

Automatic tempered posterior distributions for Bayesian inversion problems

L. Martino[†], F. Llorente*, E. Curbelo*, J. López-Santiago*, J. Míguez*

[†] Universidad rey Juan Carlos (URJC), Madrid, Spain.

* Universidad Carlos III de Madrid (UC3M), Madrid, Spain.

Abstract

We propose a novel adaptive importance sampling scheme for Bayesian inversion problems where the inference of the variables of interest and the power of the data noise is split. More specifically, we consider a Bayesian analysis for the variables of interest (i.e., the parameters of the model to invert), whereas we employ a maximum likelihood approach for the estimation of the noise power. The whole technique is implemented by means of an iterative procedure, alternating sampling and optimization steps. Moreover, the noise power is also used as a tempered parameter for the posterior distribution of the variables of interest. Therefore, a sequence of tempered posterior densities is generated, where the tempered parameter is automatically selected according to the actual estimation of the noise power. A complete Bayesian study over the model parameters and the scale parameter can be also performed. Numerical experiments show the benefits of the proposed approach.

1 Introduction

The estimation of unknown parameters from noisy observations is an essential problem in signal processing, statistics and machine learning [1, 2, 3]. Within the Bayesian signal processing framework, these problems are addressed by constructing posterior probability distributions of the unknowns. Given the posterior, one often wants to make inference about the unknowns, e.g., if we are estimating parameters, finding the values that maximize their posterior, or the values that minimize some cost function given the uncertainty of the parameters. Unfortunately, obtaining closed-form solutions, usually expressed as integrals of the posterior, is infeasible in most practical applications. Hence, developing approximate computational techniques (such as importance sampling and MCMC algorithms) are often required [4, 5, 6].

The so-called *tempering of the posterior* is a well-known procedure for improving the performance of the Monte Carlo (MC) algorithms [7, 8, 9, 10]. The tempering is obtained by modulating an artificial

scale parameter or by sequentially including new data. The reasons of the improvement in the performance are several: improving mixing, discovering modes, foster the exploration of the inference space, etc. In the first iterations of the MC scheme, a posterior density with a bigger scale is considered. The artificial scale parameter (often called “temperature”) is reduced during the iterations, until considering the true posterior distribution. However, the user should decide a *temperature schedule*, i.e., a decreasing rule for the scale parameter, which is usually chosen in an heuristic way. In the literature, the tempering procedure has gained a particular attention for the estimation of the marginal likelihood (a.k.a., Bayesian model evidence) [9, 11, 12].

Furthermore, the joint inference of parameters (denoted as θ) of observation models, $\mathbf{f}(\theta)$, and scale parameters of the likelihood function (that, in the scalar case, is usually denoted as σ) can be an hard task. Indeed, “wrong choices” of σ values can easily jeopardize the sampling of θ . In this work, we introduce a procedure to tackle this problem.

More specifically, in this work, we design an adaptive importance sampling (AIS) scheme [13, 14] for Bayesian inversion problems, where an automatic tempering procedure is implemented. We consider that the vector of observations \mathbf{y} is obtained by a nonlinear transformation $\mathbf{f}(\theta)$ of the variables of interest θ , perturbed by additive Gaussian noise with unknown power σ^2 . The nonlinear mapping $\mathbf{f}(\theta)$ usually represents a complex physical model or a computer code etc. The resulting posterior densities are usually highly multimodal and complex distributions. Furthermore, the inference task in the joint space $[\theta, \sigma]$ is particularly challenging. We proposed a split strategy to tackle this problem. We consider an optimization approach over σ and a sampling scheme for θ . More specifically, we design an iterative procedure where this two tasks are alternated. Additionally, the actual maximum likelihood (ML) estimation of the noise power, $\widehat{\sigma}_{\text{ML}}^2$, is employed as a tempering parameter, starting from high values and then “cooling down” according to the ML estimations at each iterations. Therefore, the proposed scheme deals with a sequence of tempered posteriors according to the current estimation $\widehat{\sigma}_{\text{ML}}^2$. It is important to observe that, given a fixed vector θ , the ML estimation $\widehat{\sigma}_{\text{ML}}^2$ can be obtained analytically.

Furthermore, the complete Bayesian analysis regarding the joint posterior of θ and σ is also possible (as discussed in Section 5). This is obtained by implementing a proper re-weighting of the samples of generated by ATAIS, without any additional evaluations of the observation model. An approximation of the marginal posterior of σ is provided as well. The advantages of the proposed scheme are shown in two numerical experiments, one of them considering a complex astronomical model.

2 Problem Statement

Let us denote the observed measurements as $\mathbf{y} = [y_1, \dots, y_K]^\top \in \mathbb{R}^K$, and the variable of interest that we desire to infer, as $\theta = [\theta_1, \dots, \theta_M]^\top \in \Theta \subseteq \mathbb{R}^M$. Furthermore, let us consider the observation model

$$\mathbf{y} = \mathbf{f}(\theta) + \mathbf{v}, \quad (1)$$

where we have a nonlinear mapping,

$$\mathbf{f}(\theta) = [f_1(\theta), \dots, f_K(\theta)]^\top : \Theta \subseteq \mathbb{R}^M \rightarrow \mathbb{R}^K, \quad (2)$$

and a Gaussian perturbation noise,

$$\mathbf{v} = [v_1, \dots, v_K]^\top \sim \mathcal{N}(\mathbf{v}|\mathbf{0}, \sigma^2 \mathbf{I}_K), \quad (3)$$

with $\sigma > 0$, and we have denoted the K -dimensional unit matrix as \mathbf{I}_K . The noise variance σ^2 is unknown, in general. The mapping $\mathbf{f}(\boldsymbol{\theta})$ could be analytically unknown: the only assumption is that we are able to evaluate it pointwise. The likelihood function is

$$\ell(\mathbf{y}|\boldsymbol{\theta}, \sigma) = \frac{1}{(2\pi\sigma^2)^{K/2}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{f}(\boldsymbol{\theta})\|^2\right), \quad (4)$$

$$= \frac{1}{(2\pi\sigma^2)^{K/2}} \exp\left(-\frac{1}{2\sigma^2} \sum_{k=1}^K (y_k - f_k(\boldsymbol{\theta}))^2\right). \quad (5)$$

Note that we have two types of variables of interest: the vector $\boldsymbol{\theta}$ contains the parameters of the non-linear mapping $\mathbf{f}(\boldsymbol{\theta})$, whereas σ is a scale parameter of the likelihood function.

Given the vector of measurements \mathbf{y} , we desire to make inferences regarding the hidden parameters $\boldsymbol{\theta}$ and the noise power σ^2 , obtaining at least some point estimators $\widehat{\boldsymbol{\theta}}$ and $\widehat{\sigma}^2$. We are also interested in performing uncertainty and correlation analysis among the components of $\boldsymbol{\theta}$. Furthermore, we aim to perform model selection, i.e., to compare, select or properly average different models.

Bayesian inference in the complete space. We consider prior densities $g_\theta(\boldsymbol{\theta})$ and $g_\sigma(\sigma)$ over the unknowns. Hence, the complete posterior density is

$$p(\boldsymbol{\theta}, \sigma|\mathbf{y}) = \frac{1}{p(\mathbf{y})} p(\boldsymbol{\theta}, \sigma, \mathbf{y}) = \frac{1}{p(\mathbf{y})} \ell(\mathbf{y}|\boldsymbol{\theta}, \sigma) g_\theta(\boldsymbol{\theta}) g_\sigma(\sigma), \quad (6)$$

The marginal likelihood $Z = p(\mathbf{y})$ is

$$Z = p(\mathbf{y}) = \int_{\mathbb{R}^+} \int_{\Theta} \ell(\mathbf{y}|\boldsymbol{\theta}, \sigma) g_\theta(\boldsymbol{\theta}) g_\sigma(\sigma) d\boldsymbol{\theta} d\sigma, \quad (7)$$

This quantity is often needed for model selection. Since $Z(\mathbf{y})$ is generally unknown, we can usually evaluate pointwise the unnormalized posterior $\pi(\boldsymbol{\theta}, \sigma|\mathbf{y}) = \ell(\mathbf{y}|\boldsymbol{\theta}, \sigma) g_\theta(\boldsymbol{\theta}) g_\sigma(\sigma)$ (i.e., $p(\boldsymbol{\theta}, \sigma|\mathbf{y}) \propto \pi(\boldsymbol{\theta}, \sigma|\mathbf{y})$). More generally, the computation of integrals of the form

$$I = \int_{\mathbb{R}^+} \int_{\Theta} h(\boldsymbol{\theta}, \sigma) p(\boldsymbol{\theta}, \sigma|\mathbf{y}) d\boldsymbol{\theta} d\sigma, \quad (8)$$

where $h(\cdot) : \Theta \times \mathbb{R}^+ \rightarrow \mathbb{R}$ is an integrable function, is usually required. We consider a Monte Carlo quadrature approach for approximating the integral above and, more generally, provide a particle approximation of the joint posterior $p(\boldsymbol{\theta}, \sigma|\mathbf{y})$.

Main observation. Generally, generating random samples from a complicated posterior in Eq. (6) and computing efficiently the integrals as in Eqs. (7)-(8) is a hard task. Moreover, this task becomes often more difficult when we try to perform a joint inference where scale parameters are involved,

i.e., σ , and parameters of the nonlinearity, i.e., $\boldsymbol{\theta}$. Indeed, “wrong choices” of σ values can easily jeopardize the sampling of $\boldsymbol{\theta}$. In the next section, we describe a strategy that we propose to tackle this problem. Before, we need to recall some additional definitions.

Conditional and marginal posteriors. In order to design efficient computational schemes is often useful to consider the conditional posteriors, for instance,

$$\begin{aligned} p(\boldsymbol{\theta}|\mathbf{y}, \sigma) &= \frac{p(\boldsymbol{\theta}, \mathbf{y}, \sigma)}{p(\mathbf{y}, \sigma)} = \frac{\ell(\mathbf{y}|\boldsymbol{\theta}, \sigma)g_{\boldsymbol{\theta}}(\boldsymbol{\theta})g_{\sigma}(\sigma)}{p(\mathbf{y}|\sigma)g_{\sigma}(\sigma)}, \\ &= \frac{\ell(\mathbf{y}|\sigma, \boldsymbol{\theta})g_{\boldsymbol{\theta}}(\boldsymbol{\theta})}{p(\mathbf{y}|\sigma)}. \end{aligned} \quad (9)$$

In the next section, we will see that the idea underlying the proposed scheme is to split the space $[\boldsymbol{\theta}, \sigma]$, restricting the sampling problem only with respect to $\boldsymbol{\theta}$ and considering an optimization problem for with respect to σ . The conditional marginal likelihood is obtained by integrating out one of the two variables, i.e.,

$$Z(\sigma) = p(\mathbf{y}|\sigma) = \int_{\boldsymbol{\theta}} \ell(\mathbf{y}|\boldsymbol{\theta}, \sigma)g_{\boldsymbol{\theta}}(\boldsymbol{\theta})d\boldsymbol{\theta}. \quad (10)$$

The integral above cannot be computed analytically, in general. We can also consider marginal posteriors, for instance, the marginal posterior of σ is

$$p(\sigma|\mathbf{y}) = \frac{p(\mathbf{y}|\sigma)g_{\sigma}(\sigma)}{p(\mathbf{y})} = \frac{Z(\sigma)g_{\sigma}(\sigma)}{Z}. \quad (11)$$

Note that the joint posterior in Eq. (6) can be also written as

$$p(\boldsymbol{\theta}, \sigma|\mathbf{y}) = p(\boldsymbol{\theta}|\mathbf{y}, \sigma)p(\sigma|\mathbf{y}). \quad (12)$$

Underlying idea. The underlying idea of this work is to divide the inference study in two parts. In the first part (Sections 3 and 4), we focus on the study of the conditional posterior $p(\boldsymbol{\theta}|\mathbf{y}, \sigma)$ given a fixed σ . Then, in the second part (Section 5), we also estimate the marginal posterior $p(\sigma|\mathbf{y})$. Finally, using (12), we can obtain a final approximation of the complete posterior $p(\boldsymbol{\theta}, \sigma|\mathbf{y})$. Estimations of $Z(\sigma)$ and Z are also obtained.

3 Key observations and proposed approach

In the first part of work, we assume a proper (or improper) uniform prior over $\boldsymbol{\theta}$, i.e., $g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \propto 1$ in Θ . The possible use of a general choice of $g_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is discussed in Section 4.1. Let $\boldsymbol{\theta}_{\text{MAP}}$ denote the MAP estimator of $p(\boldsymbol{\theta}|\mathbf{y}, \sigma)$. Generally, $\boldsymbol{\theta}_{\text{MAP}}$ should be a function of σ , i.e., $\boldsymbol{\theta}_{\text{MAP}} = \boldsymbol{\theta}_{\text{MAP}}(\sigma)$. However, due to the choice of likelihood function (and the uniform prior) considered in this paper, we have that

$$\begin{aligned} \boldsymbol{\theta}_{\text{MAP}} &= \arg \max_{\boldsymbol{\theta}} \log p(\boldsymbol{\theta}|\mathbf{y}, \sigma) \quad \text{with} \quad g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \propto 1, \text{ for } \boldsymbol{\theta} \in \Theta, \\ &= \arg \min_{\boldsymbol{\theta}} \|\mathbf{y} - \mathbf{f}(\boldsymbol{\theta})\|^2, \end{aligned}$$

which does not depend on σ , i.e., we have that θ_{MAP} maximizes the conditional posterior $p(\theta|\mathbf{y}, \sigma)$ for any σ . See App. A for further details.

Furthermore, the variance of the conditional posterior $p(\theta|\mathbf{y}, \sigma)$ grows when σ increases. In this sense, with larger σ , the density $p(\theta|\mathbf{y}, \sigma)$ is ‘wider’, hence it is easier for Monte Carlo methods to explore the space (namely, we have a *tempering effect*). Based on these considerations, we can run Monte Carlo schemes (specifically IS algorithms) on $p(\theta|\mathbf{y}, \sigma_0)$ with a large value σ_0 for estimating θ_{MAP} more efficiently. Furthermore, apart from estimating θ_{MAP} , we are also interested in studying the conditional posterior $p(\theta|\mathbf{y}, \sigma_{\text{ML}})$ where

$$\sigma_{\text{ML}} = \arg \max_{\sigma} \ell(\mathbf{y}|\theta_{\text{MAP}}, \sigma).$$

The value σ_{ML} can be obtained in closed-form (see App. A). In fact, for any θ , we have

$$\ell(\mathbf{y}|\theta, \sigma) \propto \left(\frac{1}{\sigma^2} \right)^{\frac{K}{2}} \exp \left(-\frac{\|\mathbf{y} - \mathbf{f}(\theta)\|^2}{2\sigma^2} \right), \quad (13)$$

has the form of an *Inverse Gamma* density for σ^2 , and it has a *unique* mode at $\sqrt{\frac{1}{K}\|\mathbf{y} - \mathbf{f}(\theta)\|^2}$. Hence, finally we have

$$\sigma_{\text{ML}} = \sqrt{\frac{1}{K}\|\mathbf{y} - \mathbf{f}(\theta_{\text{MAP}})\|^2}.$$

This can serve as a point estimator of the noise power in the system, and also as possible value to stop the tempering of the conditional posterior, as we show in the following section.

3.1 Suggested iterative approach

Consider we start with a large value σ_0 , which can be viewed as a coarse approximation of σ_{ML} , so we denote it $\sigma_0 = \widehat{\sigma}_{\text{ML}}^{(0)}$. Let $\widehat{\theta}_{\text{MAP}}^{(1)}$ denote the IS estimate of θ_{MAP} obtained by working w.r.t. $p(\theta|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(0)})$. We use this current estimation to obtain the next value of σ , i.e., $\widehat{\sigma}_{\text{ML}}^{(1)} = \sqrt{\frac{1}{K}\|\mathbf{y} - \mathbf{f}(\widehat{\theta}_{\text{MAP}}^{(1)})\|^2}$. In general, $\widehat{\sigma}_{\text{ML}}^{(1)}$ will be a better estimate of σ_{ML} than $\widehat{\sigma}_{\text{ML}}^{(0)}$. We can iterate this procedure: for $t = 1, \dots, T$:

- 1 Estimate $\widehat{\theta}_{\text{MAP}}^{(t)}$ by Monte Carlo (e.g., by an IS scheme) working with respect to $p(\theta|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(t-1)})$.
- 2 Compute

$$\widehat{\sigma}_{\text{ML}}^{(t)} = \sqrt{\frac{1}{K}\|\mathbf{y} - \mathbf{f}(\widehat{\theta}_{\text{MAP}}^{(t)})\|^2}.$$

With this iterative scheme, we have that $\widehat{\sigma}_{\text{ML}}^{(T)} \rightarrow \sigma_{\text{ML}}$ as T grows, hence we eventually perform IS with respect to the density of interest $p(\theta|\mathbf{y}, \sigma_{\text{ML}})$. Furthermore, a non-increasing sequence of values $\widehat{\sigma}_{\text{ML}}^{(0)} \geq \widehat{\sigma}_{\text{ML}}^{(1)} \geq \dots \geq \widehat{\sigma}_{\text{ML}}^{(T)}$ is produced, which facilitates the estimation of θ_{MAP} , and ensures the IS estimation of $p(\theta|\mathbf{y}, \sigma_{\text{ML}})$ is performed efficiently by using the set of intermediate, tempered (i.e., wider) distributions $p(\theta|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(t)})$ for $t = 0, 1, \dots, T$. Finally, a particle approximation of $p(\theta|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(T)})$ is obtained, i.e.,

$$p(\theta|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(T)}) = \sum_{t=1}^T \sum_{n=1}^N \widetilde{w}_t^{(n)} \delta(\theta - \theta_t^{(n)}),$$

where $\sum_{t=1}^T \sum_{n=1}^N \tilde{w}_t^{(n)} = 1$.

4 Automatic Tempering Adaptive Importance Sampling (ATAIS)

In this section, we describe an adaptive importance sampler with an *automatic tempering* approach which follows the procedure given above. At each iteration t of the algorithm, we have an ML approximation of σ , i.e., $\widehat{\sigma}_{\text{ML}}^{(t-1)}$. Considering Eq. (9), we define the unnormalized *tempered conditional posterior* at the t -th iteration,

$$\pi_t(\boldsymbol{\theta}) = \ell(\mathbf{y}|\boldsymbol{\theta}, \widehat{\sigma}_{\text{ML}}^{(t-1)})g_{\boldsymbol{\theta}}(\boldsymbol{\theta}), \quad (14)$$

where we assume $g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \propto 1$ in Θ . For other generic choice of $g_{\boldsymbol{\theta}}(\boldsymbol{\theta})$, see the discussion in Section 4.1. At each iteration, we consider $p(\boldsymbol{\theta}|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(t-1)}) \propto \pi_t(\boldsymbol{\theta})$ as the target distribution. The dependence on the iteration t is due to $\widehat{\sigma}_{\text{ML}}^{(t)}$ varies with t . The ATAIS algorithm is outlined in Table 1. The resulting scheme is an adaptive IS algorithm which combines sampling schemes and stochastic optimization. It is important to remark that, if $\widehat{\sigma}_{\text{ML}}^{(0)}$ is bigger than the true ML value, we generate a non-increasing sequence of $\widehat{\sigma}_{\text{ML}}^{(t)}$, i.e., $\widehat{\sigma}_{\text{ML}}^{(0)} \geq \widehat{\sigma}_{\text{ML}}^{(1)} \geq \dots \widehat{\sigma}_{\text{ML}}^{(t)} \geq \widehat{\sigma}_{\text{ML}}^{(t+1)}$, etc. Note that this is true since we have assumed a uniform prior $g_{\boldsymbol{\theta}}(\boldsymbol{\theta})$.

IS steps. A set of N samples $\{\boldsymbol{\theta}_t^{(n)}\}_{n=1}^N$ are drawn from a (normalized) proposal density $q(\boldsymbol{\theta}|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ with mean $\boldsymbol{\mu}_t$ and a covariance matrix $\boldsymbol{\Sigma}_t$. An importance weight

$$w_t^{(n)} = \frac{\pi_t(\boldsymbol{\theta}_t^{(n)})}{q(\boldsymbol{\theta}_t^{(n)}|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)},$$

is assigned to each sample.

Proposal adaptation. A particle estimation of the conditional MAP estimator of $\boldsymbol{\theta}$ is given by $\widehat{\boldsymbol{\theta}}_t = \arg \max_n \pi_t(\boldsymbol{\theta}_t^{(n)})$. The value of current MAP approximation $\pi_t(\widehat{\boldsymbol{\theta}}_t)$ is then compared with the value of global MAP estimator obtained so far denoted as π_{MAP} . If $\pi_t(\widehat{\boldsymbol{\theta}}_t) \geq \pi_{\text{MAP}}$, all the global MAP estimators are updated and the proposal pdf is moved at $\widehat{\boldsymbol{\theta}}_t$, i.e., we set

$$\widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(t)} = \widehat{\boldsymbol{\theta}}_t, \quad \pi_{\text{MAP}} = \pi_t(\widehat{\boldsymbol{\theta}}_t), \quad \boldsymbol{\mu}_t = \widehat{\boldsymbol{\theta}}_t. \quad (15)$$

Otherwise, we keep the previous values $\widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(t)} = \widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(t-1)}$, π_{MAP} and $\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t-1}$. The covariance matrix $\boldsymbol{\Sigma}_t$ is adapted by considering the empirical covariance of the weighted samples. Note that, we set $\boldsymbol{\mu}_t = \widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(t)}$ instead of using the empirical mean of the samples (as in other classical AIS schemes). This is because we have noticed that this choice provides better and more robust results, specially as the dimension of the problem grows.

Automatic tempering. As we show in the previous section, the current ML estimator of σ can be obtained analytically as

$$\widehat{\sigma}_t = \sqrt{\frac{1}{K} \|\mathbf{y} - \widehat{\mathbf{r}}_t\|^2}, \quad (16)$$

where $\widehat{\mathbf{r}}_t = \mathbf{f}(\widehat{\boldsymbol{\theta}}_t)$. If the current ML estimator $\widehat{\sigma}_t$ is smaller than the current global one $\widehat{\sigma}_{\text{ML}}^{(t-1)}$, i.e., $\widehat{\sigma}_t < \widehat{\sigma}_{\text{ML}}^{(t-1)}$, then we update $\widehat{\sigma}_{\text{ML}}^{(t)} = \widehat{\sigma}_t$, Otherwise, we keep the value of $\widehat{\sigma}_{\text{ML}}^{(t)} = \widehat{\sigma}_{\text{ML}}^{(t-1)}$. Actually, with a uniform prior $g_{\theta}(\boldsymbol{\theta})$, every time that we update $\widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(t)}$ we also update $\widehat{\sigma}_{\text{ML}}^{(t)}$.

ATAIS outputs. After T iterations, a final correction of the weights is needed, i.e.,

$$\widetilde{w}_t^{(n)} = w_t^{(n)} \frac{\pi_{T+1}(\boldsymbol{\theta}_t^{(n)})}{\pi_t(\boldsymbol{\theta}_t^{(n)})}, \quad (17)$$

in order to obtain a particle approximation of the measure of the final conditional posterior $p(\boldsymbol{\theta}|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(T)}) \propto \pi_{T+1}(\boldsymbol{\theta})$. Thus, the algorithm returns the final estimators $\widehat{\boldsymbol{\theta}}_{\text{MAP}}^{(T)}$, $\widehat{\sigma}_{\text{ML}}^{(T)}$, and all the weighted samples $\{\boldsymbol{\theta}_t^{(n)}, \widetilde{w}_t^{(n)}\}$, for all $n = 1, \dots, N$ and $t = 1, \dots, T$. Other outputs can be obtained with a post-processing of the weighted samples, as shown below. Note that Eq. (17) does not require any additional evaluation of the model, and the error $e_t^{(n)} = \|\mathbf{y} - \mathbf{f}(\boldsymbol{\theta}_t^{(n)})\|^2$. Moreover, we can also use $e_t^{(n)}$ and $\{\boldsymbol{\theta}_t^{(n)}\}$ for building a particle approximation of any other conditional posterior $p(\boldsymbol{\theta}|\mathbf{y}, \sigma)$. This allows the study of the marginal posterior $p(\sigma|\mathbf{y})$ and provides the complete Bayesian inference, as we show in the next section.

4.1 For a generic prior $g_{\theta}(\boldsymbol{\theta})$

The ATAIS algorithm is based on the fact that $\boldsymbol{\theta}_{\text{MAP}}$ does not depend on σ . This allows us to progressively estimate it by targeting the sequence of tempered posteriors $\pi_t(\boldsymbol{\theta}) \propto p(\boldsymbol{\theta}|\mathbf{y}, \widehat{\sigma}_{\text{ML}}^{(t)})$ that share all the same MAP. However, in the case $g_{\theta}(\boldsymbol{\theta})$ is not uniform, we generally have one $\boldsymbol{\theta}_{\text{MAP}}(\sigma)$ for each $p(\boldsymbol{\theta}|\mathbf{y}, \sigma)$, and we could have that the sequence of $\boldsymbol{\theta}_{\text{MAP}}(\widehat{\sigma}_{\text{ML}}^{(t)})$ will not approach $\boldsymbol{\theta}_{\text{MAP}}(\sigma_{\text{ML}})$.

If the data are informative and the prior $g_{\theta}(\boldsymbol{\theta})$ is chosen such it is vague with respect to the likelihood, the position of $\boldsymbol{\theta}_{\text{MAP}}(\sigma)$ is not very sensitive to the value of σ . Namely, we have $\boldsymbol{\theta}_{\text{MAP}}(\widehat{\sigma}_{\text{ML}}^{(1)}) \approx \boldsymbol{\theta}_{\text{MAP}}(\widehat{\sigma}_{\text{ML}}^{(2)}) \approx \dots \approx \boldsymbol{\theta}_{\text{MAP}}(\sigma_{\text{ML}})$, and hence our algorithm can be applied in this context.

Table 1: ATAIS: AIS with automatic tempering

1. **Initializations:** Choose N , μ_1 , Σ_1 , and obtain an initialization for $\widehat{\sigma}_{\text{ML}}^{(0)}$, and set $\pi_{\text{MAP}} = 0$.

2. **For** $t = 1, \dots, T$:

(a) **Sampling:**

- i. Draw $\theta_t^{(1)}, \dots, \theta_t^{(N)} \sim q(\theta | \mu_t, \Sigma_t)$.
- ii. Assign to each sample the weights

$$w_t^{(n)} = \frac{\pi_t(\theta_t^{(n)})}{q(\theta_t^{(n)} | \mu_t, \Sigma_t)}, \quad n = 1, \dots, N. \quad (18)$$

(b) **Current maximum estimations:**

- i. Obtain $\widehat{\theta}_t = \arg \max_n \pi_t(\theta_t^{(n)})$, and compute $\widehat{\mathbf{r}}_t = \mathbf{f}(\widehat{\theta}_t)$
- ii. Compute $\widehat{\sigma}_t = \sqrt{\frac{1}{K} \|\mathbf{y} - \widehat{\mathbf{r}}_t\|^2}$.

(c) **Global maximum estimations:**

- i. If $\widehat{\sigma}_t \leq \widehat{\sigma}_{\text{ML}}^{(t-1)}$, then set $\widehat{\sigma}_{\text{ML}}^{(t)} = \widehat{\sigma}_t$. Otherwise, set $\widehat{\sigma}_{\text{ML}}^{(t)} = \widehat{\sigma}_{\text{ML}}^{(t-1)}$.
- ii. If $\pi_t(\widehat{\theta}_t) \geq \pi_{\text{MAP}}$, then set $\widehat{\theta}_{\text{MAP}}^{(t)} = \widehat{\theta}_t$ and $\pi_{\text{MAP}} = \pi_t(\widehat{\theta}_t)$. Otherwise, $\widehat{\theta}_{\text{MAP}}^{(t)} = \widehat{\theta}_{\text{MAP}}^{(t-1)}$ and keep the value of π_{MAP} .

(d) **Adaptation:** Set

$$\mu_t = \widehat{\theta}_{\text{MAP}}^{(t)}, \quad (19)$$

$$\Sigma_t = \sum_{n=1}^N \bar{w}_t^{(n)} (\theta_t^{(n)} - \bar{\theta}_t)^{\top} (\theta_t^{(n)} - \bar{\theta}_t) + \delta \mathbf{I}_M, \quad (20)$$

where $\bar{w}_t^{(n)} \frac{w_t^{(n)}}{\sum_{i=1}^N w_t^{(i)}}$ are the normalized weights, $\bar{\theta}_t = \sum_{n=1}^N \bar{w}_t^{(n)} \theta_t^{(n)}$ and $\delta > 0$.

3. **Output:** Return the final estimators $\widehat{\theta}_{\text{MAP}}^{(T)}$, $\widehat{\sigma}_{\text{ML}}^{(T)}$, and all the weighted samples $\{\theta_t^{(n)}, \bar{w}_t^{(n)}\}$, for all t and n , with the corrected weights

$$\widetilde{w}_t^{(n)} = w_t^{(n)} \frac{\pi_{T+1}(\theta_t^{(n)})}{\pi_t(\theta_t^{(n)})}. \quad (21)$$

5 Complete Bayesian inference with ATAIS

Let us assume we have a proper prior $g_{\theta}(\theta)$ and we introduce another proper prior $g_{\sigma}(\sigma)$ for σ . The outputs of the ATAIS algorithm can serve to approximate the normalizing constant of the joint pos-

terior $p(\boldsymbol{\theta}, \sigma | \mathbf{y}) \propto \ell(\mathbf{y} | \boldsymbol{\theta}, \sigma) g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) g_{\sigma}(\sigma)$, i.e., the so-called marginal likelihood or Bayesian model evidence, given by

$$Z = \int_{\mathbb{R}^+} \int_{\Theta} \ell(\mathbf{y} | \boldsymbol{\theta}, \sigma) g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) g_{\sigma}(\sigma) d\boldsymbol{\theta} d\sigma = \int_{\mathbb{R}^+} Z(\sigma) g_{\sigma}(\sigma) d\sigma, \quad (22)$$

where we have denoted $Z(\sigma) = \int_{\Theta} \ell(\mathbf{y} | \boldsymbol{\theta}, \sigma) g_{\boldsymbol{\theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta}$, usually called conditional marginal likelihood. The quantity Z is useful for model selection purposes. Furthermore, a complete Bayesian study of the joint posterior $p(\boldsymbol{\theta}, \sigma | \mathbf{y})$ can be provided as well.

Approximation of $Z(\sigma) = p(\mathbf{y} | \sigma)$. After the T iterations of ATAIS, we can also approximate the conditional marginal likelihood $Z(\sigma) = p(\mathbf{y} | \sigma)$ without additional evaluations of the target function. Indeed, saving the error values at each particles obtained for the computation of the likelihood function during ATAIS,

$$e_t^{(n)} = \|\mathbf{y} - \mathbf{f}(\boldsymbol{\theta}_t^{(n)})\|^2,$$

we can compute the IS weights,

$$\rho_t^{(n)}(\sigma) = \frac{\frac{1}{(2\pi\sigma^2)^{\frac{K}{2}}} \exp\left(-\frac{e_t^{(n)}}{2\sigma^2}\right) g_{\boldsymbol{\theta}}(\boldsymbol{\theta})}{q(\boldsymbol{\theta}_t^{(n)} | \boldsymbol{\mu}_t, \Sigma_t)},$$

for a generic value of σ . Thus, the IS estimator of the conditional marginal likelihood $Z(\sigma)$ is given by the arithmetic mean of the weights $\rho_t^{(n)}(\sigma)$,

$$\widehat{Z}(\sigma) = \widehat{p}(\mathbf{y} | \sigma) = \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N \rho_t^{(n)}(\sigma). \quad (23)$$

Approximation of Z . Drawing $\sigma^{(r)} \sim g_{\sigma}(\sigma)$, for $r = 1, \dots, R$, (or considering a deterministic grid), we can approximate the global marginal likelihood Z by applying simple Monte Carlo to the integral in Eq. (22),

$$\widehat{Z} = \frac{1}{R} \sum_{r=1}^R \widehat{Z}(\sigma^{(r)}). \quad (24)$$

Approximation of $p(\sigma | \mathbf{y})$. An approximation of the marginal posterior $p(\sigma | \mathbf{y}) = \frac{p(\mathbf{y} | \sigma) g_{\sigma}(\sigma)}{p(\mathbf{y})}$ can be also obtained as

$$p(\sigma | \mathbf{y}) \approx \widehat{p}(\sigma | \mathbf{y}) = \frac{\widehat{Z}(\sigma) g_{\sigma}(\sigma)}{\widehat{Z}}, \quad (25)$$

which can be used to approximate, e.g., the MAP of $p(\sigma | \mathbf{y})$ by $\sigma_{\text{MAP-marg}} \approx \arg \max_{\sigma} \widehat{Z}(\sigma) g_{\sigma}(\sigma)$. Other different moments of $p(\sigma | \mathbf{y})$ can be computed by a deterministic quadrature (since the problem is now one-dimensional) or applying noisy Monte Carlo approaches.

Complete Bayesian analysis. We can approximate the integral of interest as

$$I = \int_{\mathbb{R}^+} \int_{\Theta} h(\theta, \sigma) p(\theta, \sigma | \mathbf{y}) d\theta d\sigma, \quad (26)$$

$$= \int_{\mathbb{R}^+} \int_{\Theta} h(\theta, \sigma) p(\theta | \mathbf{y}, \sigma) p(\sigma | \mathbf{y}) d\theta d\sigma \quad (27)$$

$$\approx \frac{1}{J} \sum_{j=1}^J \sum_{t=1}^T \sum_{n=1}^N \bar{\rho}_t^{(n)}(\sigma^{(j)}) h(\theta_t^{(n)}, \sigma^{(j)}), \quad (28)$$

where

$$\bar{\rho}_t^{(n)}(\sigma^{(j)}) = \frac{\rho_t^{(n)}(\sigma^{(j)})}{\sum_{\tau=1}^T \sum_{i=1}^N \rho_{\tau}^{(i)}(\sigma^{(j)})},$$

and $\sigma^{(j)}$ are generated by applying a noisy MCMC with invariant density $\widehat{p}(\sigma | \mathbf{y}) \propto \widehat{Z}(\sigma) g_{\sigma}(\sigma)$. Note that the samples $\theta_t^{(n)}$ do not depend on the index j (they do not change) since we are *recycling* the particles generated by ATAIS and reusing evaluations $e_t^{(n)} = \|\mathbf{y} - \mathbf{f}(\theta_t^{(n)})\|^2$.

6 Simulations

We test the proposed scheme in two numerical examples. The first numerical experiment is a simple bidimensional example (which is easy to be reproduced). The second experiment considers a real-world application, i.e., a radial velocity models of exoplanet systems which is often employed in astronomy applications (with a dimension of the inference problem of 6 and 11).

6.1 First numerical analysis

For the sake of simplicity, let us consider $\theta \in \mathbb{R}$ and an observation model given by the equation

$$y_k = \theta^2 - \log(|\sin(10\theta)|) + v_k,$$

so that $f(\theta) = \theta^2 + \log(|\sin(10\theta)|)$, and $v_k \sim \mathcal{N}(0, \sigma^2)$. We consider $\theta_{\text{true}} = 2.5$, and $\sigma_{\text{true}} = 4$. We generate $K = 8$ observations from the model above. We also consider a uniform prior for θ in $[0, 20]$. The conditional posterior $p(\theta | \mathbf{y}, \sigma_{\text{true}})$ is shown in Figure 1(c). We can observe that $p(\theta | \mathbf{y}, \sigma_{\text{true}})$ is highly multimodal. Figure 1 also depicts the conditional posteriors $p(\theta | \mathbf{y}, \sigma)$ with $\sigma \in \{10, 20\}$. Considering also a uniform prior over σ in $[0, 20]$, we have also a bidimensional joint posterior over $[\theta, \sigma]$, which is depicted in Fig. 2(a).

In this bidimensional example, it is possible to obtain the ground-truths using an expensive thin grid. The expected value and variance of the conditional posterior $p(\theta | \mathbf{y}, \sigma_{\text{ML}})$, approximated by a thin grid, are $E[\theta | \mathbf{y}, \sigma_{\text{ML}}] = 2.48$ and $\text{var}[\theta | \mathbf{y}, \sigma_{\text{ML}}] = 0.11$. The MAP estimator of the conditional posterior $p(\theta | \mathbf{y}, \sigma_{\text{ML}})$ is $\theta_{\text{MAP}} = 2.56$. The expected value and variance of the marginal posteriors of θ and σ are $E[\theta | \mathbf{y}] = 2.46$, $\text{var}[\theta | \mathbf{y}] = 0.18$, $E[\sigma | \mathbf{y}] = 4.32$ and $\text{var}[\sigma | \mathbf{y}] = 2.43$ (these values coincide with the expected value and the diagonal of the covariance matrix of the complete posterior). The MAP estimator provided by the complete posterior is $[\theta_{\text{MAP-joint}}, \sigma_{\text{MAP-joint}}] = [2.56, 3.23]$. Since the prior

over σ is uniform, the maximum likelihood of σ is $\sigma_{\text{ML}} = \sigma_{\text{MAP-joint}} = 3.23$. These values can be obtained sampling in the 2D space $[\theta, \sigma]$ and then considering the components of the sampled vectors. The MAP estimator of the marginal posteriors in Eq. (11) are $\theta_{\text{MAP-marg}} = 2.56$ and $\sigma_{\text{MAP-marg}} = 3.46$. The two marginal posteriors are shown in Figures 2(b)-(c).

We apply ATAIS with the goal of estimating the expected value and the variance of the posterior density with respect to θ . We consider a Gaussian proposal $q(\theta|\mu_t, \lambda_t)$ with $\mu_0 = 10$ and a starting variance of $\lambda_0 = 4$. Note that μ_0 is located in a region that does not contain modes. We also start with $\widehat{\sigma}_{\text{ML}}^{(0)} = 20$ and $\pi_{\text{MAP}} = 0$ (initial conditions). The Mean Square Error (MSE) of ATAIS, averaged over 500 runs, in estimation of different moments and modes as function of N (and with $T = 10$), is given in Table 2. The ML estimation $\widehat{\sigma}_{\text{ML}}^{(t)}$ as function of the iteration t (with $N = 5$) for different runs, is given in Figure 3(a). The approximation of the marginal posterior $p(\sigma|y)$, denoted $\widehat{p}(\sigma|y)$, is obtained as in Eq. (25) in one specific run, with different $N \in \{10, 100, 500\}$ and $T = 10$.




Figure 1: Conditional posteriors corresponding to different values of σ : (a) $\sigma = 20$, (b) $\sigma = 10$, (c) $\sigma = \sigma_{\text{true}} = 4$.




Figure 2: The bidimensional complete posterior $p(\theta, \sigma|y)$ and the two marginal posteriors $p(\theta|y)$, $p(\sigma|y)$ in Eq. (11), computed by using a thin grid approximation.




Figure 3: **(a)** The ML estimation $\widehat{\sigma}_{ML}^{(t)}$ (different runs) versus the iteration t , with $N = 5$. **(b)** The true marginal posterior $p(\sigma|y)$ and different approximations, in one specific run, $\widehat{p}(\sigma|y)$ obtained as in Eq. (25) with different $N \in \{10, 100, 500\}$ and $T = 10$ (hence, the total number of samples are NT).

Table 2: MSE of ATAIS (averaged over 500 runs), in the estimation of different moments and modes as function of N and $T = 10$.

Moments-Modes	$N = 10$	$N = 100$	$N = 1000$	$N = 5000$
$E[\theta y, \sigma_{ML}]$	0.0311	0.0098	0.0034	0.0024
$\text{var}[\theta y, \sigma_{ML}]$	0.0474	0.0370	0.0298	0.0201
θ_{MAP}	0.0410	0.0337	0.0285	0.0127
$E[\sigma y]$	0.9233	0.0785	0.0097	0.0023
$\text{var}[\sigma y]$	6.1869	0.2640	0.0035	0.0010
$\sigma_{\text{MAP-marg}}$	0.0056	0.0004	0.0001	$3 \cdot 10^{-5}$
σ_{ML}	$8 \cdot 10^{-5}$	$2 \cdot 10^{-5}$	$5 \cdot 10^{-7}$	$6 \cdot 10^{-9}$

6.2 Radial velocity curves of exoplanets and binary systems

In this example, we consider an application in an astronomical model. In recent years, the problem of revealing objects orbiting other stars has acquired large attention. Different techniques have been proposed to discover exo-objects but, nowadays, the radial velocity technique is still the most used [15, 16, 17, 18]. The problem consists in fitting a model (the so-called radial velocity curve) to data acquired at different moments spanning during long time periods (up to years). The model is highly non-linear and it is costly in terms of computation time (specially, for certain sets of parameters). Obtaining a value to compare to a single observation involves numerically integrating a differential equation in time or an iterative procedure for solving to a non-linear equation. Typically, the iteration is performed until a threshold is reached or 10^6 iterations are performed. The problem of radial velocity curve fitting is applied in several related applications.

Observation model - likelihood. When analysing radial velocity data of an exoplanetary system, it

is commonly accepted that the *wobbling* of the star around the centre of mass is caused by the sum of the gravitational force of each planet independently and that they do not interact with each other. Each planet follows a Keplerian orbit and the radial velocity of the host star is given by

$$y_{r,t} = V_0 + \sum_{i=1}^S K_i [\cos(u_{i,t} + \omega_i) + e_i \cos(\omega_i)] + \xi_t, \quad (29)$$

with $t = 1, \dots, T$ and $r = 1, \dots, R$. The number of objects in the system is S , that is consider known in this experiment (for the sake of simplicity). Both $y_{r,t}$, $u_{i,t}$ depend on time t , and then ξ_t is a Gaussian noise perturbation with variance σ^2 . The likelihood function is defined by (29) and some indicator variables described below. The angle $u_{i,t}$ is the true anomaly of the planet i and it can be determined from

$$\frac{du_{i,t}}{dt} = \frac{2\pi}{P_i} \frac{(1 + e_i \cos u_{i,t})^2}{(1 - e_i)^{\frac{3}{2}}} \quad (30)$$

This equation has analytical solution. As a result, the true anomaly u_t can be determined from the mean anomaly M . However, the analytical solution contains a non linear term that needs to be determined by iterating. First, we define the mean anomaly $M_{i,t}$ as

$$M_{i,t} = \frac{2\pi}{P_i} (t - \tau_i), \quad (31)$$

where τ_i is the time of periastron passage of the planet i and P_i is the period of its orbit. Then, through the Kepler's equation,

$$M_{i,t} = E_{i,t} - e_i \sin E_{i,t}, \quad (32)$$

where $E_{i,t}$ is the eccentric anomaly. Equation (32) has no analytic solution and it must be solved by an iterative procedure. A Newton-Raphson method is typically used to find the roots of this equation [19]. For certain sets of parameters this iterative procedure, can be particularly slow. We also have

$$\tan \frac{u_{i,t}}{2} = \sqrt{\frac{1 + e_i}{1 - e_i}} \tan \frac{E_{i,t}}{2}, \quad (33)$$

The variable of interest is then $\boldsymbol{\theta}$ is the vector

$$\boldsymbol{\theta} = [V_0, k_{a,1}, \omega_1, e_1, P_1, \tau_1, \dots, k_{a,S}, \omega_S, e_S, P_S, \tau_S], \quad (34)$$

Then, for a single object (e.g., a planet or a natural satellite), the dimension of $\boldsymbol{\theta}$ is $M = 5 + 1 = 6$, with two objects the dimension of $\boldsymbol{\theta}$ is $M = 11$ etc.

This example consists in a synthetic radial velocity curve of a planetary system with one planet or two planets (i.e., $S = 1$ or $S = 2$). More specifically, we generate simulated data with a model with two planets. The orbital parameters of the planets are listed in Table 3, where P is the period of the orbit, k_a is the amplitude of the curve, e is the eccentricity of the orbit, ω is the argument of perigee and τ is the last periastron passage. A mean velocity $v_0 = 5 \text{ m s}^{-1}$ is assumed. A Gaussian noise perturbation is added with a standard deviation $\sigma = 3 \text{ m s}^{-1}$. To simulate observations, a total of $K = 120$ points are selected from three, random time periods (and two planets in the system). Note

that the amplitude of the radial velocity curve of the second planet is close to the noise level. We run ATAIS and a standard AIS scheme with the model with one planet and with the model with two planets. The purpose of this simulation is to check the ability of the method to detect the two planets (by approximating the model evidence).

Table 3: Main orbital parameters of the two exoplanets in the simulation.

Parameter	Planet 1	Planet 2
P	15 d	115 d
k_a	25 m s^{-1}	5 m s^{-1}
e	0.1	0.0
ω	35°	10°
τ	3 d	24 d




Figure 4: Comparison of the results of the ATAIS algorithm with the simulations (blue dots). Left panel shows, in grey, the radial velocity curve for $\widehat{\theta}_{\text{MAP}}$ using a model with one planet. Right panel is like left panel but considering a model with two planets.

We apply ATAIS and a standard AIS scheme [13] over the space $[\theta, \sigma]$ for approximating the model evidence $Z = p(\mathbf{y})$ (marginal likelihood) of both models (one planet or two planets) with the given data (generated considering two planets). Uniform priors are considered for each parameter: $P \in [0, 365]$, $k_a \in [-20, 20]$, $e \in [0, 1]$, $\omega \in [0, 2\pi]$, and $\tau \in [0, 50]$ (moreover, $\sigma \in [0, 30]$ for the standard AIS scheme). The ATAIS algorithm and the standard AIS scheme has been run with $N = 10^5$ and $T = 50$ iterations for both, the model with one and two planets. In both cases, we consider the same Gaussian proposal with a starting standard deviation of 5 for each component (note that the standard AIS scheme works in higher dimensional space due the inference over σ). To decide which model is more probable, the model evidence Z of each model is estimated. More specifically, we approximate the one-planet model $\widehat{Z}_1 = \widehat{p}_1(\mathbf{y})$, and the two-planets model $\widehat{Z}_2 = \widehat{p}_2(\mathbf{y})$ with the ATAIS algorithm and the standard AIS scheme. When $\widehat{Z}_1 > \widehat{Z}_2$ we select the first model otherwise if, $\widehat{Z}_1 < \widehat{Z}_2$, we select the second one. The true model is the two-planets model, since the simulated data were

generated from that model. After 500 independent runs, the percentage of correct detection of the true model for ATAIS is $\approx 98\%$, whereas with the standard AIS scheme is only $\approx 56\%$. This is due to the difficulty of making inference jointly over $[\theta, \sigma]$. In ATAIS, the ratio between the model evidences (averaged over the 500 runs) is $Z_2/Z_1 \approx 5 \cdot 10^3$. Therefore, for ATAIS, the model with two planets is clearly more probable than the model with one planet.

The fitted curves, corresponding to the vector of parameters $\widehat{\theta}_{\text{MAP}}$ obtained with ATAIS, are shown in Fig. 4. From the figure, it is not clear which model fits better the simulated observations (blue points), although the model with two planets seems to fit better the observations in the time period from 200 to 300 days. The values of $\widehat{\theta}_{\text{MAP}}$, obtained in one specific run by ATAIS, is given in Table 4. We notice that ω and τ are highly correlated and more iterations may be needed to obtain the actual global maximum, but the remaining parameters obtained from $\widehat{\theta}_{\text{MAP}}$ are similar to the simulated values. In addition, the amplitude of the curve of the second planet is close to the intensity of the noise, what makes difficult to derive the best fit for that planet. Summarizing, our results show the method is able to discriminate between a model with one planet (with 6 dimensions of the inference problem) and a model with two planets (with 11 dimensions of the inference problem), for this particular simulation. Finally, the evolution of the automatic tempering parameter $\widehat{\sigma}_{\text{ML}}^{(t)}$ is shown in Fig. 5. The dashed line is the evolution of $\widehat{\sigma}_{\text{ML}}^{(t)}$ for the single-planet model, whereas the continuous line is the evolution of $\widehat{\sigma}_{\text{ML}}^{(t)}$ for the model with two planets. In this second model, the tempering parameters reaches a smaller value, as expected.

Table 4: The value of $\widehat{\theta}_{\text{MAP}}$ for the 2-planets model

Parameter	Planet 1	Planet 2
P	14.99 d	110.39 d
K	23.78 m s^{-1}	3.50 m s^{-1}
e	0.05	0.00
ω	43.7°	0.39°
τ	6.8 d	7.96 d




Figure 5: Evolution of the tempering parameter $\widehat{\sigma}_{\text{ML}}^{(t)}$. The dashed line is the evolution for the model with one planet. The continuous line is the evolution of the two-planets model.

7 Conclusions

We have proposed a novel AIS scheme for Bayesian inversion problems where an automatic tempering procedure is implemented (called ATAIS). The inference of the variables of interest θ and the noise power σ^2 is divided. A sampling strategy is considered for θ and an optimization approach is employed for σ^2 . Thus, ATAIS performs an iterative procedure, alternating sampling and optimization steps. Therefore, the proposed scheme deals with a sequence of tempered posteriors according to the current estimation of the noise power. We have also discussed the possibility of approximating the marginal posterior of σ without additional evaluations of the complex model. Furthermore, the complete Bayesian analysis regarding the complete joint posterior is possible as discussed in Section 5, again without any additional evaluations of the likelihood function.

Several simulations are provided and the application to a sophisticated astronomical model has been considered, where the number of planets in the system is detected by the analysis of the marginal likelihood. The results show the benefits of the proposed scheme. For instance, in the astronomical example, the percentage of correct detection of the true model obtained by ATAIS is $\approx 98\%$, whereas with the standard AIS scheme is only $\approx 56\%$. As future research, we plan to extend the ATAIS scheme in order to deal with an observation model with correlated noise perturbations (for instance, using a Gaussian Process).

References

- [1] W. J. Fitzgerald. Markov chain Monte Carlo methods with applications to signal processing. *Signal Processing*, 81(1):3–18, January 2001.
- [2] C. Andrieu, N. de Freitas, A. Doucet, and M. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50:5–43, 2003.
- [3] L. Martino and J. Míguez. Generalized rejection sampling schemes and applications in signal processing. *Signal Processing*, 90(11):2981–2995, November 2010.
- [4] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2004.
- [5] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, 2004.
- [6] L. Martino, D. Luengo, and J. Miguez. *Independent Random Sampling methods*. Springer, 2018.
- [7] S. K. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [8] E. Marinari and G. Parisi. Simulated tempering: a new Monte Carlo scheme. *Europhysics Letters*, 19(6):451–458, July 1992.
- [9] N. Friel and A. N. Pettitt. Marginal likelihood estimation via power posteriors. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 70(3):589–607, 2008.

- [10] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.
- [11] R. M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- [12] F. Llorente, L. Martino, D. Delgado, and J. Lopez-Santiago. Marginal likelihood computation for model selection and hypothesis testing: an extensive review. *arXiv:2005.08334*, pages 1–91, 2020.
- [13] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Miguez, and P. M. Djuric. Adaptive importance sampling: The past, the present, and the future. *IEEE Signal Processing Magazine*, 34(4):60–79, 2017.
- [14] M. F. Bugallo, L. Martino, and J. Corander. Adaptive importance sampling in signal processing. *Digital Signal Processing*, 47:36–49, 2015.
- [15] Philip C. Gregory. Bayesian re-analysis of the Gliese 581 exoplanet system. *Monthly Notices of the Royal Astronomical Society*, 415(3):2523–2545, August 2011.
- [16] S. C. C. Barros, D. J. A. Brown, G. Hébrard, Y. Gómez Maqueo Chew, D. R. Anderson, P. Boumis, L. Delrez, K. L. Hay, K. W. F. Lam, J. Llama, M. Lendl, J. McCormac, B. Skiff, B. Smalley, O. Turner, M. Vanhuysse, D. J. Armstrong, I. Boisse, F. Bouchy, A. Collier Cameron, F. Faedi, M. Gillon, C. Hellier, E. Jehin, A. Liakos, J. Meaburn, H. P. Osborn, F. Pepe, I. Plauchu-Frayn, D. Pollacco, D. Queloz, J. Rey, J. Spake, D. Ségransan, A. H. M. Triaud, S. Udry, S. R. Walker, C. A. Watson, R. G. West, and P. J. Wheatley. WASP-113b and WASP-114b, two inflated hot Jupiters with contrasting densities. *Astronomy and Astrophysics*, 593:A113, 2016.
- [17] L. Affer, M. Damasso, G. Micela, E. Poretti, G. Scandariato, J. Maldonado, A. F. Lanza, E. Covino, A. Garrido Rubio, J. I. González Hernández, R. Gratton, G. Leto, A. Maggio, M. Perger, A. Sozzetti, A. Suárez Mascareño, A. S. Bonomo, F. Borsa, R. Claudi, R. Cosentino, S. Desidera, P. Giacobbe, E. Molinari, M. Pedani, M. Pinamonti, R. Rebolo, I. Ribas, and B. Toledo-Padrón. HADES RV program with HARPS-N at the TNG. IX. A super-Earth around the M dwarf Gl 686. *arXiv:1901.05338*, 622:A193, February 2019.
- [18] Trifon Trifonov, Stephan Stock, Thomas Henning, Sabine Reffert, Martin Kürster, Man Hoi Lee, Bertram Bitsch, R. Paul Butler, and Steven S. Vogt. Two Jovian Planets around the Giant Star HD 202696: A Growing Population of Packed Massive Planetary Pairs around Massive Stars? *The Astronomical Journal*, 157(3):93, March 2019.
- [19] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical recipes in C++ : the art of scientific computing*. Springer, 2002.

A On the optimization of the likelihood function

Let us set $\delta = \sigma^2$ and consider to optimize of the likelihood function

$$\ell(\boldsymbol{\theta}, \delta) = \frac{1}{(2\pi\delta)^{K/2}} \exp\left(-\frac{V(\boldsymbol{\theta})}{\delta}\right).$$

Recall that, in our model, we have $V(\boldsymbol{\theta}) = \|\mathbf{y} - \mathbf{f}(\boldsymbol{\theta})\|^2$. We desire to obtain

$$[\boldsymbol{\theta}_{\text{ML}}, \delta_{\text{ML}}] = \arg \max \ell(\boldsymbol{\theta}, \delta).$$

We can write the gradient and equal to zero,

$$\begin{cases} \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}, \delta) = -\frac{1}{\delta} \nabla_{\boldsymbol{\theta}} V(\boldsymbol{\theta}) \left[\frac{1}{(2\pi\delta)^{K/2}} \exp\left(-\frac{V(\boldsymbol{\theta})}{\delta}\right) \right] = \mathbf{0} \implies \nabla_{\boldsymbol{\theta}} V(\boldsymbol{\theta}) = \mathbf{0}, \\ \frac{\partial \ell(\boldsymbol{\theta}, \delta)}{\partial \delta} = \frac{e^{-\frac{V(\boldsymbol{\theta})}{\delta}} (2V(\boldsymbol{\theta}) - \delta K)}{2^{\frac{K}{2}+1} \delta^{\frac{K}{2}+2} \pi^{K/2}} = 0 \implies \delta = \frac{2}{K} V(\boldsymbol{\theta}). \end{cases} \quad (35)$$

We have obtained that the ML solution is defined by the system of equations,

$$\begin{cases} \nabla_{\boldsymbol{\theta}} V(\boldsymbol{\theta}_{\text{ML}}) = \mathbf{0} \\ \delta_{\text{ML}} = \frac{2}{K} V(\boldsymbol{\theta}_{\text{ML}}). \end{cases} \quad (36)$$