. (4). If $\vec{\theta}_1$ is rejected, then instead of going back to step 1 of algorithm 1, we propose a new point, $\vec{\theta}_2$, drawn from a proposal distribution $q(\vec{\theta}_2|\vec{\theta}_1, \vec{\theta}_0)$. This proposal distribution may depend only on $\vec{\theta}_1$, but we write it more generally here to allow for the case that the proposal is adapted based on the sequence of steps that have been rejected. The acceptance probability for $\vec{\theta}_2$, $\alpha_2(\vec{\theta}_0, \vec{\theta}_1, \vec{\theta}_2)$, is computed using

$$\alpha_2(\vec{\theta}_0, \vec{\theta}_1, \vec{\theta}_2) = 1 \wedge \left\{ \frac{p(\vec{\theta}_2|y)q(\vec{\theta}_1|\vec{\theta}_2)q(\vec{\theta}_0|\vec{\theta}_1, \vec{\theta}_2) [1 - \alpha_1(\vec{\theta}_2, \vec{\theta}_1)]}{p(\vec{\theta}_0|y)q(\vec{\theta}_1|\vec{\theta}_0)q(\vec{\theta}_2|\vec{\theta}_1, \vec{\theta}_0) [1 - \alpha_1(\vec{\theta}_0, \vec{\theta}_1)]} \right\}. \quad (11)$$

If $\vec{\theta}_2$ is rejected, further steps can be included and each step adds additional proposal and rejection-probability terms to the numerator and denominator of the acceptance probability. For example, the three step acceptance probability, $\alpha_3(\vec{\theta}_0, \vec{\theta}_1, \vec{\theta}_2, \vec{\theta}_3)$ is the minimum of one and

$$\frac{p(\vec{\theta}_3|y)q(\vec{\theta}_2|\vec{\theta}_3)q(\vec{\theta}_1|\vec{\theta}_2, \vec{\theta}_3)q(\vec{\theta}_0|\vec{\theta}_1, \vec{\theta}_2, \vec{\theta}_3)}{p(\vec{\theta}_0|y)q(\vec{\theta}_1|\vec{\theta}_0)q(\vec{\theta}_2|\vec{\theta}_1, \vec{\theta}_0)q(\vec{\theta}_3|\vec{\theta}_2, \vec{\theta}_1, \vec{\theta}_0)} \times \frac{[1 - \alpha_1(\vec{\theta}_3, \vec{\theta}_2)] [1 - \alpha_2(\vec{\theta}_3, \vec{\theta}_2, \vec{\theta}_1)]}{[1 - \alpha_1(\vec{\theta}_0, \vec{\theta}_1)] [1 - \alpha_2(\vec{\theta}_0, \vec{\theta}_1, \vec{\theta}_2)]} \quad (12)$$

The proposal q can be different at each step, as long as the relevant proposal density is used in eq. (11). For example, in [64] the proposal is built upon a Gaussian mixture model that tries further points in the parameter space with the aim of efficiently exploring multiple modes of the posterior distribution. As the number of steps in the DR scheme becomes arbitrarily large the acceptance probability slowly approaches zero. This algorithm is

also limited in practice by high computational requirements, since at every delayed rejection step we need to evaluate a new likelihood and compute the backwards probability (the $\alpha_1(\vec{\theta}_2, \vec{\theta}_1)$ from eq. (11)). Nevertheless, the DR scheme offers many advantages, and despite the computational cost, it is very useful when the posterior surface exhibits high dimensionality, and when acceleration techniques are available. These, for example, might include the use of Graphical Processing Units (GPUs), and/or heterodyned likelihoods [68]. In our implementation here, we follow closely the one in [64], for improving the acceptance rate of the *between-model step* of the Reversible Jump algorithm (see section II F). As already mentioned, the Reversible Jump MCMC allows for sampling dynamical parameter spaces. In the special case of nested models, such as the case of searching multiple signals in the LISA data, proposing the ‘birth’ of a signal out of a very wide prior can be very inefficient. A delayed rejection scheme alleviates this problem, by effectively performing a small search around the first set of rejections, increasing the chances of finding a good signal candidate, and thus improving the mixing of the chains.

C. Multiple Try Metropolis

The Multiple Try Metropolis (MTM) [69–72] is a subclass of the implementation of the MH algorithm, which is based on the idea of generating a number of proposals for each individual current state, and then selecting one of them based on their importance weight. In proposing a move from $\vec{\theta}_{t-1}$, a set of N possible new points, $\{y_k\}$, are drawn from a proposal distribution $q(y|\vec{\theta}_{t-1})$ and are assigned weights $w_i = w(y_i|\vec{\theta}_{t-1})$ using a weight function $w(\cdot|\vec{\theta}_{t-1})$. One of these proposed points, y_J , is selected with probability given by the normalised weight

$$\bar{w}_i = \frac{w_i}{\sum_{k=1}^N w_k}. \quad (13)$$

To compute the acceptance probability we need to draw $N - 1$ points, $\{x_i, i = 1, \dots, N - 1\}$, for the reverse move from the proposal $q(x|y_J)$, and assign weights $w(x|y_J)$. We then set $\vec{\theta}_t = y_J$ with probability

$$\alpha(\vec{\theta}_{t-1}, y_J) = 1 \wedge \left\{ \frac{w(y_J|\vec{\theta}_{t-1}) + \sum_{k=1, k \neq J}^N w(y_k|\vec{\theta}_{t-1})}{w(\vec{\theta}_{t-1}|y_J) + \sum_{k=1}^{N-1} w(x_k|y_J)} \right\}. \quad (14)$$

and set $\vec{\theta}_t = \vec{\theta}_{t-1}$ otherwise [69]. This procedure will satisfy detailed balance if the weight function is chosen such that

$$p(\vec{\theta}_0|y)q(\vec{\theta}_1|\vec{\theta}_0)w(\vec{\theta}_1|\vec{\theta}_0) = p(\vec{\theta}_1|y)q(\vec{\theta}_0|\vec{\theta}_1)w(\vec{\theta}_0|\vec{\theta}_1). \quad (15)$$

This will be satisfied by a weight function of the form

$$w(\vec{\theta}_t|\vec{\theta}_{t-1}) = p(\vec{\theta}_t|y)q(\vec{\theta}_{t-1}|\vec{\theta}_t)\xi(\vec{\theta}_{t-1}, \vec{\theta}_t), \quad (16)$$

where $\xi(\vec{\theta}_{t-1}, \vec{\theta}_t)$ is any symmetric function, i.e., $\xi(\vec{\theta}_{t-1}, \vec{\theta}_t) = \xi(\vec{\theta}_t, \vec{\theta}_{t-1})$, $\forall \vec{\theta}_t, \vec{\theta}_{t-1} \in \mathcal{D} \subseteq \mathbb{R}^d$, with d being the dimensionality of the problem at hand. The detailed balance condition can also be satisfied by a weight function of the form

$$w(\vec{\theta}_t|\vec{\theta}_{t-1}) = \frac{p(\vec{\theta}_t|y)}{q(\vec{\theta}_t|\vec{\theta}_{t-1})}. \quad (17)$$

Making this choice and additionally using a proposal function that is independent of the current point, $q(\vec{\theta}_t|\vec{\theta}_{t-1}) = q(\vec{\theta}_t)$ only, we obtain the *Independent MTM* algorithm [69]. When using the independent MTM algorithm detailed balance is maintained when the same set of points is used for the reverse proposal as for the forward proposal, which saves the evaluation of $N - 1$ posterior densities.

The base MTM is currently implemented in **Eryn** with options for the Independent MTM algorithm and symmetric proposals. For a symmetric proposal distribution, $q(\vec{\theta}_{t-1}|\vec{\theta}_t) = q(\vec{\theta}_t|\vec{\theta}_{t-1})$, eq. (15) can be satisfied using the weight function $w(\vec{\theta}_1|\vec{\theta}_0) = p(\vec{\theta}_1|y)$. In this case, we still need to draw separate samples for the reverse step (unlike in the Independent MTM case).

Generating a large number of candidate points yields certain advantages. As expected, the first advantage is the fact that there is usually very good coverage of the parameter space. The second is that the implementation of the MTM usually results in chain states with very low correlation between them. Nevertheless, as for Delayed Rejection, this algorithm requires increased computational resources, since multiple likelihoods/posterior densities have to be evaluated at each iteration of the chain. This cost can be offset in cases where the computations can be parallelized, for example using either CPU or GPU acceleration.

D. Adaptive Parallel Tempering

The concept of Parallel Tempering was introduced in order to efficiently sample surfaces with high multimodality [73–75]. The idea is based on a transformation of the posterior density to a density with a different temperature, T , defined by

$$p_T(\vec{\theta}|y) \propto p(y|\vec{\theta})^{1/T} p(\vec{\theta}). \quad (18)$$

For $T = 1$ this is the target posterior density. In the limit $T \rightarrow \infty$ it is the prior density. Intermediate temperatures “smooth out” the posterior by reducing the contrast between areas of high and low likelihood.

In parallel tempering, a set of Markov chains are constructed in parallel, each one sampling the transformed posterior for a different temperature T . These chains periodically exchange information. The idea is that the hottest chains explore the parameter space more widely, and information about areas of high likelihood that they

encounter propagate to the colder chains. Information is exchanged by proposing swaps of the states between the different chains. If two chains are sampling from target densities $p_1(\vec{\theta})$ and $p_2(\vec{\theta})$ respectively, then the transition probability for chain 1 in the swap is $p_1(\vec{\theta}_0 \rightarrow \vec{\theta}_1) = p_2(\vec{\theta}_1)\alpha(\vec{\theta}_0, \vec{\theta}_1)$. Detailed balance is thus maintained by accepting the swap with probability

$$\alpha(\vec{\theta}_0, \vec{\theta}_1) = 1 \wedge \left\{ \frac{p_1(\vec{\theta}_1)p_2(\vec{\theta}_0)}{p_1(\vec{\theta}_0)p_2(\vec{\theta}_1)} \right\}, \quad (19)$$

which for the specific case of swapping between two tempered chains i and j when doing parallel tempering is

$$\alpha_{i,j} = 1 \wedge \left\{ \left(\frac{p(y|\vec{\theta}_i)}{p(y|\vec{\theta}_j)} \right)^{\beta_j - \beta_i} \right\}, \quad (20)$$

with $\beta_i = 1/T_i$ being the inverse temperature, and $\vec{\theta}_i$ the given parameter state for the i -th chain.

The temperature ladder T_i should be chosen in order to maximize the information flow between chains of different temperatures, so as to encourage the efficient exploration of the complete parameter space. Typically, this ladder can be static or dynamically adjusted during the sampling procedure. In **Eryn** we have adopted the procedure of [75], which adapts the temperature ladder based on the swap acceptance rate calculated directly from the chains. Ideally, one should aim for equal acceptance ratio between every pair of neighboring tempered chains, thus tuning their log-temperature-difference $S_i \equiv \log(T_i - T_{i-1})$, according to the swap acceptance rate from eq. (20):

$$\frac{dS_i}{dt} = \kappa(t) [\alpha_{i,i-1}(t) - \alpha_{i+1,i}(t)], \quad (21)$$

where $\kappa(t)$ tunes the timescale of the evolution of the temperatures. The function $\kappa(t)$ can be chosen depending on the desired behavior of the procedure. In [75] a hyperbolic dependence on the t state is chosen, in order to suppress large dynamic adjustments on long timescales. This setup is the default option in **Eryn**, but it can be customized. This process is more straightforward for ensemble samplers, where multiple walkers are used, simply because one can get an estimate of the acceptance rate directly from the particular state of the walkers at any given time step t . Otherwise, the acceptance rate is computed after iterating for a predefined number of steps, which can be chosen by the user for the given problem at hand. It can be proven [75], that the temperature ladder will converge to a particular stable configuration. One should only use this scheme for the initial burn-in stage of sampling, and then continue with a stationary ladder for the rest of the analysis.

E. Marginal posterior calculation for model selection

One of the most frequently encountered problems in physics, and in science in general, is that of model or variable selection, i.e., identifying the model best supported by the observed data. Working in a Bayesian framework, comparison between different hypotheses may be done by computing their evidences or marginal posteriors [46] and evaluating the Bayes Factor:

$$B_{12} = \frac{p(\vec{y}|\mathcal{M}_1)p(\mathcal{M}_1)}{p(\vec{y}|\mathcal{M}_2)p(\mathcal{M}_2)}, \quad (22)$$

where the term $p(\mathcal{M}_i)$, is the prior probability assigned to the model \mathcal{M}_i .

The marginal posterior density, or evidence, is given by the integral of eq. (2) and is in general quite challenging to compute. For some high signal-to-noise ratio (SNR) cases it can be approximated if the covariance matrix Σ of the parameters for all candidate models \mathcal{M} are known. This approach is called the *Laplace approximation* [46, 76]. However, this is only an approximation, and it sometimes fails for models with weak support from the data [77] (in particular when the posterior cannot be approximated by a multivariate Gaussian at $\vec{\theta}_{\text{MAP}}$).

When using parallel tempering IID, it is possible to compute the evidence by a procedure known as *thermodynamic integration* [78]. We define a continuous distribution of evidences based on the target distribution for a chain with inverse temperature $\beta = 1/T$ via

$$\mathcal{Z}_{i,\beta} = \int p(y|\vec{\theta}, \mathcal{M}_i)^\beta p(\vec{\theta}) d\vec{\theta}. \quad (23)$$

For $\beta = 0$ the chain is sampling the prior and therefore $\log \mathcal{Z}_{i,0} = 0$. For $\beta = 1$ we are sampling the target distribution and $\log \mathcal{Z}_{i,1} = \log \mathcal{Z}_i$. Additionally we have

$$\begin{aligned} \frac{d \log \mathcal{Z}_\beta}{d\beta} &= \int \log[p(y|\vec{\theta}, \mathcal{M}_i)] p(y|\vec{\theta}, \mathcal{M}_i)^\beta p(\vec{\theta}) d\vec{\theta} \\ &\equiv \mathbb{E}_\beta[\log p(y|\vec{\theta}, \mathcal{M}_i)]. \end{aligned} \quad (24)$$

From this we deduce

$$\log \mathcal{Z}_i = \int_0^1 \mathbb{E}_\beta \left[\log p(y|\vec{\theta}, \mathcal{M}_i) \right] d\beta. \quad (25)$$

The expectation value is over the distribution being sampled by the chain at temperature β and so can be computed by averaging over the posterior samples [75, 78, 79]. The integral can then be evaluated using standard methods, for example the trapezium rule, using the full ladder of temperatures. This approach generates reliable evidences, with accuracy limited only by the number of temperatures being sampled, and the efficiency of the sampling of the parameter space Θ by the chains. Since its first introduction, there have been many applications of this approach, and in particular, there is extensive usage in GW astronomy [75, 80–83].

F. Reversible Jump

Another approach to the model selection problem is to follow a Reversible Jump (RJ) MCMC strategy, which can dynamically estimate the most probable hypotheses given the data [30]. The RJ-MCMC is a generalization of the MH algorithm that allows trans-dimensional proposals. Thus, the model order is considered a free parameter which is fitted together with the parameters of the individual models. The most widespread variation of the algorithm uses a two-stage procedure. The first stage or *in-model step*, uses the standard MH algorithm to update all the parameters $\vec{\theta}_k$ for the given model k . The second stage or *between-model step* proposes to update the model state k to a new model state l . Parameters $\vec{\theta}_l$ for the new model are also proposed. The newly proposed state l , is accepted with a probability defined by [84]

$$\alpha' = 1 \wedge \left\{ \frac{p(\vec{\theta}_l|l, y)g'(u_l)}{p(\vec{\theta}_k|k, y)g(u_k)} |\mathbf{J}| \right\}, \quad (26)$$

where

$$p(\vec{\theta}_k|k, y) = p(y|\vec{\theta}_k, k)p(\vec{\theta}_k, k)p(k) \quad (27)$$

with $p(y|\vec{\theta}_k, k)$ the likelihood for model k , $p(\vec{\theta}_k|k)$ the prior on the parameters $\vec{\theta}_k$ in model k , and $p(k)$ the prior for the model state k . The term $g'(u_l)|\mathbf{J}|/g(u_k)$ arises because of the need for *dimension matching* between the different model states. In general, we can define a move between model states in terms of a deterministic invertible mapping, $\vec{\theta}_k = q(\vec{\theta}_l, u_l)$ with inverse $\vec{\theta}_l = q'(\vec{\theta}_k, u_k)$, that is a function of the parameters and two sets of random variables, u_k and u_l , drawn from distributions $g(u_k)$ and $g'(u_l)$ [30, 84]. The term $|\mathbf{J}|$ is the Jacobian defined by this invertible mapping

$$|\mathbf{J}| = \left| \frac{\partial(\vec{\theta}_l, u_l)}{\partial(\vec{\theta}_k, u_k)} \right| \quad (28)$$

and the term $g'(u_l)/g(u_k)$ plays the role of the proposal ratio in the standard MH acceptance probability. Dimension mapping means that $\dim(\vec{\theta}_l) + \dim(u_l) = \dim(\vec{\theta}_k) + \dim(u_k)$.

Using RJ-MCMC introduces additional computational cost at each MCMC iteration, as well as technical challenges in implementation. Luckily, implementation can be easier when sampling *nested models*. This refers to problems where more complicated models contain their simpler counterparts. Examples of such cases are fitting polynomial models, which differ only in the highest order to be determined, or detection problems where multiple similar signals are potentially present the data. In such cases, the between-model step can always be formulated such that the Jacobian of eq. (28) becomes unity, and eq. (26) simplifies to the ratio of posteriors accounting for any differences in prior and proposal volumes [22, 85, 86].

After running RJ-MCMC, the Bayes Factor can be approximated by the ratio of the number of iterations spent within each model:

$$B_{12} = \frac{\# \text{ of iterations in model } \mathcal{M}_1}{\# \text{ of iterations in model } \mathcal{M}_2}. \quad (29)$$

This algorithm has proven to be robust for evaluating high-dimensional competing models, and has been quite successful in tackling data analysis problems in GW astronomy [22, 77, 87] as well as areas spanning physics and signal processing [e.g., 88–90]. However, designing an efficient RJ-MCMC algorithm can be quite challenging. The first challenge is to choose suitable proposal distributions, which can greatly affect the convergence of the algorithm. In situations where the models are nested, it is both tempting and convenient to take the proposal to be the same as the prior distribution of the parameters. As an illustrative example, we refer to section III C 1, which describes a toy problem of searching for Gaussian pulses in noisy data. There, the parameters of the individual pulses are the amplitude and location of the pulse described by their (x, y) coordinates. In order to search for those signals, the prior on their location must be wide enough to include the complete data set (see figure 1a). A birth proposal based on the prior would inevitably be quite inefficient, simply because the chance of proposing a good source candidate is small, especially if the proposal distribution is flat across (x, y) . We treat the above problem as a motivation to adopt efficient proposals with minimal tuning in `Eryn`, which we further discuss in section III. The second major challenge, which of course depends on the given problem at hand, arises from the samplers’ capability to explore a multi-modal dynamical parameter space. We discuss our strategy to overcome that challenge in section III.

III. ERYN: GATHERING ALL THE PIECES TOGETHER

All the different algorithms described in previous sections can be extremely useful in tackling different kinds of problems that require sampling. In Gravitational Wave Astronomy, we encounter such problems far too often, where dynamical parameter spaces require vast computational resources in order to be explored efficiently. Motivated by those problems, we have implemented a new toolbox that combines all these techniques to enhance the capabilities of an MCMC sampler. We have named this package `Eryn` [91], borrowing the name from the Tolkien mythos [92]. The analogy has its basis in the idea of a forest: within a forest you have trees which correspond to different walkers, a.k.a. *Ents*, in an ensemble MCMC sampler. On each tree there are branches that represent the various types of models used to fit the data. For example, in the case of GW global fitting for LISA, we can imagine using the Galactic binaries as one branch and massive black hole binaries as another

branch. Each branch has leaves, which represent the individual instances of each model. In the LISA example, leaves would represent the individual Galactic binaries or massive black hole binaries. And finally, to zoom out, when adding tempering capabilities, we can think of groups of walkers in each temperature taking the form of many forests (of walkers) located within different temperature climates.

We adopt the architecture of ensemble samplers, and in particular the one of `emcee` [58]. Having multiple walkers running in parallel is ideal for efficiently sampling the parameter spaces using techniques such as parallel tempering, as described in IID (also see section II A). In this setting, we evolve n_w walkers per temperature T_i , where each walker follows a Reversible Jump MCMC (see section II F), mapping a parameter space of altering dimensionality. In practice, walkers in higher temperatures sample the dynamic parameter space with fewer model components as the penalty from higher prior volume is not compensated by the smoother annealed likelihood. In other words higher temperatures have a sharper Occam’s razor: the data can be explained with models that are simpler, or lower-dimensional. The highest temperature chain samples the prior on the model space (provided that $T_{\max} = \infty$). More details will be given in section III C.

As already mentioned in section II F, Reversible Jump algorithms are extremely challenging to tune, even for simpler classes of problems. One of the major challenges is the low acceptance rate for the between-model proposal, i.e., when we propose a new state where the parameter dimensionality differs. In cases of signal search and detection (which is a nested model situation), it is convenient to set the proposal corresponding to a “birth” move to be the same as the prior distribution. In order to accommodate all possibilities for the signals present, the prior densities are usually quite wide, and thus accepting a new higher dimensional state becomes quite improbable. For that reason, within `Eryn`, we have implemented a Delayed Rejection scheme with the aim of improving this acceptance ratio. When proposing $\vec{\theta}_l$ for a higher-dimensional model l , we do not reject immediately, but rather make new delayed rejection proposals around the first rejected point $\vec{\theta}_l$, using the given *in-model step* update proposal. This, in principle, allows the sampler to explore around $\vec{\theta}_l$ before rejecting the new state [64], which in turn improves the *between-model step* acceptance rate and produces healthier MCMC chains.

The Delayed Rejection scheme, as described in section II B, requires a serial computation of the delayed rejection acceptance ratio for walkers where the newly proposed state has been rejected. This scheme of calculating costly likelihoods sequentially in a loop during the between-model step, can lead to a computational bottleneck of the MCMC process. This is especially true for the LISA Global Fit problem, where multiple binary waveform signals are present in the data-stream. Then, the computational time for each RJ-MCMC iteration is sig-




FIG. 1: Searching for 2D Gaussian pulses in the presence of Gaussian noise. Panel (a): the simulated data, which consists of injections of 25 pulses in Gaussian noise with $\sigma_n = 0.2$. Panel (b): the distribution of the model order, obtained by exploring the dynamical parameter space with **Eryn**. The true value is marked with a dashed red line. For this toy investigation, the correct number of simulated components is recovered. Panel (c): the 2D posterior densities for the parameters of the k Gaussian peaks (see text for details).

nificantly increased, since the progress will be halted until all walkers have gone through their respective Delayed Rejection process, which requires evaluation of new waveforms at each step. For the reasons summarized here, we have not used the Delayed-Rejection scheme for our analysis in section IV, and have resorted to the Multiple Try scheme. However, the Delayed-Rejection scheme, as explained in section II B, has been implemented in the **Eryn** package.

The Multiply Try scheme was essentially implemented in order to facilitate use of a parallelized likelihood framework. Parallelization is naturally compatible with Multiple Try MCMC as multiple proposals are made for each individual walker, which allows for the parallelized evaluation of proposal distributions, likelihood functions, and acceptance probabilities. Under these parallelized settings, one proposal can act as many when compared to the usual serial evaluation of proposals, allowing for better chain mixing in situations where proposals are infrequently accepted. That being said, it is still important to choose a good proposal distribution, for both the in-model and between-model RJ-MCMC steps, which we discuss further below.

A. Choosing efficient Proposal distributions

In the previous sections, we briefly discussed some of the challenges in choosing efficient proposal distributions for both the in-model and between-model steps of the RJ algorithm. For the in-model case, the challenge arises from the fact that it is sometimes impractical, or even unfeasible, to define a well-tuned proposal for each of the possible models that could represent the data. Using again the example of LISA data, one would need to te-

diouly design an effective proposal distribution for the thousands of overlapping binary signals in the data. On the other hand, for the case of the between-model step, choosing proposals from the prior distribution, especially if it is highly uninformative, can be very inefficient for Reversible Jump sampling. For **Eryn**, in order to tackle those issues, we have implemented the *Group Proposals* explained below for addressing the within-model proposals in RJ, as well as a scheme to design an efficient proposal for birth moves during the between-model step, which is based on normalizing flows.

1. Group Proposals

In section II A, the stretch-move proposal was introduced and discussed. One of the obvious advantages of such a scheme of proposing new MCMC samples is that it requires minimal tuning [58]. However, it does not extend well in its simplest form to the generalized Reversible Jump MCMC. The stretch proposal is based on the idea that the ensemble of points (X_j) is sitting on the same posterior mode as the current point (X_i). In a nested model situation where both the model count and the individual model parameters change, each point may lie on a different posterior mode representing a different set of leaves (sources) in the data. This can be alleviated by applying the stretch move to individual leaves within each branch of each walker, but there is still an issue of identifying leaves in different walkers that lie on the same posterior mode. The stretch proposal will technically still work when mixing leaves in different posterior modes, but the acceptance fraction will be negatively affected. However, within the stretch proposal formalism, the choice of X_j is customizable. The key to maintaining

detailed balance is that X_j cannot depend on X_i , and X_j cannot be updated in the same iterative step as X_i [59].

We leverage this property to design a new type of stretch move that can handle Reversible Jump setups while maintaining a small number of tuning parameters. We call this proposal the “group” proposal. The group proposal chooses a single leaf from a stationary group, $\{X_j\}$, that is fixed for many proposed updates. The stationary group is updated after a large number of sampler iterations and we make sure that detailed balance is maintained during the update. We update every leaf within every branch of every walker at each iteration and repeat many iterations between updates of the stationary group.

The appropriate stationary group varies from problem to problem. The goal is to set a group that resembles as best as possible the posterior modes of the current leaves and then draw from it strategically so that the drawn point is likely (but not guaranteed) to lie on the same mode as the leaf that is currently being updated, X_i . In the example of the LISA Galactic binaries analysis, we set our stationary group to the full set of leaves (binaries) across all walkers at a specific temperature of the sampler at the end of a given iteration. Then, at proposal time, we efficiently locate the $\sim n_w$ points in the stationary group that are closest to X_i from based on their initial frequency parameter. We then draw X_j from this group. The hope is that some percentage of the n_w drawn points will lie on the posterior mode on which X_i sits. The exact percentage will vary depending on how close the posterior modes are to each other and how many model instances exist in the sampler that include this specific mode. For low SNR binaries, for example, a source may exist in some walkers and not others, making it harder to map its posterior mode with the current group of stationary points.

The performance of group proposals is highly situation- and/or model-dependent. With individual source posterior modes that are well separated and easy to define in terms of separation, the performance will approach the performance of the base stretch proposal in non-Reversible Jump MCMC because the stationary group will well represent the specific posterior mode on which X_i is located. As the parameter space becomes more crowded and/or separation (distance) metrics become harder to define, the performance of group proposals will worsen.

2. Learning from the data

The second improvement concerns the between-model step of the RJ-MCMC. As mentioned earlier, for the case of nested models, it is often convenient to draw “birth” candidates directly from the prior distribution of parameters of the given model. This practice can be quite ineffective in terms of acceptance rate. As an example we can again use the LISA data-set case. The UCBs are

distributed within the Galactic disk, congregated mostly around its Center [93], therefore, adopting a proposal based on an uninformative uniform prior across the sky, would waste computational resources exploring a part of the parameter space with low probability mass. A proposed solution is to use an informative prior derived from the spatial distribution of binaries in the Galaxy [21]. In our work here, we have chosen an alternative data-driven route, based on the actual residual data after a burn-in period of the RJ-MCMC, which we describe below.

After a sufficient number of RJ-MCMC iterations, we can extract a subset of sources from nested models which are constantly present in almost all walkers of our cold chain. In other words, we can find and subtract the brightest sources from the data, and then allow for another burn-in period on the resulting residuals. This allows the sampler to explore the remaining parameter space more easily, thus providing a good initial estimate for the weak signals possibly buried in the noisy residual data. We can then use those samples to construct a proposal density which will help us search for good candidates for those weak signals, without excluding the rest of the parameter space. This can be accomplished by learning the posterior distribution of the parameters run on the residual data mentioned above. The most efficient way to fit to the generic distribution, is to use invertible transformations such as normalizing flows (for example [94, 95]). The methods work in the following way: we sample from the base distribution (which is usually chosen to be Normal) and we fit the transform by optimizing Kullback-Leibler divergence between the transformed distribution and the distribution that we want to estimate. After the fit has converged we can draw samples from the Normal distribution and transform these to samples from the distribution fitted to all residuals and use it as a proposal. We will cover this method in more detail in a separate paper.

B. Implementation

In this section, we discuss the main implementation details of **Eryn**. We refer the interested reader, or user, to the **Eryn** documentation for more exhaustive information and examples [96].

The goal of **Eryn** is to produce a sampler that can handle all (or most) cases of MCMC sampling ranging from basic, non-tempered, single-model type, single-model instance posterior estimation to the full reversible jump MCMC with tempering, multiple model types, and adjustable model counts, as well as everywhere in between. In the basic case, **Eryn** aims to be a close replica of **emcee** trying to maintain as much simplicity as possible. At the complicated end of the spectrum, **Eryn** attempts to provide a common interface and underlying infrastructure for the variety of problems that may arise, allowing the user to maintain usage of the majority of the code from project to project, focusing on changing only the specific

parts of the code that are difficult to implement or require special treatment for each specific problem. Since `Eryn` is effectively an enhanced version of the `emcee` package, the overall structure of `emcee` is strongly maintained. Like in `emcee`, “State” objects move coordinate and likelihood information around the ensemble sampler, storing information in a similar back end object either in memory or `HDF5` files. Additionally, the interface used for adding proposals has remained.

The various enhancements discussed in this work, including tempering, reversible jump moves, multiply try MCMC, etc., are all implemented within the `emcee`-like structure. This involved two main changes. First, the State objects have been scaled to hold information necessary for reversible jump MCMC: temperature information, prior information, and efficient and concise containers for multiple types of models with an adjustable number of individual model instances. Second, the reversible-jump proposal has been added as a proposal base, similar to the use of the ‘MH’ or ‘RedBlue’ moves within `emcee`. Beyond these main enhancements, there are also a variety of smaller, but useful, additions to `Eryn` that help the user build a variety of analysis pipelines. These include stopping or convergence functions, functions to periodically update the sampler setup while running, objects to carry special information through the sampler, and aids for coordinate transformation.

C. Toy Examples

In this section, we present a series of working examples for `Eryn`. We begin with simple problems, such as searching for simple signals in noisy data, with the aim of demonstrating the performance of this toolbox in a dynamical parameter space. The impact of the different enhancements discussed in section III will be assessed and discussed. Finally, in section IV, we will apply this machinery to more realistic problems in Gravitational-Wave astronomy.

1. Searching for pulse signals in Gaussian noise

In this first example, we explore the capabilities of `Eryn` in a simplified application, commonly encountered in physical sciences. We perform a search for Gaussian pulses in a simulated 2D data-set, in the presence of Gaussian noise with variance $\sigma_n = 0.2$. We generate 25 pulses randomly distributed on the $x - y$ plane with all pulses contained within $x, y \in [-10, 10]$ (see figure 1a), and amplitude uniformly drawn from $\mathcal{U}[0.7, 1.5]$. The amplitude of each pulse is considered a free parameter to be estimated, in addition to the Cartesian coordinates of the centres. The pulses’ width was kept fixed to $\sigma_p \times \delta_{ij}$, with $\sigma_p = 0.2$, for the sake of simplicity. Thus, we are required to estimate N_p , the total number of pulses in the data, and also estimate the parameters for each in-

dividual signal k : $\vec{\theta}_k = \{A_k, x_k, y_k\}$. The noise variance σ_n is estimated as part of the fit.

The analysis of this problem is performed using the adaptive parallel tempering scheme of section IID and the Reversible Jump MCMC proposals (section IIF). The in-model proposals for each model component are Gaussians, with a diagonal covariance matrix $\Sigma = 10^{-4}\delta_{ij}$. This proposal is not tuned during sampling. The priors for the parameters are quite wide, covering the entire range of the data, while the prior on the number of pulses k is set to $k \sim \mathcal{U}[0, 50]$. With the above settings, we obtain the results summarized in figures 1 and 2. In figure 1b we plot the most probable number of Gaussian pulses present in the data, or in other words, the most probable model given this particular data-set. It is clear that for the given level of noise, it is straightforward to recover the true number of signals. The noise variance is also estimated accurately as $\sigma_n = 0.2 \pm 2 \times 10^{-3}$. In figure 1c we plot the posterior densities for the parameters of all pulses recovered, while we also mark the true injected values. Figure 1c shows the trans-dimensional MCMC chains “stacked” over all samples of both model order and model-parameters. As already mentioned, in this simplified scenario all signals have similar value for the amplitude, thus the almost uni-modal marginal on A_k . This illustrative example is useful as an introductory application to the more complicated case of detection in Gravitational Wave astronomy presented below, in section IV.

In figure 2, three diagnostic quantities for this run are shown. In the top panel, the evolution of temperatures is presented. Following the recipe of [75], we control the distances between each temperature chain based on their in-between swap acceptance rate, computed from eq. (21). The tuning term $\kappa(t)$ is set to $\kappa(t) = t_0 / (\nu(t + t_0))$, with the *adaptation lag* $t_0 = 10^4$ and the *adaptation time* $\nu = 10^2$. The middle panel shows the evolution of the swap acceptance rate per number of walkers between the chains running at different temperatures. After $\sim 10^5$ sampler iterations, the system converges to an equilibrium, where the rate of swapping states reaches a single value across the temperature range. In the bottom panel we show the acceptance rate for the in-model step of the algorithm, for all temperatures. As expected, after temperature equilibrium at $\sim 10^5$ samples, the acceptance rate converges to a (different) value for each temperature, which is higher for higher temperatures (smoother posterior surfaces are easier to explore).

Finally, it is interesting to investigate how the sampled dimensionality of the problem varies at different temperatures. In figure 3, we plot the posterior on the number of pulses at each temperature. As expected, higher temperature chains tend to favour lower model dimensionality and the T_∞ chain samples the prior on k . This can be attributed to the choice of priors and “birth proposal” distributions for both the signal parameters and k . The likelihood is down-weighted at higher temperatures,




FIG. 2: *Top panel:* The evolution of the temperature chains running in parallel for the toy problem of searching for 2D Gaussian pulses in Gaussian noise. The different colors indicate the the initial temperature chain index. Following the parallel tempering recipe of [75], the temperature ladder is tuned according to eq. (21), and the chains start to converge after $\sim 10^4$ iterations. *Middle panel:* The evolution of the swap acceptance rate $\alpha_{i,j}$ described in eq. (20), per number of walkers n_w . For this run we have used $n_w = 10$ walkers. After 10^5 iterations, the swap acceptance rate converges to a single (different) value for every temperature chain. *Bottom panel:* The ‘‘in-model’’ acceptance rate per temperature chain, given by eq. (4).

making it harder to overcome the Occam penalty from including extra parameters in the model. This means quieter sources are less likely to be added and the preferred models have fewer sources.




FIG. 3: Posterior on the number of Gaussians, k , at each temperature T_i , for the toy problem of section III C 1. The different colors indicate the the initial temperature chain index. Darker colors corresponds to colder chains and vice versa.

2. Modelling power spectra: Searching for the optimal number of B-spline knots.

One of the most common problems in signal processing is the characterization of the spectra of the data. This is often done by adopting spectral models and fitting the spectra directly in the frequency domain. This methodology is used when the signal of interest has stochastic properties. Examples from GW astronomy, include the measurement of stochastic signals with astrophysical, or cosmological origin [97, 98]. There are many examples of possible stochastic signals for LISA [16, 20, 93, 98]. Searching for signals with stochastic properties requires flexible spectral models, both for the observatory instrumental noise, and the measured stochastic signal. For these reasons, it is sometimes convenient to adopt a versatile model, such as one that is based on B-spline interpolation schemes.

B-splines are a geometrical modeling tool, and have proven to be very useful for modelling or generating smoother representations of data. They are piecewise polynomial curves with a certain number of continuous derivatives, and can be parameterised in various ways [101]. For this application, we follow [102], and choose to work with cubic-spline interpolation, using the corresponding SCIPY library [103]. The procedure starts by selecting a number of control points, or *knots*, with a given position and amplitude, which the smooth polynomial curve crosses and at which there is a change in the first non-continuous derivative. One of the challenging problems using such methods, is to choose the optimal number of spline knots for fitting the data, without over-

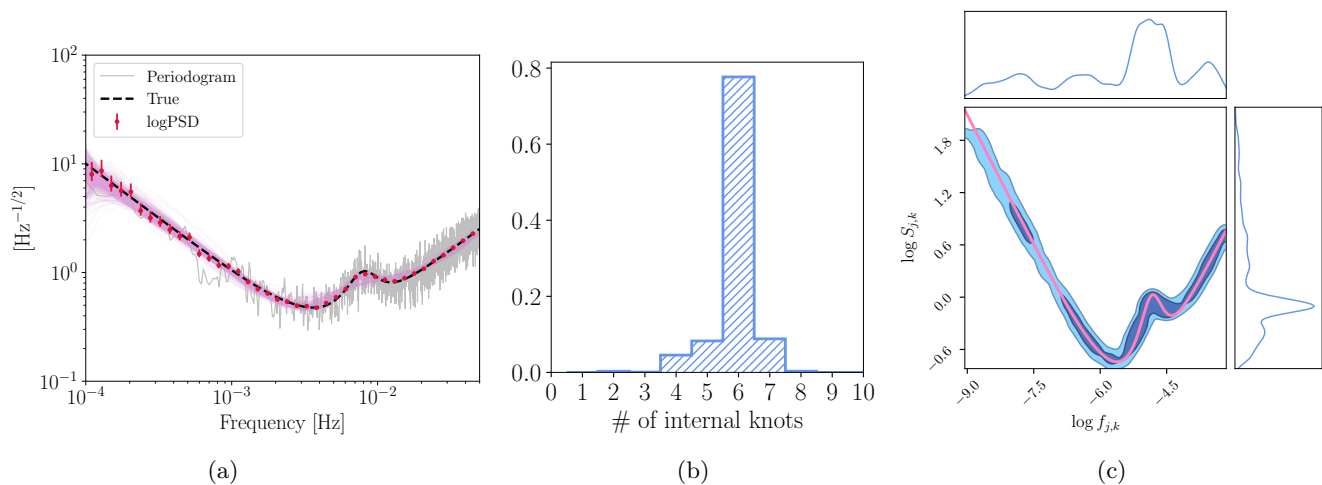


FIG. 4: Results for power spectra modelling with a shape agnostic model. (a) The simulated data (gray), generated from the theoretical model (dashed black line). The PSD computed on an equally-spaced logarithmic grid with the method of [99, 100] is also shown (red data points and their $1\text{-}\sigma$ errors). The pink solid lines represent models drawn randomly from the posterior chains. (b) The optimal B-spline knots estimated by the dynamical parameter estimation procedure. As shown from this histogram, the optimal interior knot count for this data converges to six, corresponding to eight total knots including the two edge knots. (c) Stacking the MCMC chains for all models. It is evident from this figure that we essentially “scan” the true noise shape (pink solid line), by placing knots across the frequency range (see text for more details).

fitting. This is a model selection problem that can be easily solved with dynamical algorithms such as the one presented here. For our next example, we generate time-series data directly from a theoretical model PSD. The simulated data are shown in figure 4a. We then use the machinery of `Eryn` to find the optimal dimensionality for the problem, together with the best-fit parameters for the knots. To ease the computational complexity, we compute the PSD of the time series using the methodology presented in [99]. In more detail, we begin by choosing a new frequency grid, to which we compute the PSD using the *optimal* number of averaged segments for each given frequency. We essentially split the time-series data at the maximum number of segments that the given choice of window and percentage of data overlap permits, which will allow us to estimate the PSD at each frequency bin with minimal variance. By carefully choosing the window function and distance between the data points, one can generate a numerical spectrum estimate with minimal correlations between them (see red data-points in figure 4). For more details we refer the reader to [99, 100]. Finally, we also keep two knots fixed at the edges of the spectra, allowing the sampler to estimate only their amplitude, while the rest of the knot parameters (and their number) are left to be estimated from the data.

For the spline knot positions, $\{\log f_{j,k}\}$, and amplitudes, $\{\log S_{j,k}\}$, we adopt uniform priors that cover the complete parameter space. Here, the j index corresponds to the knot number for the given model order k . We also use a ladder of 10 temperatures, with 10 walkers each, while maintaining the same settings for the adaptivity

of the temperatures as in section III C 1. Each walker is initialized at a random point on the parameter space, after drawing the dimensionality k of the model from $k \sim \mathcal{U}[1, 20]$. We adopt a Gaussian likelihood, with its logarithm written as

$$\log p(D|\vec{\theta}_k) \propto -\frac{1}{2} \sum_i N_i \left(\frac{D_i}{S_{i,k}(\vec{\theta}_k)} + \log S_{i,k}(\vec{\theta}_k) \right), \quad (30)$$

where D_i is the PSD data value for the given frequency f_i , as computed by the method presented in [99, 100], using N_i averaged segments. The $S_{i,k}(\vec{\theta}_k)$ is the spline model of order k evaluated at f_i , that depends on a parameter set

$$\vec{\theta}_k = \{\log f_{1,k}, \dots, \log f_{k,k}, \log S_0, \dots, \log S_{k,k}, \log S_{k+1}\}, \quad (31)$$

in which the $\log S_0$ and $\log S_{k+1}$ parameters refer to the logarithm of the PSD amplitude of the two fixed knots at the “edges” of the spectrum. Those two parameters correspond to our zeroth model order ($k=0$), thus they are always being explored by the walkers of `Eryn`.

The results are shown in figure 4. In particular, in figure 4b we show the histogram of the recovered number of knots for the particular data-set. It is clear that 8 spline knots are preferred, two of them being fixed at the edges of the spectrum, and the other six knots free to take any position in the given frequency range. In figure 4c we show the 2D sliced posteriors for the spline parameters, $\{\log S_{j,k}\}$ and $\{\log f_{j,k}\}$. In this figure, we again stack all the MCMC samples across model orders.

The true spectrum is indicated by the orange solid line. There is an interesting outcome of this toy investigation; while there is a preferable dimensionality of the model, there is a weak constraint on the actual positions of the knots. We find that the sampler is virtually “scanning” the PSD data, showing slightly higher preference for locations between -6 and -4 in log-frequency, where the spectrum follows a more complicated shape. Finally, in figure 4a, the data (gray solid line and red data points), is shown together with model evaluations drawn from the posterior samples (pink solid lines).

IV. EXAMPLES FROM GRAVITATIONAL WAVE ASTRONOMY

In recent years, we have witnessed the beginning of Gravitational Wave Astronomy. Since the first detection [104] dozens of waveform signatures have been measured with the current network of observatories. At the time of the writing of this paper, more than 90 events have been recorded [7], the vast majority of them are black hole (BH) binary mergers, with a few of them being binary neutron star (NS) and BH-NS mergers. At the same time, detector networks are being improved [105, 106] and there are plans to expand them with the addition of new observatories, such as the Einstein Telescope [107, 108] or Cosmic Explorer [109, 110]. Those detectors will unlock the sky to larger redshifts z , allowing access to a vast number of potential sources. In addition, space missions, such as LISA [16, 20], are predicted to be signal-dominated observatories, with many types of sources populating their data streams. In fact, we expect that source confusion will be one of the primary challenges in future data analysis efforts in gravitational wave astronomy. In a typical data-set, we expect an unknown number of signals, originating from sources that generate waveforms with different characteristics. Those range from the stellar-mass BH binaries now frequently observed by ground-based detectors, to the supermassive BH binaries, extreme mass ratio inspirals, ultra compact Galactic binaries (UCB), and stochastic GW signals from both astrophysical and possibly cosmological origin [16, 20, 98]. For this final example, we will focus on the LISA mission, and in particular on the case of discriminating UCB signals.

A. Application to LISA data and the Ultra Compact Galactic Binaries

LISA is going to measure GW signals in the mHz regime, accessing sources of all the aforementioned types. As already discussed, the most numerous of them are going to be the UCBs, which will be almost monochromatic in the LISA band. Out of the millions of sources, only $\sim \mathcal{O}(10^4)$ will be individually resolvable, and the rest will generate a confusion signal. As a consequence, for the du-




FIG. 5: *Top panel:* The simulated data used for demonstrating the capabilities of **Eryn** in tackling a high-dimensional problem. A total of 10 Ultra Compact Binaries in the vicinity of our Galaxy emitting Gravitational Wave signals at the mHz range. Here we plot the power spectral density of the A data channel of LISA. The catalogue of sources is taken from the second LISA Data Challenge [111]. Each signal is represented by a different colour. *Bottom panel:* The same data-set, now comparing the injected signal against the solution yielded by **Eryn** (see text for more details). We have plotted the shaded area by sampling the joint posterior on model order k and the corresponding parameters.

ration of the mission, we will need to disentangle tens of thousands of sources which will be overlapping in both time and frequency domains. This is no trivial task, but various different strategies have already been proposed for analyzing such challenging data-sets. For example, Gaussian Processes can be employed [24], or Swarm Optimization techniques [25], or hybrid swarm-based algorithms [112]. Pipelines based on MCMC methods have been tested extensively [21, 23, 27, 113], and have been demonstrated to be able to tackle complex cases where signals are overlapping.

Here, we will focus on the same problem, employing **Eryn** to solve a down-scaled version of the UCB challenge. It is down-scaled because we focus only on a single narrow frequency band, containing several overlapping signals. In addition, we focus solely on demonstrating the capabilities of **Eryn** on dynamic parameter estimation for UCB type sources and no other types of signals are contained in the data (e.g., chirping signatures from supermassive BH binaries). At the same time, we have access to the level of instrumental noise, which is shown in both panels of figure 5. Searching for the UCB signals across the complete LISA band requires a more elaborate imple-

mentation of this simplified pipeline. This pipeline will be focusing on solving the complete second LISA Data Challenge [111], and is going to be presented in future work [31]. We choose to work on the frequency seg-

#	f_{gw} [mHz]	ρ_{opt}
1	3.99780	9.98
2	3.99781	46.70
3	3.99784	4.55
4	3.99854	39.45
5	3.99873	13.02
6	3.99882	8.47
7	3.99919	10.88
8	3.99939	19.07
9	3.99964	20.00
10	3.99965	7.99

TABLE I: The optimal SNR ρ_{opt} for each of the 10 injected sources, computed for the given duration of the mission (see eq. (37)). The dominant emission frequency f_{gw} is also given for reference.

ment between 3.997 and 4 mHz, which contains 10 UCB objects, drawn directly from the LDC2 catalogue [111]. Those are shown in the top panel of figure 5 which shows the power spectrum of the A data channel of LISA. We use the two noise-orthogonal A , E , and T Time Delay Interferometry variables [114–116], which are linear combinations of the LISA relative frequency TDI Michelson measurements X , Y , and Z as:

$$\begin{aligned} A &= \frac{1}{\sqrt{2}}(Z - X), & E &= \frac{1}{\sqrt{6}}(X - 2Y + Z), \\ T &= \frac{1}{\sqrt{3}}(X + Y + Z). \end{aligned} \quad (32)$$

In ideal conditions (equal noises across spacecrafts, and equal LISA arms), the noise in A and E is independent, while the T data stream can be used as a signal-insensitive *null* channel, useful for instrument noise calibration. Since we perform analysis on a noise-free injection, we will be neglecting the T channel altogether. We simulated the injection data for an observation time of $T_{\text{obs}} = 1$ year.

The *optimal* SNR for each injected source, ρ_{opt} , is given in table I. The ρ_{opt} quantity refers to the SNR of each source in *isolation*, with respect only to the instrumental noise, and can be calculated as

$$\rho_{\text{opt}}^2 = \sum_C (h_C | h_C)_C, \quad (33)$$

with $C \in \{A, E\}$ the noise-orthogonal TDI channels of eq. (32), while the $(\cdot | \cdot)_C$ notation represents the noise weighted inner product expressed for two time series a

and b as

$$(a|b)_C = 2 \int_0^\infty df \left[\tilde{a}^*(f) \tilde{b}(f) + \tilde{a}(f) \tilde{b}^*(f) \right] / \tilde{S}_{n,C}(f). \quad (34)$$

The tilde represents the data in the Fourier frequency domain, and the asterisk indicates complex conjugate. The $\tilde{S}_{n,C}(f)$ is the one-sided PSD of the noise in TDI channel C . Under our assumptions $S_{n,A}(f) = S_{n,E}(f)$.

For our investigation we chose to analyze noiseless data (no noise realization), while in the likelihood we are using the PSD noise levels taken from the LISA design studies [117]. For the signals, we utilize the fast frequency-domain UCB waveform model of [14]. Then, the two polarizations of an emitting UCB can be written as

$$\begin{aligned} h_+(t) &= \frac{2\mathcal{M}}{D_L} (\pi f_{\text{gw}}(t))^{2/3} (1 + \cos^2 \iota) \cos \phi(t), \\ h_\times(t) &= -\frac{4\mathcal{M}}{D_L} (\pi f_{\text{gw}}(t))^{2/3} \cos \iota \sin \phi(t), \end{aligned} \quad (35)$$

where \mathcal{M} is the chirp mass, f_{gw} is the instantaneous gravitational wave frequency, D_L is the luminosity distance, ι is the inclination of the binary orbit, and $\phi(t)$ is the gravitational wave phase. The phase ϕ can be expressed as $\phi = \phi_0 + 2\pi \int^t f_{\text{gw}}(t') dt'$, with ϕ_0 being an initial arbitrary phase shift. For more details about the waveform model, we refer the reader to [14, 80, 118].

In our simplified scenario, each binary signal in the Solar System Barycenter is determined by a set of eight parameters. Those are the $\vec{\theta} = \{\mathcal{A}, f_{\text{gw}} [\text{mHz}], \dot{f}_{\text{gw}} [\text{Hz/s}], \phi_0, \cos \iota, \psi, \lambda, \sin \beta\}$, where \mathcal{A} is the overall amplitude, \dot{f}_{gw} is the first derivative of the gravitational-wave frequency, ψ the polarization, λ is the ecliptic longitude, and β the ecliptic latitude of the binary. The amplitude of the signal is calculated as

$$\mathcal{A} = \left(2\mathcal{M}^{5/3} \pi^{2/3} f_{\text{gw}}^{2/3} \right) / D_L, \quad (36)$$

which can be used to obtain a rough SNR estimate, via [21]

$$\rho^2 = \frac{\mathcal{A}^2 T_{\text{obs}} \sin^2(f_{\text{gw}}/f_*)}{4S_n(f_{\text{gw}})}, \quad (37)$$

with $S_n(f_{\text{gw}})$ being the instrumental noise power spectral density at frequency f_{gw} , and $f_* = 1/(2\pi L)$, where L the LISA arm length. Given eq. (36) and (37), we find it convenient to directly sample on ρ instead of \mathcal{A} , which also yields a more illustrative measure of the amplitude of each binary. We use wide uniform priors for all the rest of the binary parameters, covering essentially the complete parameter space. The exception is again the amplitude (SNR), ρ , where we adopt a prior which was first introduced in [119] and then adapted in [21]. The prior density can be expressed as

$$p(\rho) = \frac{3\rho}{4\rho_* (1 + \rho/(4\rho_*))}, \quad (38)$$

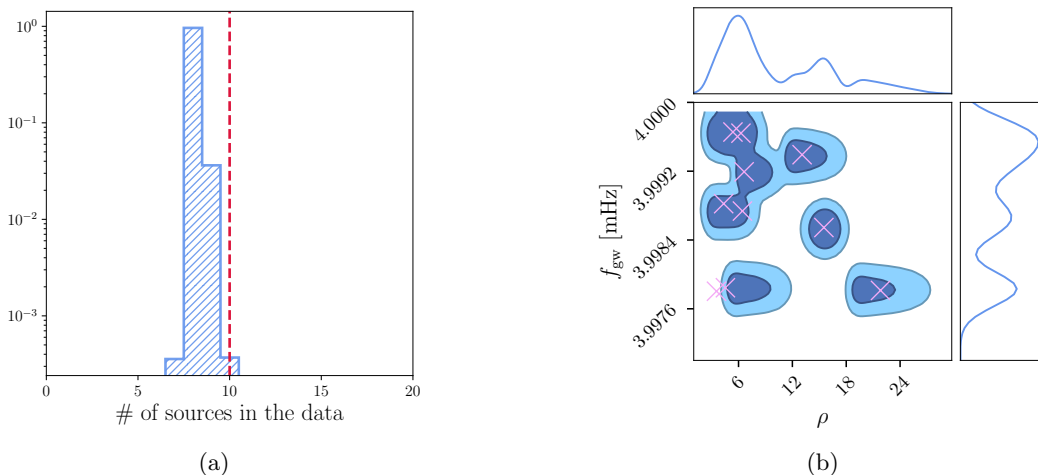


FIG. 6: *Left panel:* In this figure, we show the posterior on the number of UCB sources in the data. The true injected number is shown with the red dashed line. It is clear that, for the given measurement duration of the particular data set, we manage to confidently resolve eight binaries out of a total of ten. *Right panel:* Corner plot for two of the eight parameters characterising each UCB source. These are the amplitude, expressed as an SNR ρ , and the dominant emission (or initial) frequency, f_{gw} [mHz] (see text for more details). The violet crosses represent the injected parameter values. A corner plot for more parameters is shown in figure A.7 in the Appendix.

where ρ_* is a given constant that specifies the peak of the above distribution. This distribution is designed to prevent the proposal of sources with very small SNR in the model, as it drops sharply for $\rho \rightarrow 0$. Those weak sources do not significantly affect the likelihood, and so their inclusion must be penalised by the prior [120]. This prior choice forces the sampler to explore only potential sources with non-zero SNR, avoiding populating the chains with numerous undetectable signals. This prior performs adequately in this problem, but there are other solutions one could adopt in order to keep control of the number of very weak sources. This discussion, which sets the grounds for a *global-fit* analysis pipeline for the LISA data [21], is out of the scope of this paper, but a more detailed description will be presented in a future work [31].

a. Search Phase: Before parameter estimation, we initiate a *search phase* of our analysis, with the aim of getting the walkers to a better starting point on the posterior surface. This phase consists of an iterative brute force procedure, based on drawing a very large number of proposals, then maximizing the likelihood over the initial phase ϕ_0 , and finally perform a rapid MCMC sampling over the parameter space, using only a one-source model (therefore there is no dynamical parameter spaces). In particular, we draw 5×10^6 points in the parameter space, and after phase maximization, we use them as starting points to a parallel tempered MCMC run with $N_T = 10$ temperatures, each running with $n_w = 500$ walkers. When this step concludes, we keep the 100 best samples in terms of likelihood value and use their corresponding parameter estimates as starting points for the parameter estimation portion of the analysis. We also use the maximum Likelihood solution to subtract the

source found from the data. We then use the residual data to search for another source, and this process repeats until there is no signal found with SNR $\rho > 5$. In between successive iterations of the single-binary search, we run another MCMC over all sources found so far in order to adjust the parameters to account for correlations and overlap between sources. After convergence, we found eight sources in our data-set with an optimal SNR greater than 5. We take these found sources and add them to all walkers in the sampler at the beginning of the full MCMC run described below.

b. Parameter estimation: During this step, we perform hybrid MCMC sampling, where we both update the found sources (in-model) and dynamically search for new and weaker sources in the data employing Reversible Jump sampling. For the number of signals k , we adopt a uniform prior $k \sim \mathcal{U}[6, 20]$. For the sake of convenience in this simple application, we keep the six loudest binaries found during the search phase as *fixed*. This means that we still sample their waveform parameters, but they are not allowed to be removed by the Reversible Jump process. We chose this setup in order to accelerate the convergence of the algorithm, being confident that these sources are part of the solution. In future work, this will be adjusted to deal with the much larger complexity of the full problem.

Concerning the sampler settings, we use the adaptive parallel tempering scheme of section IID, building a temperature ladder of $N_T = 10$ temperatures, with 100 walkers for each temperature. For this run, we have also utilized the Multiple Try Metropolis algorithm (see section IIC) in order to improve the acceptance rate in the reversible jump proposal. We have also tried the Delayed-

Rejection scheme which is implemented for **Eryn**, but we found that the Multiple-Try strategy yields more efficient sampling. Finally, we have utilized the basic stretch and group stretch proposals that were described in section II A.

After convergence, the result is shown in figures 5 and 6. In figure 6a, the sampled posterior on the number of sources k is presented. In this histogram, we have added the six fixed binaries to the actual number of signals being sampled via the Reversible Jump algorithm. It is fairly obvious that we have managed to confidently resolve eight out of the ten injected binary signals. This fact that we do not favour 10 sources can be explained partly by the low SNR of the signals (see table I) and partly by confusion from source overlap (also shown in figure 5). Additionally, the result of figure 6a depends on the given observation duration. The greater the T_{obs} , the better our ability to resolve the confused sources. Thus, in that case, we should expect more Reversible Jump iterations across the higher dimensional models.

On the right panel, in figure 6b, the ensemble 2D posterior slice is shown, for two selected parameters. We call it ensemble because we are again “stacking” all the chains for these two parameters for all sources for all model orders k . We chose to show only the amplitude (the ρ parameter explained in eq. (37) above) and the dominant emission frequency f_{gw} , which illustrates the number of sources resolved, and how they overlap in frequency. A corner plot for more parameters is shown in figure A.7 in Appendix A. We also show the true injection values, marked as crosses, on top of the 2D posterior. From this plot alone, one can see that the sampler is exploring efficiently the parameter space, converging to the true values of the resolvable binaries that were injected.

V. DISCUSSION

We have implemented **Eryn**, a Bayesian sampling package capable of performing efficient trans-dimensional inference, by employing different techniques that improve its acceptance rate. These techniques are the affine invariant sampling, the adaptive parallel tempering, the delayed rejection, and multiple try metropolis, in combination with the construction of informative proposal distributions for the parameters of the models. The structure of **Eryn** is based on the widely used software `emcee` [58], enhanced with the ability of performing Reversible Jumps [30] between different model spaces. The sampler capabilities have been demonstrated with toy models that are commonly encountered in different data analysis problems. We have begun with an application to signal detection, and in particular to searching for simple signals in the form of Gaussian pulse signals in the presence of Gaussian noise (see section III C 1).

In section III C 2, we applied our algorithm to a problem of modeling power spectra with arbitrary shapes in frequency domain. In such cases, it is convenient to de-

fine models based on B-splines, which are able to faithfully capture the shape of any spectral data. However, in order to avoid over-fitting situations, the optimal order of the model (*i.e.* the optimal number of spline knots), needs to be estimated from the data. This can be done either sequentially, by trying models of different dimensionality and then comparing their performance, or dynamically, by using trans-dimensional algorithms such as **Eryn**. This class of problems is often encountered in cosmology [121, 122], where the signal of interest is stochastic in nature, and sometimes the prior knowledge on its shape is very limited. As already discussed, this is especially true for future GW observatories, which open the possibility of detecting such signals from both astrophysical and cosmological origin [16, 98, 102]. The different theoretical models produce spectra with distinct shapes, increasing the need for shape-agnostic spectral models, such as the B-spline used here.

Finally, in section IV A, we demonstrated **Eryn** in a more complicated problem, that of the analysis of ultra compact binary signals measured by the future LISA detector. These objects are going to produce the majority of the signals in the LISA data, each emitting almost monochromatic radiation. Their vast number will generate a confusion foreground, while only a few thousand of them will be resolvable from the data. We employ our tools described in this work, together with a search phase that is based on iteratively running the sampler on “static models” (no trans-dimensional moves) with phase-maximized likelihoods. We do these runs on the residuals of each iteration, with the aim of extracting all bright sources. In order to achieve faster convergence of our parameter estimation run, we choose to keep the brightest sources found during search as fixed (minimum number of model order is $k = 6$), while the the Reversible Jump algorithm is used to search for weaker signals in the data. This is purely a choice that allows quick convergence in this fully-controlled and simplified LISA dataset.

We perform this analysis for a mission duration of $T_{\text{obs}} = 1$ year and only on a single narrow frequency band around 4 mHz, which contains a total of 10 binary signals. It is worth noting here, that the synthetic data were produced assuming idealized conditions. This means that we do not consider any data irregularities, such as data gaps and glitches and spectral lines, or any other contamination originating from the mixing of signals of different types (such as supermassive BH binaries). In the end, as shown in figure 6a, we manage to recover eight out of ten injected signals. This result makes sense given the relative strength of the injections, and their waveform overlap. Many of the injected sources have an optimal SNR in isolation which is rather low (see table I), so these are more susceptible to deterioration when we account for signal overlap.

The above investigations demonstrate that the dynamical parameter estimation capabilities of **Eryn** are suitable for these types of computationally demanding problems.

Eryn has already been used in several works that have been already published [80, 102, 123], or are going to appear soon. The work presented in this paper is the initial part of our efforts toward implementing a data analysis pipeline for LISA data. This pipeline will be demonstrated on the LDC2 data-set [111], which contains multiple types of signals overlapping in both time and frequency domains. That being said, Eryn is a generic and versatile sampler, which can be used in any investigation that requires Reversible Jump sampling, and to our knowledge is one of the very few statistical tools of this kind that is not specialized to a single type of analysis (see discussion in section II F).

ACKNOWLEDGMENTS

We wish to thank S. Babak, M. Le Jeune, S. Marsat, T. Littenberg, and N. Cornish for their useful comments and very fruitful discussions. N Stergioulas and N Karnezi acknowledge support from the Gr-PRODEX 2019 funding program (PEA 4000132310).

Appendix A:

In figure A.7 we show the triangle plot of the stacked posterior points as sampled by Eryn, for the investigation of section IV. The difference to figure 6b is that here we include more parameters of the sources, but we still do not include all parameters for the sake of clarity.

-
- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [2] W. K. Hastings, Monte Carlo Sampling Methods Using Markov Chains and Their Applications, *Biometrika* **57**, 97 (1970).
- [3] D. B. Hitchcock, A history of the metropolis-hastings algorithm, *The American Statistician* **57**, 254 (2003), <https://doi.org/10.1198/0003130032413>.
- [4] W. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice*, Chapman & Hall/CRC Interdisciplinary Statistics (Taylor & Francis, 1995).
- [5] G. Ashton and C. Talbot, Bilby-MCMC: an MCMC sampler for gravitational-wave inference, *Mon. Not. Roy. Astron. Soc.* **507**, 2037 (2021), [arXiv:2106.08730 \[gr-qc\]](https://arxiv.org/abs/2106.08730).
- [6] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019), [arXiv:1807.10312 \[astro-ph.IM\]](https://arxiv.org/abs/1807.10312).
- [7] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run, (2021), [arXiv:2111.03606 \[gr-qc\]](https://arxiv.org/abs/2111.03606).
- [8] T. L. S. Collaboration and the Virgo Collaboration, A guide to ligo-virgo detector noise and extraction of transient gravitational-wave signals, *Classical and Quantum Gravity* **37**, 055002 (2020).
- [9] P. Shawhan, Ligo data analysis, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **502**, 396 (2003), proceedings of the VIII International Workshop on Advanced Computing and Analysis Techniques in Physics Research.
- [10] R. Abbott *et al.* (LIGO Scientific, Virgo), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021), [arXiv:1912.11716 \[gr-qc\]](https://arxiv.org/abs/1912.11716).
- [11] G. Ashton and R. Prix, Hierarchical multistage mcmc follow-up of continuous gravitational wave candidates, *Phys. Rev. D* **97**, 103020 (2018).
- [12] E. Thrane and C. Talbot, An introduction to bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models, *Publications of the Astronomical Society of Australia* **36**, e010 (2019).
- [13] M. Isi, W. M. Farr, and K. Chatziioannou, Comparing Bayes factors and hierarchical inference for testing general relativity with gravitational waves, *Phys. Rev. D* **106**, 024048 (2022), [arXiv:2204.10742 \[gr-qc\]](https://arxiv.org/abs/2204.10742).
- [14] N. J. Cornish and T. B. Littenberg, Tests of bayesian model selection techniques for gravitational wave astronomy, *Phys. Rev. D* **76**, 083006 (2007).
- [15] A. Gelman, G. O. Roberts, and W. R. Gilks, Efficient metropolis jumping rules, in *Bayesian Statistics*, edited by J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Oxford University Press, Oxford, 1996) pp. 599–608.
- [16] P. Amaro-Seoane *et al.*, Laser Interferometer Space Antenna, *arXiv e-prints*, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786) (2017), [arXiv:1702.00786 \[astro-ph.IM\]](https://arxiv.org/abs/1702.00786).
- [17] J. Luo, L.-S. Chen, H.-Z. Duan, Y.-G. Gong, S. Hu, J. Ji, Q. Liu, J. Mei, V. Milyukov, M. Sazhin, C.-G. Shao, V. T. Toth, H.-B. Tu, Y. Wang, Y. Wang, H.-C. Yeh, M.-S. Zhan, Y. Zhang, V. Zharov, and Z.-B. Zhou, Tianqin: a space-borne gravitational wave detector, *Classical and Quantum Gravity* **33**, 035010 (2016).
- [18] S. Kawamura *et al.*, The japanese space gravitational wave antenna - decigo, *Journal of Physics: Conference Series* **122**, 012006 (2008).
- [19] Z. Ren *et al.*, Taiji Data Challenge for Exploring Gravitational Wave Universe, (2023), [arXiv:2301.02967 \[gr-qc\]](https://arxiv.org/abs/2301.02967).
- [20] P. Amaro-Seoane *et al.*, Low-frequency gravitational-wave science with eLISA/NGO, *Classical and Quantum Gravity* **29**, 124016 (2012), [arXiv:1202.0839 \[gr-qc\]](https://arxiv.org/abs/1202.0839).
- [21] T. B. Littenberg, N. J. Cornish, K. Lackeos, and T. Robson, Global analysis of the gravitational wave signal from galactic binaries, *Phys. Rev. D* **101**, 123021 (2020).
- [22] T. Littenberg, N. Cornish, K. Lackeos, and T. Robson, Global Analysis of the Gravitational Wave Signal from Galactic Binaries, *Phys. Rev. D* **101**, 123021 (2020),

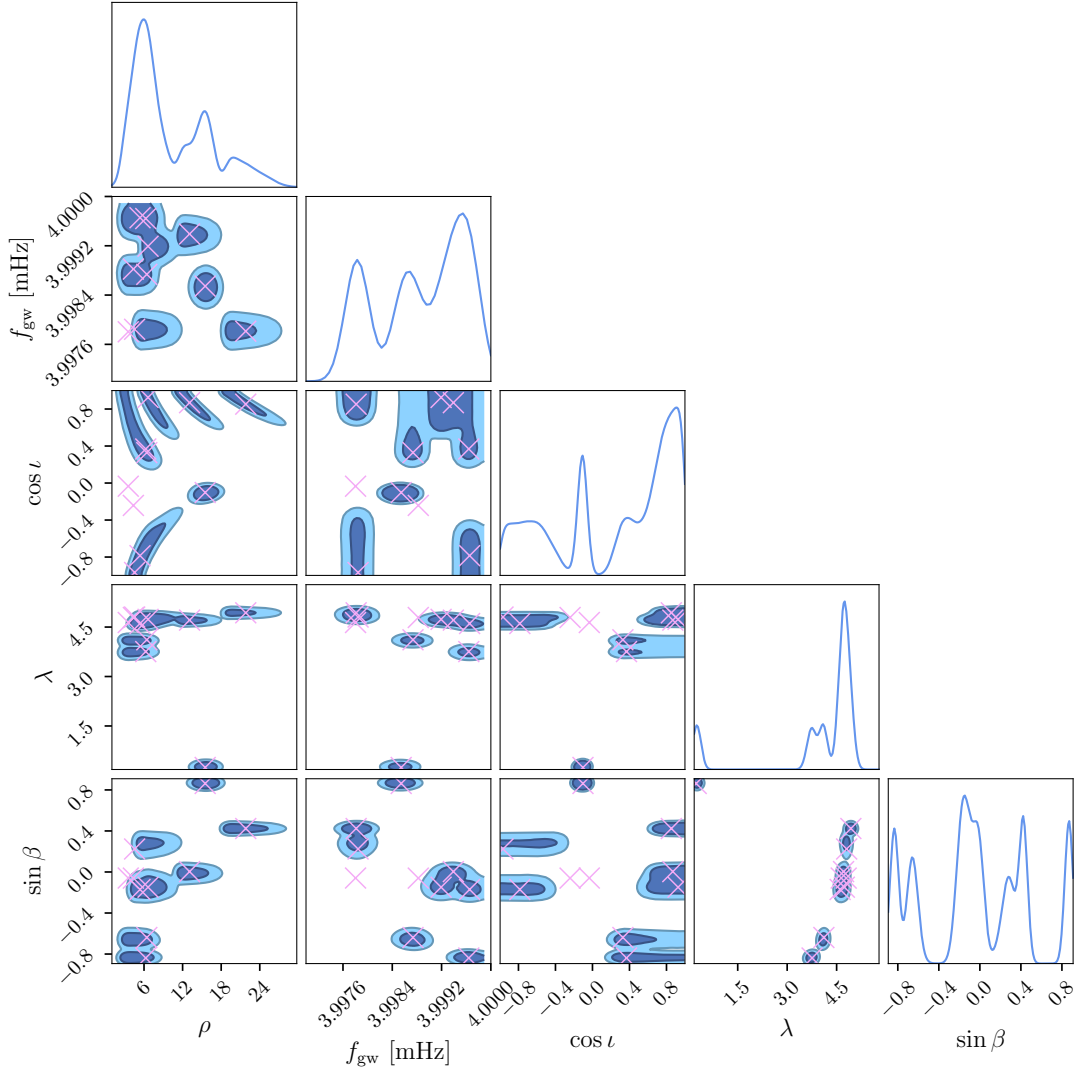


FIG. A.7: A triangle plot showing the 2D posterior slices for the application of section IV A, but for a greater selection of parameters than figure 6b. The rest of the parameters, if plotted stacked in the same manner, result in surfaces that cannot be so easily interpreted, and therefore have been left out. The true injected values are marked with crosses.

- [arXiv:2004.08464 \[gr-qc\]](#).
- [23] T. B. Littenberg and N. J. Cornish, Prototype Global Analysis of LISA Data with Multiple Source Types, (2023), [arXiv:2301.03673 \[gr-qc\]](#).
- [24] S. H. Strub, L. Ferraioli, C. Schmelzbach, S. C. Stähler, and D. Giardini, Bayesian parameter-estimation of Galactic binaries in LISA data with Gaussian Process Regression, (2022), [arXiv:2204.04467 \[astro-ph.IM\]](#).
- [25] X.-H. Zhang *et al.*, Resolving Galactic binaries in LISA data using particle swarm optimization and cross-validation, *Phys. Rev. D* **104**, 024023 (2021), [arXiv:2103.09391 \[gr-qc\]](#).
- [26] N. Karnesis, S. Babak, M. Pieroni, N. Cornish, and T. Littenberg, Characterization of the stochastic signal originating from compact binary populations as measured by LISA, *Phys. Rev. D* **104**, 043019 (2021), [arXiv:2103.14598 \[astro-ph.IM\]](#).
- [27] J. Crowder and N. J. Cornish, Solution to the galactic foreground problem for lisa, *Phys. Rev. D* **75**, 043008 (2007).
- [28] K. Breivik, S. Coughlin, M. Zevin, C. L. Rodriguez, K. Kremer, C. S. Ye, J. J. Andrews, M. Kurkowski, M. C. Digman, S. L. Larson, and F. A. Rasio, COSMIC Variance in Binary Population Synthesis, *ApJ* **898**, 71 (2020), [arXiv:1911.00903 \[astro-ph.HE\]](#).
- [29] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart, The gravitational wave signal from the Galactic disk population of binaries containing two compact objects, *A&A* **375**, 890 (2001), [arXiv:astro-ph/0105221 \[astro-ph\]](#).
- [30] P. J. Green, Reversible jump Markov chain Monte Carlo computation and Bayesian model determination, *Biometrika* **82**, 711 (1995).
- [31] M. L. Katz et al., in prep, .

- [32] M. P. Hobson, A. H. Jaffe, A. R. Liddle, P. Mukherjee, and D. Parkinson, *Bayesian Methods in Cosmology* (2009).
- [33] D. W. Hogg and D. Foreman-Mackey, Data analysis recipes: Using markov chain monte carlo, *The Astrophysical Journal Supplement Series* **236**, 11 (2018).
- [34] S. Sharma, Markov Chain Monte Carlo Methods for Bayesian Data Analysis in Astronomy, *ARA&A* **55**, 213 (2017), [arXiv:1706.01629](https://arxiv.org/abs/1706.01629) [astro-ph.IM].
- [35] M. Johannes and N. Polson, *MCMC Methods for Financial Econometrics* (2009).
- [36] H. Haario, E. Saksman, and J. Tamminen, An adaptive metropolis algorithm, *Bernoulli* **7**, 223 (2001).
- [37] D. Frenkel, Waste-recycling monte carlo, in *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, edited by M. Ferrario, G. Ciccotti, and K. Binder (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006) pp. 127–137.
- [38] C. Andrieu and J. Thoms, A tutorial on adaptive mcmc, *Statistics and Computing* **18**, issue 4, 343 (2008), publisher: Springer Netherlands.
- [39] M. Vallisneri, Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects, *Phys. Rev. D* **77**, 042001 (2008), [arXiv:gr-qc/0703086](https://arxiv.org/abs/gr-qc/0703086) [gr-qc].
- [40] R. Neal, MCMC Using Hamiltonian Dynamics, in *Handbook of Markov Chain Monte Carlo* (2011) pp. 113–162.
- [41] M. J. Betancourt and M. Girolami, Hamiltonian Monte Carlo for Hierarchical Models, [arXiv e-prints](https://arxiv.org/abs/1312.0906), [arXiv:1312.0906](https://arxiv.org/abs/1312.0906) (2013), [arXiv:1312.0906](https://arxiv.org/abs/1312.0906) [stat.ME].
- [42] M. D. Hoffman and A. Gelman, The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo, [arXiv e-prints](https://arxiv.org/abs/1111.4246), [arXiv:1111.4246](https://arxiv.org/abs/1111.4246) (2011), [arXiv:1111.4246](https://arxiv.org/abs/1111.4246) [stat.CO].
- [43] C. Ritter and M. Tanner, Facilitating the Gibbs Sampler: The Gibbs Stopper and the Griddy-Gibbs Sampler *Journal of the American Statistical Association* **87**, 861 (1992).
- [44] P. Muller, *A generic approach to posterior integration and Gibbs sampling*, Tech. Rep. 91-09 (Purdue University, 1991).
- [45] W. R. Gilks and P. Wild, Adaptive rejection sampling for gibbs sampling, *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **41**, 337 (1992).
- [46] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*, 2nd ed. (Chapman and Hall/CRC, 2004).
- [47] S. P. Brooks, Markov chain monte carlo method and its application, *Journal of the Royal Statistical Society. Series D (The Statistician)* **47**, 69 (1998).
- [48] C. Robert and G. Casella, A short history of markov chain monte carlo: Subjective recollections from incomplete data, *Statistical Science* **26**, 102 (2011).
- [49] A. Joseph, Markov Chain Monte Carlo Methods in Quantum Field Theories: A Modern Primer (Springer, 2019) [arXiv:1912.10997](https://arxiv.org/abs/1912.10997) [hep-th].
- [50] S. Sharma, Markov chain monte carlo methods for bayesian data analysis in astronomy, *Annual Review of Astronomy and Astrophysics* **55**, 213 (2017), <https://doi.org/10.1146/annurev-astro-082214-122339>.
- [51] W. Kendall, F. Liang, and J. Wang, *Markov Chain Monte Carlo: Innovations and Applications*, Institute for Mathematical Sciences lecture notes series (World Scientific, 2005).
- [52] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, Oxford science publications (OUP Oxford, 2006).
- [53] D. Sorensen and D. Gianola, *Likelihood, Bayesian, and MCMC Methods in Quantitative Genetics*, Statistics for Biology and Health (Springer New York, 2007).
- [54] M. Johannes and N. Polson, Chapter 13 - mcmc methods for continuous-time financial econometrics, in *Handbook of Financial Econometrics: Applications*, Handbooks in Finance, Vol. 2, edited by Y. AÏT-SAHALIA and L. P. HANSEN (Elsevier, San Diego, 2010) pp. 1–72.
- [55] G. Baio, *Bayesian Methods in Health Economics*, Chapman & Hall/CRC Biostatistics Series (CRC Press, 2012).
- [56] J. Hermans, V. Begy, and G. Louppe, Likelihood-free mcmc with amortized approximate ratio estimators (2019).
- [57] H. Wu, J. Köhler, and F. Noé, Stochastic normalizing flows (2020).
- [58] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC Hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013), [arXiv:1202.3665](https://arxiv.org/abs/1202.3665) [astro-ph.IM].
- [59] J. Goodman and J. Weare, Ensemble samplers with affine invariance, *Communications in Applied Mathematics and Computational Science* **5**, 65 (2010).
- [60] P. Virtanen *et al.*, SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nature Methods* **17**, 261 (2020), [arXiv:1907.10121](https://arxiv.org/abs/1907.10121) [cs.MS].
- [61] K. C. Wong *et al.*, H0LiCOW - XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes, *MNRAS* **498**, 1420 (2020), [arXiv:1907.04869](https://arxiv.org/abs/1907.04869) [astro-ph.CO].
- [62] P. J. McMillan, The mass distribution and gravitational potential of the Milky Way, *MNRAS* **465**, 76 (2017), [arXiv:1608.00971](https://arxiv.org/abs/1608.00971) [astro-ph.GA].
- [63] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, Tidal Deformabilities and Radii of Neutron Stars from the Observation of GW170817, *Phys. Rev. Lett.* **121**, 091102 (2018), [arXiv:1804.08583](https://arxiv.org/abs/1804.08583) [astro-ph.HE].
- [64] M. Trias, A. Vecchio, and J. Veitch, Delayed rejection schemes for efficient Markov-Chain Monte-Carlo sampling of multimodal distributions, (2009), [arXiv:0904.2207](https://arxiv.org/abs/0904.2207) [stat.ME].
- [65] A. Mira, On metropolis-hastings algorithms with delayed rejection, *Metron - International Journal of Statistics* **LIX**, 231 (2001).
- [66] P. J. Green and A. Mira, Delayed rejection in reversible jump metropolis-hastings, *Biometrika* **88**, 1035 (2001).
- [67] H. Haario, M. Laine, A. Mira, and E. Saksman, Dram: Efficient adaptive mcmc, *Stat Comp*, 339.
- [68] N. J. Cornish, Heterodyned likelihood for rapid gravitational wave parameter inference, *Phys. Rev. D* **104**, 104054 (2021), [arXiv:2109.02728](https://arxiv.org/abs/2109.02728) [gr-qc].
- [69] L. Martino, A review of multiple try mcmc algorithms for signal processing, *Digital Signal Processing* **75**, 134 (2018).
- [70] J. S. Liu, F. Liang, and W. H. Wong, The multiple-try method and local optimization in metropolis sampling, *Journal of the American Statistical Association* **95**, 121 (2000).
- [71] L. Martino, V. P. Del Olmo, and J. Read, A multi-point Metropolis scheme with generic weight functions, [arXiv:1808.08888](https://arxiv.org/abs/1808.08888).

- e-prints , arXiv:1112.4048 (2011), [arXiv:1112.4048 \[stat.CO\]](#).
- [72] M. Bédard, R. Douc, and E. Moulines, Scaling analysis of multiple-try mcmc methods, *Stochastic Processes and their Applications* **122**, 758 (2012).
- [73] R. H. Swendsen and J.-S. Wang, Replica monte carlo simulation of spin-glasses, *Phys. Rev. Lett.* **57**, 2607 (1986).
- [74] K. Hukushima and K. Nemoto, Exchange monte carlo method and application to spin glass simulations, *Journal of the Physical Society of Japan* **65**, 1604 (1996), <https://doi.org/10.1143/JPSJ.65.1604>.
- [75] W. Vousden, W. Farr, and I. Mandel, Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations, *MNRAS* **455**, 1919 (2016), [arXiv:1501.05823 \[astro-ph.IM\]](#).
- [76] R. E. Kass and A. E. Raftery, Bayes factors, *Journal of the American Statistical Association* **90**, 773 (1995).
- [77] N. J. Cornish and T. B. Littenberg, Tests of bayesian model selection techniques for gravitational wave astronomy, *Phys. Rev. D* **76**, 083006 (2007).
- [78] N. Lartillot and H. Philippe, Computing Bayes Factors Using Thermodynamic Integration, *Systematic Biology* **55**, 195 (2006), <https://academic.oup.com/sysbio/article-pdf/55/2/195/26557316/10635150500433722.pdf>.
- [79] P. M. Goggans and Y. Chi, Using Thermodynamic Integration to Calculate the Posterior Probability in Bayesian Model Selection Problems, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, American Institute of Physics Conference Series, Vol. 707 (2004) pp. 59–66.
- [80] M. L. Katz, C. Danielski, N. Karnesis, V. Korol, N. Tamanini, N. J. Cornish, and T. B. Littenberg, Bayesian Characterisation of Circumbinary Exoplanets with LISA, arXiv e-prints , arXiv:2205.03461 (2022), [arXiv:2205.03461 \[astro-ph.EP\]](#).
- [81] P. Maturana-Russel, R. Meyer, J. Veitch, and N. Christensen, Stepping-stone sampling algorithm for calculating the evidence of gravitational wave models, *Phys. Rev. D* **99**, 084006 (2019).
- [82] T. B. Littenberg and N. J. Cornish, A Bayesian Approach to the Detection Problem in Gravitational Wave Astronomy, *Phys. Rev. D* **80**, 063007 (2009), [arXiv:0902.0368 \[gr-qc\]](#).
- [83] T. B. Littenberg and N. J. Cornish, Separating Gravitational Wave Signals from Instrument Artifacts, *Phys. Rev. D* **82**, 103007 (2010), [arXiv:1008.1577 \[gr-qc\]](#).
- [84] S. J. Godsill, On the Relationship Between Markov Chain Monte Carlo Methods for Model Uncertainty, *Journal of Computational and Graphical Statistics* **10**, 230 (2001).
- [85] P. Dellaportas, J. J. Forster, and I. Ntzoufras, On Bayesian model and variable selection using MCMC, *Statistics and Computing* , 27 (2002).
- [86] H. F. Lopes and M. West, Bayesian model assessment in factor analysis, *Statistica Sinica* **14**, 41 (2004).
- [87] N. Karnesis *et al.*, Bayesian model selection for lisa pathfinder, *Phys. Rev. D* **89**, 062001 (2014).
- [88] A. Marrs, An application of reversible-jump mcmc to multivariate spherical gaussian mixtures, in *Advances in Neural Information Processing Systems*, Vol. 10, edited by M. Jordan, M. Kearns, and S. Solla (MIT Press, 1997).
- [89] Z. Khan, T. Balch, and F. Dellaert, Mcmc-based particle filtering for tracking a variable number of interacting targets, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**, 1805 (2005).
- [90] X. Yu, X. Wang, C. Lu, and Y. Wang, An adaptive reversible jump MCMC inversion algorithm for airborne time-domain electromagnetic data, *IOP Conference Series: Earth and Environmental Science* **660**, 012124 (2021).
- [91] Eryn Vorn (Sindarin for Blackwood) was a wooded cape in Eriador, and a region of dark pine trees. Located in western Minhiriath, Eryn Vorn (likely named so by the Númenóreans) was originally part of the vast ancient treescape that covered most of north-western Middle-earth.
- [92] Tolkien Gateway, http://tolkiengateway.net/wiki/Eryn_Vorn (), accessed: 2022-06-30.
- [93] P. Amaro-Seoane *et al.*, Astrophysics with the Laser Interferometer Space Antenna, (2022), [arXiv:2203.06016 \[gr-qc\]](#).
- [94] L. Dinh, J. Sohl-Dickstein, and S. Bengio, *Density estimation using real nvp* (2016).
- [95] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, *Neural spline flows* (2019).
- [96] Eryn Documentation , mikekatz04.github.io/Eryn/ (), accessed: 2023-01-30.
- [97] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), Upper limits on the isotropic gravitational-wave background from advanced ligo and advanced virgo’s third observing run, *Phys. Rev. D* **104**, 022004 (2021).
- [98] P. Auclair *et al.* (LISA Cosmology Working Group), Cosmology with the Laser Interferometer Space Antenna, (2022), [arXiv:2204.05434 \[astro-ph.CO\]](#).
- [99] M. Tröbs and G. Heinzel, Improved spectrum estimation from digitized time series on a logarithmic frequency axis, *Measurement* **39**, 120 (2006).
- [100] M. Armano *et al.*, Beyond the required lisa free-fall performance: New lisa pathfinder results down to 20 μ Hz, *Phys. Rev. Lett.* **120**, 061101 (2018).
- [101] L. Piegl and W. Tiller, *The NURBS Book*, 2nd ed. (Springer-Verlag, New York, NY, USA, 1996).
- [102] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, Uncovering gravitational-wave backgrounds from noises of unknown shape with LISA, (2023), [arXiv:2302.12573 \[gr-qc\]](#).
- [103] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, *et al.*, SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nature Methods* **17**, 261 (2020), publisher: Nature Research.
- [104] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [105] The LIGO Scientific Collaboration, *Gravitational wave astronomy with ligo and similar detectors in the next decade* (2019).
- [106] B. P. Abbott *et al.*, Prospects for observing and localizing gravitational-wave transients with advanced ligo, advanced virgo and kagra, *Living reviews in relativity* **23**, 1 (2020).
- [107] M. Maggiore *et al.*, Science Case for the Einstein Telescope, *JCAP* **03**, 050, [arXiv:1912.02622 \[astro-ph.CO\]](#).

- [108] M. Punturo *et al.*, The Einstein Telescope: a third-generation gravitational wave observatory, *Classical and Quantum Gravity* **27**, 194002 (2010).
- [109] M. Evans *et al.*, A Horizon Study for Cosmic Explorer: Science, Observatories, and Community, arXiv e-prints, arXiv:2109.09882 (2021), arXiv:2109.09882 [astro-ph.IM].
- [110] B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, K. Ackley, C. Adams, P. Addesso, R. X. Adhikari, V. Adya, C. Affeldt, *et al.*, Exploring the sensitivity of next generation gravitational wave detectors, *Classical and Quantum Gravity* **34**, 044001 (2017).
- [111] LISA Data Challenges Working Group, *LISA Data Challenges* (2022).
- [112] Y. Bouffanais and E. K. Porter, Detecting compact galactic binaries using a hybrid swarm-based algorithm, *Phys. Rev. D* **93**, 064020 (2016), arXiv:1509.08867 [gr-qc].
- [113] T. B. Littenberg, Detection pipeline for Galactic binaries in LISA data, *Phys. Rev. D* **84**, 063009 (2011), arXiv:1106.6355 [gr-qc].
- [114] M. Tinto and S. V. Dhurandhar, Time-delay interferometry, *Living Reviews in Relativity* **8**, 4 (2005).
- [115] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong, Lisa optimal sensitivity, *Phys. Rev. D* **66**, 122002 (2002).
- [116] Q. Baghi, J. I. Thorpe, J. Slutsky, and J. Baker, Statistical inference approach to time-delay interferometry for gravitational-wave detection, *Phys. Rev. D* **103**, 042006 (2021).
- [117] LISA Science Study Team, *LISA Science Requirements Document*, Tech. Rep. ESA-L3-EST-SCI-RS-001 (ESA, 2018).
- [118] T. Robson, N. J. Cornish, N. Tamanini, and S. Toonen, Detecting hierarchical stellar systems with lisa, *Phys. Rev. D* **98**, 064012 (2018).
- [119] N. J. Cornish and T. B. Littenberg, Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches, *Classical and Quantum Gravity* **32**, 135012 (2015), arXiv:1410.3835 [gr-qc].
- [120] We remind that the prior is also used as our “birth” proposal here.
- [121] Planck Collaboration, Aghanim, N., *et al.*, Planck 2018 results - i. overview and the cosmological legacy of planck, *A&A* **641**, A1 (2020).
- [122] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. V. CMB power spectra and likelihoods, *A&A* **641**, A5 (2020), arXiv:1907.12875 [astro-ph.CO].
- [123] M. L. Katz, Fully automated end-to-end pipeline for massive black hole binary signal extraction from LISA data, *Phys. Rev. D* **105**, 044055 (2022), arXiv:2111.01064 [gr-qc].