

Probabilities of exoplanet signals from posterior samplings

Mikko Tuomi^{1,2*} and Hugh R. A. Jones¹

¹ University of Hertfordshire, Centre for Astrophysics Research, Science and Technology Research Institute, College Lane, AL10 9AB, Hatfield, UK

² University of Turku, Tuorla Observatory, Department of Physics and Astronomy, Väisäläntie 20, FI-21500, Piikkiö, Finland

Received XX.XX.2011 / Accepted XX.XX.XXXX

ABSTRACT

Aims. Estimating the marginal likelihoods is an essential feature of model selection in the Bayesian context. It is especially crucial to have good estimates when assessing the number of planets orbiting stars when the models explain the noisy data with different numbers of Keplerian signals. We introduce a simple method for approximating the marginal likelihoods in practice when a statistically representative sample from the parameter posterior density is available.

Methods. We use our truncated posterior mixture estimate to receive accurate model probabilities for models with differing number of Keplerian signals in radial velocity data. We test this estimate in simple scenarios to assess its accuracy and rate of convergence in practice when the corresponding estimates calculated using deviance information criterion can be applied to receive trustworthy results for reliable comparison. As a test case, we determine the posterior probability of a planet orbiting HD 3651 given Lick and Keck radial velocity data.

Results. The posterior mixture estimate appears to be a simple and an accurate way of calculating marginal integrals from posterior samples. We show, that it can be used to estimate the marginal integrals reliably in practice, given a suitable selection of parameter λ , that controls its accuracy and convergence rate. It is also more accurate than the one block Metropolis-Hastings estimate and can be used in any application because it is not based on assumptions on the nature of the posterior density nor the amount of data or parameters in the statistical model.

Key words. Methods: Statistical, Numerical – Techniques: Radial velocities – Stars: Individual: HD 3651

1. Introduction

The selection between a collection of candidate models is of significant in all fields of astronomy but especially so, when the purpose is to extract weak planetary signals from noisy data. The ability to tell whether a signal is present in data as reliably as possible is essential in several searches for low-mass exoplanets orbiting nearby stars, whether made using the Doppler spectroscopy method; e.g. the Anglo-Australian Planet Search (e.g. Tinney et al., 2001; Jones et al., 2002, and references therein), High-Accuracy Radial Velocity Planet Searcher (e.g. Mayor et al., 2003; Lovis et al., 2011, and references therein), High Resolution Echelle Spectrometer (e.g. Vogt et al., 1994, 2010, and references therein); by searching photometric transits; e.g. Convection Rotation and Planetary Transits (e.g. Barge et al., 2007; Hébrard et al., 2011, and references therein), WASP (e.g. Collier Cameron et al., 2007; Faedi et al., 2011, and references therein); or other possible techniques, such as astrometry (e.g. Benedict et al., 2002; Pravdo & Shaklan, 2009) and transit timing (e.g. Holman & Murray, 2005) or other current or future methods.

Using Bayesian tools, it is possible to determine the relative probabilities for each statistical model in some selected collection of models to assess their relative performance, or relative ability to explain the data in a probabilistic manner. This is also important to be able to assess their inability to explain several data sets in terms of the model inadequacy of Tuomi et al. (2011). Especially, when different statistical models contain dif-

ferent numbers of planets orbiting the target star, assessing their relative posterior probabilities given the measurements is extremely important to detect all the signals in the data (e.g. Gregory, 2005, 2007a,b; Tuomi & Kotiranta, 2009) and to avoid the detection of false positives (e.g. Bean et al., 2010; Tuomi, 2011). However, determining the posterior probabilities require the ability to calculate marginal integrals that are complicated multidimensional integrals of likelihood functions and priors over the whole parameter space. While there are several methods of estimating the values of these integrals, those that are computationally simple and easy to implement are more often than not the poorest ones with respect to their accuracy and convergence properties (e.g. Kass & Raftery, 1995; Clyde et al., 2007; Ford & Gregory, 2007). Similarly, there are more complicated methods for estimating multidimensional integrals but they may provide more difficult computational problems themselves than typical data analyses are, which makes it difficult to use them in practice.

Because of these difficulties, we introduce a simple method for estimating the marginal integrals in practice if a statistically representative sample from the parameter posterior density exists. As such a sample is usually calculated when assessing the posterior densities of model parameters using posterior sampling algorithms (e.g. Metropolis et al., 1953; Hastings, 1970; Haario et al., 2001), the ability to use the very same sample in determining the marginal integral is extremely useful in practice. There are methods of taking advantage of the posterior sample in this manner (e.g. Newton & Raftery, 1994; Kass & Raftery, 1995; Chib & Jeliazkov, 2001; Clyde et al., 2007) but their performance, despite some studies (e.g. Kass & Raftery, 1995;

* The corresponding author, e-mail: m.tuomi@herts.ac.uk; mikko.tuomi@utu.fi

Ford & Gregory, 2007), is not generally well known in astronomical problems and some of them may also require sampling from other densities simultaneously, such as the prior density or the proposal density of the Metropolis-Hastings (M-H) output, making their application difficult.

In this article, we introduce a simple method that can be used to receive accurate estimates for the marginal integral. We test our estimate, called the truncated posterior mixture (TPM) estimate in scenarios where the marginal integral can be calculated accurately using simple existing methods. The deviance information criterion (DIC, Spiegelhalter et al., 2002) is asymptotically an accurate estimate when the sample size, i.e. the sample drawn from the posterior density, increases and can be used if the posterior is a multivariate Gaussian. Therefore, we compare our estimate with the DIC estimate in such cases to test its accuracy in practice. If accurate, our estimate is applicable whenever a statistically representative sample from the posterior is available because we do not make any assumptions regarding the shape of the posterior density when deriving the TPM estimate. The only assumption is, that such a sample exists and it is statistically representative. We also calculate the marginal likelihoods using the simple Akaike information criterion (AIC) for small sample size (Akaike, 1973; Burnham & Anderson, 2002), the harmonic mean (HM) estimate that is a special case of the TPM with poor convergence properties, and the One Block Metropolis-Hastings (OBMH) method of Chib & Jeliazkov (2001) that requires the simultaneous sampling of posterior and proposal densities. Kass & Raftery (1995) and Clyde et al. (2007) give detailed summaries of different methods.

2. Estimating marginal integrals

In the Bayesian context, the models in some *a priori* selected collection can be equipped with relative numbers representing the probabilities of having observed the data m if the model was a correct one. Therefore, for k different models $\mathcal{M}_1, \dots, \mathcal{M}_k$, these probabilities are calculated as

$$P(\mathcal{M}_i|m) = \frac{P(m|\mathcal{M}_i)P(\mathcal{M}_i)}{\sum_{j=1}^k P(m|\mathcal{M}_j)P(\mathcal{M}_j)}, \quad (1)$$

where the marginal integrals, sometimes called the marginal likelihoods, are defined as

$$P(m|\mathcal{M}_i) = \int l(m|\theta_i, \mathcal{M}_i)\pi(\theta_i|\mathcal{M}_i)d\theta_i \quad (2)$$

and l denotes the likelihood function and $\pi(\theta|\mathcal{M}_i)$ is the prior density of the parameters.

The truncated posterior mixture estimate that approximates the marginal integral is defined as (see Appendix)

$$\hat{P}_{TPM} = \left[\sum_{i=1}^N \frac{l_i}{(1-\lambda)l_i + \lambda l_{i-h}} \right] \left[\sum_{i=1}^N \frac{1}{(1-\lambda)l_i + \lambda l_{i-h}} \right]^{-1}, \quad (3)$$

where l_i is the value of the likelihood function at θ_i , and $\lambda \in [0, 1]$ and $h \in \mathbb{N}$ are parameters that control the convergence and accuracy properties of the estimate. While it is easy to select h – it only needs to be large enough such that θ_i and θ_{i-h} are independent – selecting parameter λ is more difficult. If λ is too large, the sample from the posterior is not close to the sample from the importance sampling function g in the Eq. (A.6) in the Appendix, and the resulting estimate for the marginal is biased. Conversely, too small values of λ , while making the

estimate more accurate, decrease its convergence rate because the estimate approaches asymptotically the HM estimate that is known to have extremely poor convergence properties (see the Appendix and Kass & Raftery, 1995). Therefore, we test different values of λ to find the best choice in applications. We note, however, that when θ_i and θ_{i-h} are independent, i.e. when h is large enough given the mixing properties of the Markov chain used to draw a sample from the posterior density, the TPM can converge to the marginal integral. The reason is that it is clear from Eq. (3) that occasional very small values of l_i , that consequently have a large impact on the sums in the estimate, do not slow down the convergence as much as they would in the HM estimate because it is unlikely that l_{i-h} is also small at the same time. This is the key feature in the TPM estimate that ensures its relatively rapid convergence in practice.

We estimate the integral in Eq. (2) using five methods. The HM estimate (see Appendix), the truncated posterior mixture estimate introduced here, the DIC, AIC, and the OBMH method of Chib & Jeliazkov (2001). While the DIC is a reasonably practical estimate in certain cases, it requires that the posterior is unimodal and symmetric and can be approximated as a multivariate Gaussian density, which is only rarely the case in applications. It can be easily calculated by using the average of the likelihoods and the likelihood of the parameter mean, which also reveals why the posterior needs to be unimodal and symmetric for reliable results. These means do not reflect the properties of the posterior in the cases of skewness and multiple modes, not to mention nonlinear correlations between some parameters in vector θ . The DIC is asymptotically accurate when the sample size becomes large (Spiegelhalter et al., 2002). We do not consider the HM estimate to be a trustworthy one but calculate its value because it is a special case of the truncated posterior mixture estimate when $\lambda = 0$ (or 1). The AIC could provide a reasonably accurate estimate in practice, and therefore we compare its performance in various scenarios. However, it relies on the maximum likelihood parameter estimate, and does not therefore take into account the prior information on the model parameters. Its accuracy also decreases as the amount of parameters in the model increases or the number of measurements decreases. Finally, we calculate the OBMH estimate (Chib & Jeliazkov, 2001). While this estimate appears to provide reliable results, e.g. the number of companions orbiting Gliese 581 determined in Tuomi (2011) was supported by additional data (Forveille et al., 2011), its performance has not been studied thoroughly with examples. It is also computationally more expensive than the TPM estimate, and indeed the other estimates compared here, because it requires the simultaneous sampling from the proposal density of the M-H algorithm.

When assessing the convergence of our TPM estimate given some selection of λ , we say that it has converged if the estimate at the i th member of the Markov chain, namely $\hat{P}_{TPM}(i)$, satisfies $|\hat{P}_{TPM}(i+k) - \hat{P}_{TPM}(i)| < r$ for all $k > 0$ and some small number r – in accordance of the standard definition of convergence. However, in practice, we use the logarithms of $\hat{P}_{TPM}$ and a value of $r = 0.1$ on the logarithmic scale for simplicity. We also approximate the estimate as having converged if the convergence condition holds for $0 < k < 10^5$ for practical reasons. While all the estimates except the AIC (which is based only on the maximum likelihood value) converge the better the greater sample they are based on, we only plot this convergence for the TPM estimate. For DIC, HM, and OBMH, we calculate the final estimate using the mean and standard deviation of values from several samplings.

3. Comparison of estimates: radial velocities of HD 3651

To assess the performance of the TPM estimate for the marginal integral, we compare its performance with different selections of parameter λ in simple cases where the marginal integral can be calculated reliably using the DIC, i.e. when the model parameters receive close-Gaussian posteriors and the sample size is large. Therefore, we select as test cases radial velocity time-series made using several telescope-instrument combinations that have different velocity offsets and different noise levels. The simple model without any Keplerian signals provides a suitable scenario where the DIC is known to be accurate and the accuracy of our estimate can be assessed in practice.

The nearby K0 V dwarf HD 3651 has been reported to be a host to a $0.20 M_{\text{Jup}}$ exoplanet with an orbital period of 62.23 ± 0.03 days and an orbital eccentricity of 0.63 ± 0.04 (Fischer et al., 2003). The radial velocity variations of HD 3651 have been observed using the HIRES at the Keck I telescope (Fischer et al., 2003; Butler et al., 2006) and the Shane and CAT telescopes at the Lick observatory (Fischer et al., 2003; Butler et al., 2006). These datasets contain measurements at 42 and 121 epochs, respectively. The reason we chose these data is that they enable us to investigate several scenarios reliably. The fact that the planet orbiting HD 3651 is on an eccentric orbit and there is plenty of data available make it possible to assess the accuracy of the TPM estimate in several scenarios by enabling the comparison to the DIC estimate that is accurate as long as the posterior density is Gaussian. Therefore, we investigate the accuracy and convergence properties of the TPM in various scenarios: with high and low numbers of data compared to the number of model parameters, and when the marginal likelihoods of two models are close to each other and as far from each other as possible given the available data.

We analyse the radial velocities of HD 3651 made using the HIRES and Lick exoplanet surveys and calculate the marginal likelihoods of models with 0 and 1 Keplerian signals using the methods based on DIC, AIC, TPM, HM, and OBMH. We denote these estimates of Eq. (2) as $\hat{P}_{\text{DIC}}$, $\hat{P}_{\text{AIC}}$, $\hat{P}_{\text{TPM}}$, $\hat{P}_{\text{HM}}$, and $\hat{P}_{\text{OBMH}}$, respectively.

3.1. Case 1. HIRES data

The HIRES data with 42 epochs reveals some interesting differences between the five estimates for marginal integrals. The log-marginal integrals are plotted in Fig. 1 as a function of Markov chain length. The estimated uncertainties of DIC and OBMH estimates represent the standard deviations of 6 different Markov chains. The DIC estimate can be considered a reliable one in this case, because the posterior density is very close to a multivariate Gaussian. It can be seen that the AIC is biased because of the low number of measurements (42) compared to the number of parameters of the statistical model (7). Also, the OBMH estimate gives the 1-planet model a greater marginal likelihood than DIC. However, the TPM is similarly biased for $\lambda = 0.5, 0.1, 10^{-2}, 10^{-3}$ but converges to the DIC estimate for $\lambda = 10^{-4}, 10^{-5}$. The HM estimate is not shown in the Fig. because its extremely poor convergence properties – it receives values between -130 and -140 on the logarithmic scale of Fig. 1.

When using as small values of λ for the TPM as possible such that it converges in the sense that it approaches some limiting value, we calculate the Bayes factors (B) in favour of the one-Keplerian model against the model without Keplerian signals. These values are shown in Table 1. The TPM estimate con-




Fig. 1. The marginal integrals of the 1-planet model given the HIRES data (case 1): DIC and its 3σ uncertainty (black dashed line and black dotted lines), AIC (blue dashed), OBMH and its 3σ uncertainty (red dashed and red dotted), and the TPM estimates with $\lambda = 0.5, 0.1, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}$ (black, grey, blue, purple, pink, and red curves).

Table 1. The Bayes factors in favour of the one-Keplerian model given the HIRES data (case 1).

Estimate	B
TPM	$1.1 \times 10^{14} \pm 1.2 \times 10^{13}$
DIC	$1.1 \times 10^{14} \pm 1.1 \times 10^{13}$
AIC	3.3×10^{15}
OBMH	$2.8 \times 10^{16} \pm 5.6 \times 10^{15}$
HM	$3.3 \times 10^{13} \pm 6.5 \times 10^{13}$

verges to the same value as DIC, which is known accurate in this case because the posterior densities of both models are very close to Gaussian. However, the AIC and OBMH overestimate the posterior probability of the model containing a Keplerian signal. Also, the problems of the HM estimate are clear because its uncertainty becomes greater than the estimated value.

3.2. Case 2. Combined HIRES and Lick data

Increasing the number of measurements likely makes the AIC yield a more accurate estimate for the marginal likelihood. However, to see how this affects the other estimates, we again compare them to the DIC which is reliable because of the close-Gaussianity of the posterior density. The inclusion of additional Lick data also makes the posterior probability of the one-Keplerian model much greater than that of the model without Keplerian signals, and enables us to investigate the accuracy and convergence of the TPM in such a scenario. Therefore, we study the properties of the different estimates for marginal integrals using the combined HIRES and Lick data of HD 3651 with 163 epochs.

The TPM converges to the DIC estimate when $\lambda = 10^{-3}$ for the model without any Keplerians, whereas its convergence takes place for $\lambda = 10^{-5}$ for the one-Keplerian model (Fig. 2, pink curve). Clearly, the AIC is indeed closer to the DIC estimate because of the greater number of data but the OBMH is also consistent with the DIC estimate. We note that the HM estimate is again omitted from the Fig. 2 because it receives significantly lower values than the other estimates.




Fig. 2. As in Fig. 1 but for the combined data (case 2) and the TPM estimates with $\lambda = 0.1, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}$ (black, grey, blue, purple, pink, and red curves).

Table 2. The Bayes factors in favour of the one-Keplerian model given the combined HIRES and Lick data (case 2).

Estimate	B
TPM	$2.0 \times 10^{38} \pm 1.0 \times 10^{37}$
DIC	$2.2 \times 10^{38} \pm 1.9 \times 10^{37}$
AIC	5.7×10^{38}
OBMH	$2.0 \times 10^{41} \pm 9.0 \times 10^{39}$
HM	$1.4 \times 10^{38} \pm 1.7 \times 10^{38}$

Now, we calculate the Bayes factors in favour of the one-Keplerian model and present them in Table 2. The TPM estimate is again very close to the DIC estimate and the AIC is close to these, providing slightly greater support for the one-Keplerian model. The OBMH again overestimates the one-Keplerian model and the HM estimate, while this time being rather accurate, has an uncertainty in excess of the estimate itself. Clearly, the TPM estimate can be used to receive reliable estimates for the marginal integral in this case as well, because the posterior density is again very close to Gaussian and the DIC estimate is therefore a reliable one in assessing the integral.

3.3. Case 3. Partial HIRES data

As a third test, we calculate the different estimates for marginal integral given only 20 epochs of HIRES data – the first 20 epochs between 366 and 2602 JD-2450000 – to see their relative performance when the number of parameters is comparable to the number of measurements. We find that the TPM converges to the marginal integral very accurately with $\lambda = 10^{-3}$ for both models and yields very reliable estimates for these integrals. It is again very close to the DIC estimate, making it reliable because of the Gaussianity of the posterior density for both models and the consequent reliability of the DIC estimate. It is not surprising that the AIC overestimates the Bayes factor and therefore also the posterior odds of the one-Keplerian model because of the low number of data. However, the OBMH overestimates it as well as was in fact found to be the case in test cases 1 and 2 as well.

Table 3. The Bayes factors in favour of the one-Keplerian model given the partial HIRES data (case 3).

Estimate	B
TPM	$3.0 \times 10^5 \pm 5.1 \times 10^4$
DIC	$2.8 \times 10^5 \pm 8.1 \times 10^4$
AIC	1.2×10^9
OBMH	$1.4 \times 10^6 \pm 3.3 \times 10^5$
HM	$4.3 \times 10^3 \pm 8.1 \times 10^3$

4. Conclusions

Calculating the marginal integral for model selection purposes is generally a challenging computational problem. While there are several good estimates for these integrals, they are usually only applicable under certain limiting assumptions about the nature of the posterior density, the amount of parameters in the statistical model or the number of measurements available. Therefore, we have introduced a new method for estimating these integrals in practice. Given the availability of a sample from the posterior density of model parameters, our truncated posterior mixture estimate is a reasonably accurate and very easily calculated in practice. We have only assumed that a statistically representative sample drawn from the posterior density exists when deriving our posterior mixture estimate (see Appendix). Therefore, it is applicable to any model comparison problems in astronomy and other fields of scientific inquiry and is not restricted to problems where the posterior has a certain shape and dimension.

The comparisons of different estimates given the radial velocities of HD 3651 revealed that the TPM yields estimates very close to the DIC estimate, which is known to be a reliable one in case of Gaussian posterior density. In fact, we chose the HD 3651 as an example star because of the planet orbiting it is known to have an eccentric orbit that enables the Gaussianity of the probability distributions of eccentricity and the two angular parameters of the Keplerian model, namely, longitude of pericentre and mean anomaly. However, the simple small-sample version of the AIC proved reasonably accurate as well when the number of measurements well exceeded the number of free parameters of the model (e.g. Table 2). We also note that the OBMH estimate of Chib & Jeliazkov (2001), while converging rapidly, tends to yield somewhat biased results that exaggerate the posterior probability of the more complicated model, making it possibly – at least in the test cases of the current work – prone to detections of false positives.

In practice, the TPM can be used by calculating its value directly from the sample drawn from the posterior density of the model parameters. Selecting a suitable value for parameter λ is then of essence when calculating its value in practice. In all the three different test cases studied in this article, a choice of $\lambda = 10^{-4}$ yielded estimates that converged rapidly for all the models in all the test cases and resulted in posterior probabilities that differed little from those calculated using the DIC estimate. When the difference between the two models was the smallest (case 3.), there was practically no bias in the TPM estimate with respect to the DIC. Also, when the posterior odds of the one-Keplerian model was the greatest (case 2.), the TPM, with $\lambda = 10^{-4}$, overestimated the posterior probability of the one-Keplerian model by a factor of 10, though, in that case, the Bayes factor used in model selection was already so heavily in favour of the one-Keplerian model that this overestimation is not significant in practice in terms being able to select the best model.

Because of these possible biases caused by too large λ , it would then be convenient in practice to calculate the TPM estimate using few different values of parameter λ . With the sample from the posterior density available, this could be done with little computational cost. Then, it would be possible to use the lowest value for λ that still converges to receive a trustworthy TPM estimate and correspondingly trustworthy model selection results.

Acknowledgements. M. Tuomi is supported by RoPACS (Rocky Planets Around Cools Stars), a Marie Curie Initial Training Network funded by the European Commission's Seventh Framework Programme.

Appendix A: Marginal integrals from importance sampling

In the context of Bayesian model selection, the marginal integral needed to assess the relative probabilities of different model is

$$P(m|\mathcal{M}) = \int_{\theta \in \Theta} l(m|\theta, \mathcal{M})\pi(\theta|\mathcal{M})d\theta, \quad (\text{A.1})$$

where $\mathcal{M}$ is a model with parameter vector θ constructed to model the measurements m using the likelihood function l . Function $\pi(\theta|\mathcal{M})$ is the prior probability density of the model parameters. This quantity is essential in calculating the posterior probabilities of different models in Eq. (1).

Importance sampling can be used to receive estimates for the integral in Eq. (A.1). Choosing functions g and w such that $\pi(\theta) = w(\theta)g(\theta)$ and dropping the model from the notation, the marginal integral can be written using the expectation with respect to density g as

$$\mathbb{E}_g[w(\theta)l(m|\theta)] = \int_{\theta \in \Theta} g(\theta)w(\theta)l(m|\theta)d\theta = P(m). \quad (\text{A.2})$$

where $g(\theta)$ is usually called the importance sampling function. Now, the idea of importance sampling is that if we draw a sample of N values from g and denote $\theta_i \sim g(\theta)$ for all $i = 1, \dots, N$, we can calculate a simple estimate for the marginal integral as (e.g. Kass & Raftery, 1995)

$$\hat{P} = \left[\sum_{i=1}^N \frac{\pi(\theta_i)l(m|\theta_i)}{g(\theta_i)} \right] \left[\sum_{i=1}^N \frac{\pi(\theta_i)}{g(\theta_i)} \right]^{-1}. \quad (\text{A.3})$$

All there remains is to choose g such that it is easy to draw a sample from it and that the estimate in Eq. (A.3) converges rapidly to the marginal integral.

Some simple choices of g would be the prior density or the posterior density. In these cases, the resulting estimates would be called the mean estimate and the harmonic mean estimate (Newton & Raftery, 1994; Kass & Raftery, 1995). We denote these estimates as $\hat{P}_M$ and $\hat{P}_{HM}$ and write

$$\hat{P}_M = \frac{1}{N} \sum_{i=1}^N l(m|\theta_i) \quad (\text{A.4})$$

and

$$\hat{P}_{HM} = N \left[\sum_{i=1}^N \frac{1}{l(m|\theta_i)} \right]^{-1}. \quad (\text{A.5})$$

Though easily computed in practice, these estimates have some undesirable properties. For instance, the mean estimate requires drawing a sample from the prior density and computation of the corresponding likelihoods. However, because the prior

contains less information and is therefore much broader density than the posterior, most of the values in this sample correspond to very low likelihoods and the convergence of this estimate is slow. The resulting value is also dominated by few high likelihoods, which can make it too biased to be useful in applications.

Also, while converging to the desired value, the harmonic mean estimate does it extremely slowly in practice (Kass & Raftery, 1995) and its usage cannot be recommended. In applications, this estimate doesn't generally converge to the marginal integral within the limited sample available from the posterior. The reason is that occasional small values of $l(m|\theta_i)$ have a large impact on the sum which makes its convergence extremely slow. For these reasons, better estimates are needed to approximate the marginal integrals in model selection problems.

A.1. The posterior mixture estimate

To construct a better estimate for the marginal integral, we start by assuming that a statistically representative sample has been drawn from the posterior density using some posterior sampling algorithm. Therefore, we have a collection of N vectors $\theta_i \sim \pi(\theta|m)$, for all $i = 1, \dots, N$. These values form a Markovian chain with N members. Selecting integer $h > 0$, the value of the posterior $\pi(\theta_{i-h}|m)$ is available if the value corresponding to θ_i is available given $i > h > 0$. Here we can denote $\pi_i = \pi(\theta_i|m)$ and see that if θ_i is a random vector then π_i is some random number corresponding to the value of the posterior at θ_i . Using the notation similarly for g_i , and setting $\lambda \in [0, 1]$, we can set

$$g_i = (1 - \lambda)\pi_i + \lambda\pi_{i-h}. \quad (\text{A.6})$$

Now, if λ is a small number, it follows that $g_i \approx \pi_i$ – the importance sampling function g is close to the posterior but not exactly equal. We call it a truncated posterior mixture (TPM) function. The sample from the posterior is close to a sample from g – a desired property because a sample from the posterior can be calculated rather readily with posterior sampling algorithms (e.g. Metropolis et al., 1953; Hastings, 1970; Haario et al., 2001). Therefore, assuming that the parameter space Θ is a bounded subset of $\mathbb{R}^K$ and that the prior is a proper constant prior – actually, it is sufficient to assume that the prior is a broad density that has almost the same value in the set $\mathcal{D}_\delta$ of the posterior – the estimate in Eq. (A.3) can now be calculated. We denote $l_i = l(m|\theta_i)$ and write the resulting posterior mixture estimate as

$$\hat{P}_{TPM} = \left[\sum_{i=1}^N \frac{l_i}{(1 - \lambda)l_i + \lambda l_{i-h}} \right] \left[\sum_{i=1}^N \frac{1}{(1 - \lambda)l_i + \lambda l_{i-h}} \right]^{-1}. \quad (\text{A.7})$$

If the Markov chain has good mixing properties such that the value θ_i has already become independent of θ_{i-h} , the likelihoods of these values are also independent. When comparing this estimate with $\hat{P}_{HM}$ in Eq. (A.5), it can be seen that occasional small values of l_i do not have such a large effect on the sum in the denominator because it is unlikely that the corresponding value of l_{i-h} is also small at the same time.

References

- Akaike, H. 1973. pp. 267 in Petrov, B. N. & Csaki, F. (eds.). Second International Symposium on Information Theory. Akadémiai Kiadó.
- 1
- Barge, P., Baglin, A., Auvergne, M., et al. 2007, A&A, 482, L17
- Bean, J. L., Seifahrt, A., Hartman, H., et al. 2010, The ApJ, 711, L19
- Benedict, G. F., McArthur, B. E., Forveille, T., et al. 2002, ApJ, 581, L115

- Burnham, K. P. & Anderson D. R. 2002, Model selection and multimodel inference: A practical information-theoretic approach, Springer-Verlag
- Butler, R. P., Wright, J. T., Marcy, G. W., et al. 2006, ApJ, 646, 505
- Chib S. & Jeliazkov I. 2001, J. Am. Stat. Ass., 96, 270
- Clyde, M. A., Berger, J. O., Bullard, F., et al. 2007, Statistical Challenges in Modern Astronomy IV, Babu, G. J. & Feigelson, E. D. (eds.), ASP Conf. Ser., 371, 224
- Collier Cameron, A., Bouchy, F., Hébrard, G., et al. 2007, MNRAS, 375, 951
- Faedi, F., Barros, S. C. C., Anderson, D. R., et al. 2011, A&A, 531, A40
- Fischer, D. A., Butler, R. P., Marcy, G. W., et al. 2003, ApJ, 590, 1081
- Ford, E. B. & Gregory, P. C. 2007, Statistical Challenges in Modern Astronomy IV, Babu, G. J. & Feigelson, E. D. (eds.), ASP Conf. Ser., 371, 189
- Forveille, T., Bonfils, X., Delfosse, X., et al. 2011, A&A, submitted (arXiv:1109.2505 [astro-ph.EP])
- Gregory, P. C. 2005, ApJ, 631, 1198
- Gregory, P. C. 2007a, MNRAS, 381, 1607
- Gregory, P. C. 2007b, MNRAS, 374, 1321
- Haario, H., Saksman, E., & Tamminen, J. 2001, Bernoulli, 7, 223
- Hastings, W. 1970, Biometrika 57, 97
- Hébrard, G., Evans, T. M., Alonso, R., et al. 2011. A&A, in press (arXiv:1107.2032 [astro-ph.EP])
- Holman, M. J. & Murray, N. W. 2005, Science, 307, 1288
- Jones, H. R. A., Butler, R. P., Marcy, G. W., et al. 2002, MNRAS, 337, 1170
- Kass, R. E. & Raftery, A. E. 1995, J. Am. Stat. Ass., 430, 773
- Kullback, S. & Leibler, R. A. 1951, Ann. Math. Stat., 22, 76
- Lovis, C., Ségransan, D., Mayor, M., et al. 2011, A&A, 528, A112
- Mayor, M., Pepe, F., Queloz, D., et al. 2003, Messenger, 114, 20
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., et al. 1953, J. Chem. Phys., 21, 1087
- Newton, M. A. & Raftery, A. E. 1994, J. Roy. Stat. Soc. B, 56, 3
- Pravdo, S. H. & Shaklan, S. B. 2009, ApJ, 700, 623
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. 2002, JRSS B, 64, 583
- Tinney, C. G., Butler, R. P., Marcy, G. W., et al. 2001, ApJ, 551, 507
- Tuomi, M. & Kotiranta, S. 2009, A&A, 496, L13
- Tuomi, M. 2011, A&A, 528, L5
- Tuomi, M., Pinfield, D., & Jones, H. R. A. 2011, A&A, 532, A116
- Vogt, S. S., Allen, S. L., Bigelow, B. C., et al. 1994, SPIE Instrumentation in Astronomy VIII, Crawford, D. L., Craine, E. R. (eds.), 2198, p. 362
- Vogt, S. S., Butler, R. P., Rivera, E. J., et al. 2010, ApJ, 723, 954