

Learned harmonic mean estimation of the marginal likelihood with normalizing flows

Alicja Polanska^{1,*}, Matthew A. Price¹, Alessio Spurio Mancini¹, and Jason D. McEwen^{1,2}

¹ Mullard Space Science Laboratory, University College London, Dorking, RH5 6NT, UK;

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

* Correspondence: alicja.polanska.22@ucl.ac.uk;

Abstract: Computing the marginal likelihood (also called the Bayesian model evidence) is an important task in Bayesian model selection, providing a principled quantitative way to compare models. The learned harmonic mean estimator solves the exploding variance problem of the original harmonic mean estimation of the marginal likelihood. The learned harmonic mean estimator learns an importance sampling target distribution that approximates the optimal distribution. While the approximation need not be highly accurate, it is critical that the probability mass of the learned distribution is contained within the posterior in order to avoid the exploding variance problem. In previous work a bespoke optimization problem is introduced when training models in order to ensure this property is satisfied. In the current article we introduce the use of normalizing flows to represent the importance sampling target distribution. A flow-based model is trained on samples from the posterior by maximum likelihood estimation. Then, the probability density of the flow is concentrated by lowering the variance of the base distribution, i.e. by lowering its “temperature”, ensuring its probability mass is contained within the posterior. This approach avoids the need for a bespoke optimisation problem and careful fine tuning of parameters, resulting in a more robust method. Moreover, the use of normalizing flows has the potential to scale to high dimensional settings. We present preliminary experiments demonstrating the effectiveness of the use of flows for the learned harmonic mean estimator. The `harmonic` code implementing the learned harmonic mean, which is publicly available, has been updated to now support normalizing flows.

Keywords: Bayesian model selection; harmonic mean estimator; normalizing flows.

1. Introduction

Model selection is the task of evaluating which of the statistical models under consideration best describes observed data. In the Bayesian formalism this involves computing the marginal likelihood (also called the Bayesian model evidence), which gives a way to quantitatively compare the suitability of models for a given problem. Model selection is relevant in a range of fields such as astronomy, biostatistics, economics, medical research and many more. However, in practice Bayesian model selection is often difficult as computing the marginal likelihood requires evaluating a high-dimensional integral, which can be a challenging task. A number of methods to compute the marginal likelihood have been proposed (for reviews see [1,2]), such as nested sampling [3,4].

The learned harmonic mean estimator was proposed recently by some of the authors of the current article as an effective technique to compute the marginal likelihood [5]. The estimator requires only samples from the posterior and so is agnostic to the method used to generate samples, in contrast to nested sampling. Thus, the learned harmonic mean can be easily coupled with various MCMC sampling techniques or variational inference approaches. This property also allows the estimator to be adapted to address Bayesian model selection for simulation-based inference (SBI) [6], where an explicit likelihood is unavailable or infeasible. The learned harmonic mean estimator is implemented in the `harmonic` Python code¹.

¹ <https://github.com/astro-informatics/harmonic>



Citation: Polanska, A.; Price, M. A.; Spurio Mancini, A.; McEwen, J. D. Learned harmonic mean estimation of the marginal likelihood with normalizing flows. *Preprints* 2023, 1, 0. <https://doi.org/>



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

In this work we introduce the use of normalizing flows [7] for the learned harmonic mean estimator, which addresses some limitations of the models used previously. The use of normalizing flows eliminates the need for bespoke training, resulting in a more robust and scalable approach, which is now implemented in the `harmonic` code. We first review the learned harmonic mean estimator, before describing how normalizing flows may be used for the estimator and their main advantages in this context. We then present a number of experiments that demonstrate the effectiveness of the use of flows in the learned harmonic mean estimator.

2. The harmonic mean estimator

Bayesian model selection requires the computation of the marginal likelihood given by

$$z = p(y | M) = \int d\theta p(y | \theta, M) p(\theta | M) = \int d\theta \mathcal{L}(\theta) \pi(\theta), \quad (1)$$

where y denotes observed data, θ the parameters of interest, and M the model under consideration. We adopt the shorthand notation for the likelihood of $\mathcal{L}(\theta) = p(y | \theta, M)$ and prior of $\pi(\theta) = p(\theta | M)$.

The harmonic mean estimator was first proposed by [8], who showed that the marginal likelihood z can be estimated from the harmonic mean of the likelihood, given posterior samples. This follows by considering the expectation of the reciprocal of the likelihood with respect to the posterior distribution, leading to the following estimator of the reciprocal of the marginal likelihood $\rho = 1/z$:

$$\hat{\rho} = \frac{1}{N} \sum_{i=1}^N \frac{1}{\mathcal{L}(\theta_i)}, \quad \theta_i \sim p(\theta | y), \quad (2)$$

where N specifies the number of samples θ_i drawn from the posterior $p(\theta | y)$. The marginal likelihood can then be estimated from its reciprocal straightforwardly [5]. It was soon realized that the original harmonic mean estimator can fail catastrophically [9], as it can suffer from an exploding variance.

The estimator can also be interpreted as importance sampling. Consider the reciprocal marginal likelihood, which may be expressed in terms of the prior and posterior as:

$$\rho = \int d\theta \frac{1}{\mathcal{L}(\theta)} p(\theta | y) = \int d\theta \frac{1}{z} \frac{\pi(\theta)}{p(\theta | y)} p(\theta | y). \quad (3)$$

It is clear the estimator has an importance sampling interpretation where the importance sampling target distribution is the prior $\pi(\theta)$, while the sampling density is the posterior $p(\theta | y)$, in contrast to typical importance sampling scenarios.

For importance sampling to be effective, one requires the sampling density to have fatter tails than the target distribution, i.e. to have greater probability mass in the tails of the distribution. Typically the prior has fatter tails than the posterior since the posterior updates our initial understanding of the underlying parameters θ that are encoded in the prior, in the presence of new data y . For the harmonic mean estimator the importance sampling density (the posterior) typically does not have fatter tails than the target (the prior) and so importance sampling is not effective. This explains why the original harmonic mean estimator can be problematic.

In [10] an arbitrary density $\varphi(\theta)$ is introduced to relate the reciprocal of the marginal likelihood to the likelihood through the following expectation:

$$\rho = \mathbb{E}_{p(\theta | y)} \left[\frac{\varphi(\theta)}{\mathcal{L}(\theta) \pi(\theta)} \right]. \quad (4)$$

The above expression motivates the estimator:

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

The normalised density $\varphi(\theta)$ can be interpreted as an alternative importance sampling target distribution, hence we refer to this approach as the re-targeted harmonic mean estimator. Note that the original harmonic mean estimator is recovered for the target distribution $\varphi(\theta) = \pi(\theta)$.

The learned harmonic mean estimator is introduced in [5], where the target density $\varphi(\theta)$ is learned by machine learning techniques. It is shown in [5] that the optimal target distribution is the posterior. Since the target must be normalized, the normalized posterior is clearly not accessible since its normalizing constant is precisely the term of interest. The learned harmonic mean approximates the optimal target of the posterior with a learned model that is normalized. While the approximation need not be highly accurate, it is critical that the probability mass of the learned distribution is contained within the posterior in order to avoid the exploding variance problem. In [5] a bespoke optimization problem is introduced when training models in order to ensure this property is satisfied. Specifically, the model is fitted by minimizing the variance of the resulting estimator, while ensuring it is also unbiased, and with possible regularization. Such an approach requires a careful selection of an appropriate model and its hyperparameters for a problem at hand, determined by cross-validation. Furthermore, only simple classical machine learning models were considered in [5], which in many cases struggle to scale to high-dimensional settings.

3. Learning the target distribution using normalizing flows

In this paper we learn the target distribution of the learned harmonic mean estimator [5] using normalizing flows. Using normalizing flows renders the previous bespoke approach to training no longer necessary since it provides an elegant way to ensure the probability mass of the learned distribution is contained within the posterior, thereby resulting in a learned harmonic mean estimator that is more flexible and robust. Furthermore, normalizing flows also offer the potential to scale to higher dimensional settings. We first introduce normalizing flows, before describing how they may be used for the learned harmonic mean estimator and their main advantages in this context.

3.1. Normalizing flows

Normalizing flows are a class of probabilistic models that allow one to evaluate the density of and sample from a learned probability distribution (for a review see [7]). They consist of a series of transformations that are applied to a simple base distribution. A vector θ of an unknown distribution $p(\theta)$, can be expressed through a transformation T of a vector z sampled from a base distribution $q(z)$:

$$\theta = T(z), \text{ where } z \sim q(z). \quad (6)$$

Typically the base distribution is chosen so that its density can be evaluated simply and that it can be sampled from easily. Often a Gaussian is used for the base distribution. The unknown distribution can then be recovered by the change of variables formula:

$$p(\theta) = q(z) |\det J_T(z)|^{-1}, \quad (7)$$

where $J_T(z)$ is the Jacobian corresponding to transformation T . In a flow-based model T consists of a series of learned transformations that are each invertible and differentiable, so that the full transformation is also invertible and differentiable. This allows us to compose multiple simple transformations with learned parameters, into what is called a flow, obtaining a normalized approximation of the unknown distribution that we can




Figure 1. Diagram illustrating the concentration of the probability density of a normalizing flow. The flow is trained on samples from the posterior, giving us a normalized approximation of the posterior distribution. The temperature of the base distribution $T \in (0, 1)$ is reduced, which concentrates the probability density of the transformed distribution ensuring that it is contained within the posterior. The concentrated flow can then be used as the target distribution for the learned harmonic mean estimator, avoiding the exploding variance issue of the original harmonic mean estimator.

sample from and evaluate. Careful attention is given to construction of the transformations such that the determinant of the Jacobian can be computed easily.

A relatively simple example of a normalizing flow is the real-valued non-volume preserving (real NVP) flow introduced in [11]. It consists of a series of bijective transformations given by affine coupling layers. Consider the D dimensional input z , split into elements up to and following d , respectively, $z_{1:d}$ and $z_{d+1:D}$, for $d < D$. Given input z , the output y of an affine couple layer is calculated by

$$y_{1:d} = z_{1:d}; \quad (8)$$

$$y_{d+1:D} = z_{d+1:D} \odot \exp(s(z_{1:d})) + t(z_{1:d}), \quad (9)$$

where \odot denotes Hadamard (elementwise) multiplication. The scale s and translation t are typically represented by neural networks with learnable parameters that take as input $z_{1:d}$. This construction is easily invertible and ensures the Jacobian is a lower-triangular matrix, making its determinant efficient to calculate.

3.2. Concentrating the probability density for the learned harmonic mean estimator

Normalizing flows meet the core requirements of the learned target distribution of the learned harmonic mean estimator: namely, they provide a normalized probability distribution for which one can evaluate probability densities. In this work we use them to introduce an elegant way to ensure the probability mass of the learned distribution is contained within the posterior. We thereby avoid the exploding variance issue of the original harmonic mean estimator and can evaluate the marginal likelihood accurately without the need for fine-tuning.

Reducing the variance of the base distribution, or equivalently lowering its “temperature” in a statistical mechanics perspective, clearly concentrates the probability density of the base distribution. This has the effect of also concentrating the probability density of the transformed distribution due to the continuity and differentiability of the flow. Consequently, once a flow is trained to approximate the posterior, by lowering the temperature of the base distribution (i.e. reducing its variance) we can concentrate the learned distribution to ensure its probability mass is contained within the posterior, as illustrated in Figure 1.

The learned distributions considered previously for the learned harmonic mean estimator [5] required the introduction of a bespoke optimization problem for training in




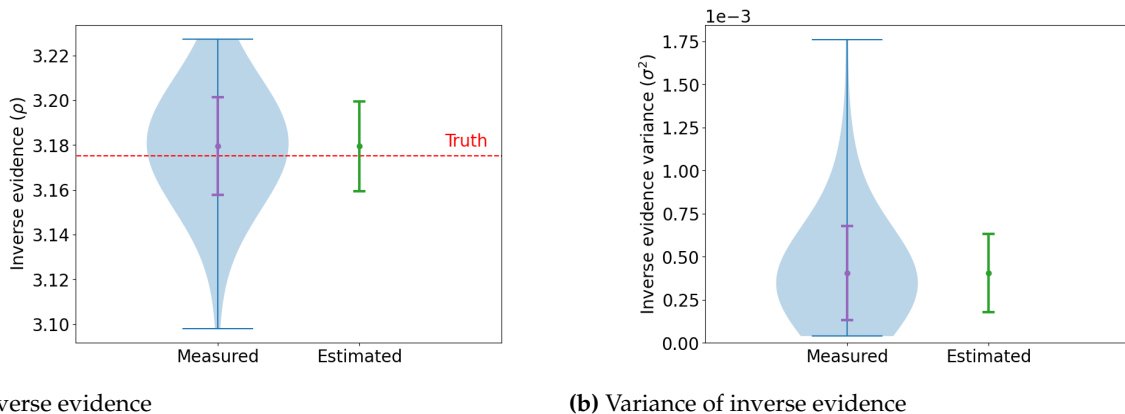
Figure 2. Corner plot of samples from the posterior (red) and real NVP flow with temperature $T = 0.9$ (blue) for the Rosenbrock benchmark problem. The target distribution given by the concentrated flow is contained within the posterior and has thinner tails, as required for the learned harmonic mean estimator.

order to ensure the learned target is contained within the posterior. This requires careful selection of an appropriate model and its hyperparameters, determined by cross-validation. The introduced normalizing flow approach renders bespoke training no longer necessary. Instead, we train a flow in the usual manner, based on maximum likelihood estimation, before concentrating its probability density. There is only one parameter to consider, the temperature $T \in (0, 1)$. Moreover, we expect a common value of $T \sim 0.9$ to be suitable for most problems. Once we have a flow with its probability density concentrated for $\varphi(\theta)$, the learned harmonic mean estimator can be computed in the usual manner [5]. Using normalizing flows with the learned harmonic mean thus provides a much more robust method. Furthermore, an added benefit of using flows is that we can draw samples from the flow distribution efficiently, in order to easily visualize the concentrated target distribution and compare it to the posterior.

In this preliminary work we consider real NVP flows only, as described above, which are implemented in the harmonic code. In the future we will consider more expressive and scalable flows. The use of normalizing flows for the learned distribution therefore has the potential to extend the learned harmonic mean estimator to problems with complex, high-dimensional posteriors.

4. Experiments

To validate the effectiveness of the method described in Section 3, we repeat a number of the numerical experiments carried out in [5] but using normalizing flows as the target distribution of the learned harmonic mean estimator. In the experiments that follow we consider a real NVP flow where the scale and translation networks of the affine coupling layers are given by two layer dense neural networks with a leaky ReLU in between. The scaling layers additionally include a preceding softplus activation. We typically consider a flow with six coupling layers, where typically only the first two include scaling, and permute elements of the vector between coupling layers to ensure the flow transforms all elements. We consider a Gaussian base distribution with unit variance. We use the `emcee` package [12] to generate MCMC samples from the posterior. We then train the real NVP flow on half of the samples by maximum likelihood and calculate the marginal likelihood using the remaining samples by the learned harmonic mean estimator with the flow concentrated to temperature T . In all experiments we consider an identical temperature of $T = 0.9$, which works well throughout, demonstrating that T does not require fine-tuning. We



(a) Inverse evidence

(b) Variance of inverse evidence

Figure 3. Marginal likelihood computed by the learned harmonic mean estimator with a concentrated flow for the Rosenbrock benchmark problem. 100 experiments are repeated to recover empirical estimates of the statistics of the estimator. In panel (a) the distribution of marginal likelihood values are shown (measured) along with the estimate of the standard deviation computed by the error estimator (estimated). The ground truth is indicated by the red dashed line. In panel (b) the distribution of the variance estimator is shown (estimated) along with the standard deviation computed by the variance-of-variance estimator (estimated). The learned harmonic mean estimator and its error estimators are highly accurate.

consider a relatively simple flow in this preliminary work and a small number of simple experiments. In future work we will consider more expressible and scalable flows, and further experiments to thoroughly evaluate the robustness and scalability of the method.

4.1. Rosenbrock

A common benchmark problem to test methods that compute the marginal likelihood is a likelihood specified by the Rosenbrock function, which exhibits a narrow curving degeneracy. We consider the Rosenbrock likelihood in $d = 2$ dimensions and a simple uniform prior with $x_0 \in [-10, 10]$ and $x_1 \in [-5, 15]$. We sample the resulting posterior distribution, drawing 5,000 samples for 200 chains, with burn in of 2,000 samples, yielding 3,000 posterior samples per chain. Figure 2 shows a corner plot of the training samples from the posterior (red) and from the normalizing flow (blue) at temperature $T = 0.9$. It can be seen that the concentrated flow approximates the posterior well and has thinner tails, as required for the marginal likelihood estimate to be stable and accurate.

This process is repeated 100 times and the marginal likelihood is computed for each trial. Figure 3 shows a summary of the estimates across all the runs. The dashed red line in Figure 3a indicates the ground truth computed through numerical integration, which is tractable in two dimensions. It can be seen that the learned harmonic mean estimator using a real NVP flow provides an accurate and unbiased estimate of the marginal likelihood.

4.2. Normal-Gamma

We consider the Normal-Gamma example for which the marginal likelihood can be computed analytically [2,5]. It was found that the marginal likelihood values computed by the original harmonic mean estimator do not vary with a varying prior [2], highlighting this example as a pathological failure of the original harmonic mean estimator. We consider the same pathological example here and demonstrate that our learned harmonic mean estimator with normalizing flows is highly accurate (as is the learned harmonic mean with other models; [5]). We consider the Normal-Gamma model [5,13] with data $y_i \sim N(\mu, \tau^{-1})$, for $i \in \{1, \dots, n\}$, with mean μ and precision (inverse variance) τ . A normal prior is assumed for μ and a Gamma prior for τ :

$$\mu \sim N(\mu_0, (\tau_0 \tau)^{-1}), \quad \tau \sim \text{Ga}(a_0, b_0), \quad (10)$$




Figure 4. Corner plot of samples from the posterior (red) and real NVP flow trained on the posterior samples with temperature $T = 0.9$ (blue) for the Normal-Gamma example with $\tau = 0.001$. The target distribution given by the concentrated flow is contained within the posterior and has thinner tails, as required for the learned harmonic mean estimator.

with mean $\mu_0 = 0$, shape $a_0 = 10^{-3}$ and rate $b_0 = 10^{-3}$. The precision scale factor τ_0 is varied to observe the impact of changing prior on the computed marginal likelihood.

We draw 1,500 samples for 200 chains, with burn in of 500 samples, yielding 1,000 posterior samples per chain. Figure 4 shows a corner plot of the training samples from the posterior for $\tau = 0.001$ (red) and from the normalizing flow (blue) at temperature $T = 0.9$. Again, it can be seen the concentrated learned target is close to the posterior but with thinner tails, as expected. We consider priors with $\tau \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$. Figure 5 shows the relative accuracy of the marginal likelihood computed by the learned harmonic mean estimator using normalizing flows, that is the ratio of the estimated marginal likelihood to the analytic ground truth. We additionally consider a concentrated flow with $T = 0.95$ to demonstrate that accuracy is not highly dependent on the temperature parameter. It can be seen that the estimate remains accurate and is indeed sensitive to the prior for both temperatures. The estimates for the flow with T closer to one have a slightly lower variance, as one would expect, since the broader target φ makes more efficient use of samples.

4.3. Logistic regression models: Pima Indian example

We consider the comparison of two logistic regression models using the Pima Indians data, which is another common benchmark problem for comparing estimators of the marginal likelihood. The original harmonic mean estimator has been shown to fail catastrophically for this example [2]. The Pima Indians data [14], originally from the National Institute of Diabetes and Digestive and Kidney Diseases, were compiled from a study of indicators of diabetes in $n = 532$ Pima Indian women aged 21 or over. Seven primary predictors of diabetes were recorded, including: number of prior pregnancies (NP); plasma glucose concentration (PGC); diastolic blood pressure (BP); triceps skin fold thickness (TST); body mass index (BMI); diabetes pedigree function (DP); and age (AGE). The probability of diabetes p_i for person $i \in \{1, \dots, n\}$ can be modelled by the logistic function. An independent multivariate Gaussian prior with precision $\tau = 0.01$ is assumed for parameters θ . Two different logistic regression models are compared, with different subsets of covariates:

Model M_1 : covariates = {NP, PGC, BMI, DP} (and bias);

Model M_2 : covariates = {NP, PGC, BMI, DP, AGE} (and bias).




Figure 5. Ratio of marginal likelihood values computed by the learned harmonic mean estimator with a concentrated flow to those computed analytically for the Normal-Gamma problem. Errors bars corresponding to the estimated standard deviation of the learned harmonic estimator are also shown. Notice that the marginal likelihood values computed by the learned harmonic mean estimator are highly accurate and are indeed sensitive to changes in the prior. Predictions made with flow at temperature $T = 0.9$ (blue) and $T = 0.95$ (green) are shown, which are slightly offset for ease of visualization, demonstrating accuracy is not highly sensitive to the choice of T .

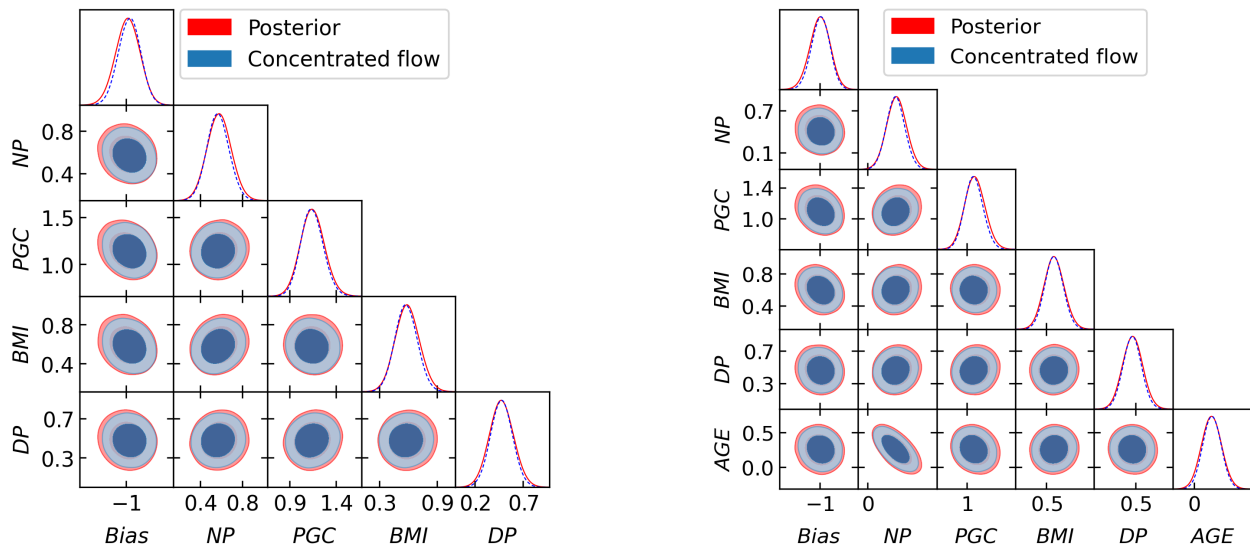
A reversible jump algorithm [15] is used by [2] to compute a benchmark Bayes factor BF_{12} of 13.96 ($\log \text{BF}_{12} = 2.6362$) which is treated as ground truth.

We draw 5,000 samples from for 200 chains, with burn in of 1,000 samples, yielding 4,000 posterior samples per chain. We train a flow consisting of 6 scaled layers followed by 2 unscaled ones. Figure 6 shows a corner plot of the training samples from the posterior (red) and from the normalizing flow (blue) at temperature $T = 0.9$. Again, it can be seen the concentrated learned target is close to the posterior but with thinner tails, as expected. We compute the marginal likelihood for Model 1 and Model 2 using our learned harmonic mean estimator. The log evidence found for Model 1 and 2 is -257.2300 ± 0.0020 and -259.8602 ± 0.0031 respectively, resulting in the estimate $\log \text{BF}_{12} = 2.6302 \pm 0.0051$, which is in close agreement with the benchmark.

5. Conclusions

In this work we propose using normalizing flows for the learned harmonic mean estimator of the marginal likelihood. The flow may be fitted to posterior samples by the usual maximum likelihood estimation. Its probability density may then be concentrated by lowering the temperature of the base distribution, ensuring the probability mass of the transformed distribution is contained within the posterior to avoid the exploding variance issue of the original harmonic mean estimator. The use of flows therefore results in a more robust learned harmonic mean estimator. We perform a number of experiments to compute the marginal likelihood with the proposed approach, using a real NVP flow, finding excellent agreement with ground truth values. In this preliminary work we consider only simple real NVP flows and a simple set of experiments. In a follow-up article we will consider more expressive and scalable flows to address problems with complex, high-dimensional posteriors. We will also perform a more extensive set of numerical experiments to thoroughly assess performance. This preliminary work nevertheless suggests the learned harmonic mean estimator with normalising flows provides an effective technique to compute the marginal likelihood for Bayesian model selection. Furthermore, it is applicable for any MCMC sampling technique or variational inference approach.

Author Contributions: Conceptualization, A.P., M.A.P. and J.D.M.; methodology, A.P., M.A.P. and J.D.M.; software, A.P., M.A.P. and J.D.M.; validation, A.P., M.A.P. and J.D.M.; investigation, A.P.; resources, J.D.M.; writing—original draft preparation, A.P. and J.D.M.; writing—review and editing, A.P., M.A.P., A.S.M. and J.D.M.; visualization, A.P.; supervision, M.A.P., A.S.M. and J.D.M.; funding acquisition, J.D.M. All authors have read and agreed to the published version of the manuscript.



(a) Model 1

(b) Model 2

Figure 6. Corner plot of the samples from the posterior (red) and real NVP flow trained on the posterior samples with temperature $T = 0.9$ (blue) for the Pima Indian benchmark problem for $\tau = 0.01$. The target distribution given by the concentrated flow is contained within the posterior and has thinner tails, as required for the learned harmonic mean estimator.

Funding: A.P. is supported by the UCL Centre for Doctoral Training in Data Intensive Science (STFC grant number ST/W00674X/1). A.S.M. acknowledges support from the MSSS STFC Consolidated Grant ST/W001136/1. This research was also funded by EPSRC grant number EP/W007673/1.

Data Availability Statement: The harmonic code and experiments are available at <https://github.com/astro-informatics/harmonic>.

References

- Clyde, M.; Berger, J.; Bullard, F.; Ford, E.; Jefferys, W.; Luo, R.; Paulo, R.; Lored, T. Current challenges in Bayesian model choice. In Proceedings of the Statistical challenges in modern astronomy IV, 2007, Vol. 371, p. 224.
- Friel, N.; Wyse, J. Estimating the evidence—a review. *Statistica Neerlandica* **2012**, *66*, 288–308.
- Skilling, J. Nested sampling for general Bayesian computation. *Bayesian Analysis* **2006**, *1*, 833–859.
- Ashton, G.; Bernstein, N.; Buchner, J.; Chen, X.; Csányi, G.; Fowlie, A.; Feroz, F.; Griffiths, M.; Handley, W.; Habeck, M.; et al. Nested sampling for physical scientists. *Nature Reviews Methods Primers* **2022**, *2*, 39.
- McEwen, J.D.; Wallis, C.G.; Price, M.A.; Docherty, M.M. Machine learning assisted Bayesian model comparison: learnt harmonic mean estimator. *Statistics & Computing*, submitted, **2021**, [arXiv:2111.12720].
- Spurio Mancini, A.; Docherty, M.M.; Price, M.A.; McEwen, J.D. Bayesian model comparison for simulation-based inference. *RASTI*, submitted **2022**, [arXiv:2207.04037].
- Papamakarios, G.; Nalisnick, E.; Rezende, D.J.; Mohamed, S.; Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research* **2021**, *22*, 2617–2680.
- Newton, M.A.; Raftery, A.E. Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society: Series B (Methodological)* **1994**, *56*, 3–26.
- Neal, R.M. Contribution to the discussion of “Approximate Bayesian inference with the weighted likelihood bootstrap” by Newton MA, Raftery AE. *JR Stat Soc Ser A (Methodological)* **1994**, *56*, 41–42.
- Gelfand, A.E.; Dey, D.K. Bayesian model choice: asymptotics and exact calculations. *Journal of the Royal Statistical Society: Series B (Methodological)* **1994**, *56*, 501–514.
- Dinh, L.; Sohl-Dickstein, J.; Bengio, S. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803* **2016**.
- Foreman-Mackey, D.; Hogg, D.W.; Lang, D.; Goodman, J. emcee: The MCMC Hammer. *PASP* **2013**, *125*, 306–312, [arXiv:1202.3665].
- Bernardo, J.; Smith, A. Bayesian theory. *New York* **1994**.
- Smith, J.W.; Everhart, J.E.; Dickson, W.; Knowler, W.C.; Johannes, R.S. Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In Proceedings of the Proceedings of the annual symposium on computer application in medical care. American Medical Informatics Association, 1988, p. 261.
- Green, P.J. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **1995**, *82*, 711–732.