
Improving Gradient-Guided Nested Sampling for Posterior Inference

Pablo Lemos^{1 2 3 4 5} Nikolay Malkin^{1 3} Will Handley^{6 7 8} Yoshua Bengio^{1 3 9} Yashar Hezaveh^{1 2 3 4 10 11}
Laurence Perreault-Levasseur^{1 2 3 4 10 11}

Abstract

We present a performant, general-purpose gradient-guided nested sampling algorithm, GGNS, combining the state of the art in differentiable programming, Hamiltonian slice sampling, clustering, mode separation, dynamic nested sampling, and parallelization. This unique combination allows GGNS to scale well with dimensionality and perform competitively on a variety of synthetic and real-world problems. We also show the potential of combining nested sampling with generative flow networks to obtain large amounts of high-quality samples from the posterior distribution. This combination leads to faster mode discovery and more accurate estimates of the partition function.

1. Introduction

Bayesian parameter estimation and model comparison are key to most scientific disciplines and remain challenging problems, especially in high-dimensional and multimodal settings. While traditionally Markov chain Monte Carlo (MCMC) methods have been used to perform Bayesian inference, differentiable programming has enabled the development of new, more efficient algorithms, such as variational inference (MacKay, 2003), Hamiltonian Monte Carlo (Duane et al., 1987; Neal et al., 2011), and Langevin dynamics (Besag, 1994; Roberts & Tweedie, 1996; Roberts & Rosenthal, 1998); as well as more recent learning-based methods such as the Path Integral Sampler (Zhang & Chen, 2022) and generative flow networks (Bengio et al., 2023; Lahlou et al., 2023).

¹Mila – Québec Artificial Intelligence Institute ²Ciela Institute
³Université de Montréal ⁴CCA – Flatiron Institute ⁵Dreamfold
⁶Cavendish Laboratory ⁷Kavli Institute for Cosmology ⁸Gonville and Caius College, University of Cambridge ⁹Canadian Institute for Advanced Research (CIFAR) ¹⁰Perimeter Institute
¹¹Trottier Space Institute. Correspondence to: Pablo Lemos <pablo@dreamfold.ai>.

From the perspective of differentiable programming, less attention has been paid in recent years to nested sampling (Skilling, 2006; Buchner, 2021; Ashton et al., 2022), which is a widely used algorithm for Bayesian parameter inference and model comparison. Nested sampling has been used in a range of applications in the natural sciences, from cosmology (Mukherjee et al., 2006; Handley et al., 2015a) and astrophysics (Lavie et al., 2017; Günther & Daylan, 2021) to particle physics (Yallup et al., 2022) and biology (Russel et al., 2018). Furthermore, it provides both samples from the posterior distribution and an estimate of the Bayesian evidence, which can be used for model comparison (Marshall et al., 2006) or to test compatibility between datasets (Handley & Lemos, 2019).

The key challenge in implementing a nested sampling algorithm is constructing a method that generates samples drawn from the prior, subject to a hard likelihood constraint. Whilst there is a wide variety of publicly available implementations for doing so (see Buchner, 2021; Ashton et al., 2022, for an exhaustive list), these methods are only capable of scaling to hundreds of dimensions (Handley, 2023).

One way to improve the performance of nested sampling algorithms is to use the information about the gradient of the likelihood to propose new points. However, gradient-guided sampling on constrained domains is not straightforward. Whilst the materials science and chemistry literature has made extensive use of gradient-guided nested sampling (Baldock et al., 2017; Pártay et al., 2021; Habeck, 2015; Nielsen, 2013), these methods are often generally bespoke to their physical problems of interest and are not suitable as general-purpose Bayesian samplers. Previous works, such as Betancourt (2011); Speagle (2020) have shown the potential of reflective slice sampling (Neal, 2003), also known as Galilean Monte Carlo (GMC) (Feroz & Skilling, 2013) or Hamiltonian slice sampling (HSS) (Zhang et al., 2016; Bloem-Reddy & Cunningham, 2016), for general-purpose sampling, including the case of non-smooth functions (Mohasel Afshar & Domke, 2015). However, these approaches have found HSS to be “*substantially less efficient (and in general less reliable) than other gradient-based approaches*”.¹ Another alternative is recent work in proximal nested sampling (Cai

¹Verbatim from the documentation of the `dynesty` algorithm.

et al., 2022; McEwen et al., 2023), which uses proximal operators to propose new points. However, this method only works for log-concave likelihoods.

In this work, we combine ideas from across the nested sampling literature and learning-based samplers and create a new gradient-guided nested sampling (GGNS) algorithm that benefits from both algorithmic changes and improved implementation. The five key differences with previous work are 1) the use of self-tuning HSS in combination with gradients calculated through differentiable programming to propose new points, 2) incorporating parallel updates using ideas from dynamic nested sampling (Higson et al., 2019; Speagle, 2020; Handley et al., 2015b) to increase the speed of calculations, 3) A novel termination criterion, and 4) cluster identification to avoid mode collapse, 5) The addition of techniques to more efficiently use trajectories, particularly sampling along accepted points in the trajectory, perturbing the trajectory with random noise during the dynamics, and the detection and "pruning" of points that have left the sampling region. We show that with these changes in combination, our GGNS algorithm scales to significantly higher-dimensional problems without necessitating a proportional increase in the number of live points with respect to dimensionality. This allows GGNS to perform fast and reliable inference for high-dimensional problems.

We show that the GGNS method presented in this work can be used to perform inference in a wide range of problems and that it can be used to improve the performance of existing nested sampling algorithms. Furthermore, we compare our method to existing algorithms for posterior inference and show that it outperforms them, particularly when dealing with highly multimodal distributions. One of the main advantages of the proposed approach is that it requires little hyperparameter tuning and can be used out-of-the-box in a wide range of problems.

Finally, we show the potential of combining nested sampling with generative flow networks (GFlowNets, Bengio et al., 2021; 2023), which are policy learning algorithms that are trained to generate samples from a target distribution and can flexibly be trained off-policy (Malkin et al., 2023). We show how we can use nested sampling to guide GFlowNet training, leading to faster mode finding and convergence of evidence estimates than with traditional GFlowNets. Conversely, we also show how the amortization achieved by GFlowNets can be used to obtain large amounts of high-quality samples from the posterior distribution.

2. Background and Related Work

2.1. Nested Sampling

Nested sampling is used for estimating the marginal likelihood, also known as the evidence, in Bayesian inference

problems (Skilling, 2006):

$$Z = \int \mathcal{L}(\theta)\pi(\theta)d\theta, \tag{1}$$

where $\mathcal{L}(\theta)$ is the likelihood function and $\pi(\theta)$ is the prior distribution. This integral is often intractable due to the high dimensionality and complexity of modern statistical models. In the process of calculating this integral, nested sampling also produces samples from the posterior distribution.

At its core, nested sampling transforms the evidence integral into a one-dimensional nested sequence of likelihood-weighted prior mass, allowing for efficient exploration of the parameter space. The key idea is to enclose the region of high likelihood within a series of nested iso-likelihood contours. This is achieved by introducing a set of *live points* distributed within the prior space and successively updating this set by iteratively replacing the point with the lowest likelihood with a new point drawn from the prior while ensuring the likelihood remains above a likelihood threshold.

As nested sampling progresses, it adaptively refines the prior volume containing higher-likelihood regions. By constructing a sequence of increasing likelihood thresholds, nested sampling naturally focuses on the most informative regions of parameter space. Consequently, nested sampling offers several advantages, including robustness to multimodality in posterior distributions, convergence guarantees, and the ability to estimate posterior probabilities and model comparison metrics. A more detailed review of the algorithm can be found in §A.

The number of nested sampling likelihood evaluations scales as (Skilling, 2006; Handley, 2023):

$$n_{\text{like}} \propto n_{\text{live}} \times f_{\text{sampler}} \times \mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi), \tag{2}$$

where n_{live} is the number of live points, f_{sampler} is the efficiency of the live point generation method (the average number of likelihood evaluations required to generate each new sample), and $\mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi)$ is the Kullback-Leibler divergence between the posterior and the prior.

We should consider the three terms separately to understand the scaling of nested sampling with dimensionality. Here, $\mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi)$ is fixed by the problem at hand (so cannot be modified without substantial adjustment of the meta-algorithm (Petrosyan & Handley, 2022)), and is usually assumed to scale linearly with the number of dimensions. n_{live} for most algorithms scales linearly with dimensionality for two independent reasons: First since the uncertainty in the log-evidence estimation is approximately (Skilling, 2006)

$$\sigma(\log Z) \approx \sqrt{\mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi)/n_{\text{live}}}, \tag{3}$$

if we wish to keep this constant we must scale n_{live} with $\mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi)$, which as discussed before scales linearly with

dimension. Second, most practical live point generation methods require a minimum number of points to tune their internal parameters (such as ellipsoidal/cholesky decompositions or neural network training), and this minimum number scales with dimensionality. The next section describes f_{sampler} scaling.

2.2. Previous Work

The key difficulty in nested sampling is that to generate a new point, one needs to sample points from the prior subject to a hard likelihood constraint:

$$\{\theta \sim \pi : \mathcal{L}(\theta) > \mathcal{L}_*\}. \quad (4)$$

Broadly, the mechanisms for achieving this fall into two classes: *region sampling* and *step sampling* (Ashton et al., 2022). Region samplers have excellent performance in low dimensions, but have a computational cost that scales exponentially with dimensionality $f_{\text{sampler}} \sim \mathcal{O}(e^{d/d_0})$, where $d_0 \sim \mathcal{O}(10)$ is both method and problem dependent. Step samplers have a live point generation cost that scales linearly with dimensionality $f_{\text{sampler}} \sim \mathcal{O}(d)$, so are less efficient in low dimensions.

Region samplers use the current set of live points to define a proxy that encapsulates the likelihood-constrained region Eq. 4, and then appropriately samples from this proxy. For example `MultiNest` (Feroz & Hobson, 2008; Feroz et al., 2009; Feroz et al., 2019) achieves this with an ellipsoidal decomposition fit to the current set of live points, `nessai` (Williams et al., 2021; 2023) trains a normalising flow and `ultranest` (Buchner, 2021) places ellipsoidal kernels on each live point.

Step samplers run a Markov chain starting from one of the current live points, terminating when one has decorrelated from the initial point and then using the final point of the chain as new point. Whilst Skilling (Skilling, 2006) originally envisaged a Metropolis Hastings step, in practice on its own this is a poor choice for sampling from hard-bounded regions. `proxnest` (Cai et al., 2022; McEwen et al., 2023) uses `prox`-guided Langevin diffusion, `DNest` (Brewer & Foreman-Mackey, 2016) offers a flexible framework for programming one’s own stepper, `neuralnest` (Moss, 2020) uses normalizing flow guided Metropolis steps and `PolyChord` (Handley et al., 2015a;b) uses slice sampling. Finally, `dynesty` (Speagle, 2020) and `ultranest` (Buchner, 2021) offer Python re-implementations of many of the above within a single package, with a default dimensionality-dependent switching between region and path sampling.

Dynamic nested sampling (Higson et al., 2019; Speagle, 2020) is a variant of nested sampling which proposes eliminating and replacing multiple points at each iteration.

It was initially implemented in the `dyPolyChord`² and `dynesty`³ packages, but now is common to many implementations (Ashton et al., 2022). It has two main use-cases; increasing the number of posterior samples generated by nested sampling, and implementing parallelization schemes.

2.3. Hamiltonian Slice Sampling

HSS was first introduced in the context of slice sampling (Neal, 2003), as a variant of Hamiltonian Monte Carlo. As in slice sampling, the algorithm initially selects an *initial point* from the current set of live points and a direction. An initial *momentum* variable p_{ini} , which is a d -dimensional array (where d is the dimension of the space), is also defined, typically by randomly sampling a unit vector. The algorithm then proceeds by simulating the *trajectory* of a particle located at the initial point with the chosen initial velocity integrated with some time step Δt , such that at each step the position of the particle is updated according to $x' = x + p\Delta t$. When the particle goes beyond the slice, it is *reflected* back into the slice. This reflection is performed by updating the momentum from p to p' using the equation

$$p' = p - 2(p \cdot n)n, \quad n := \nabla \mathcal{L}(\theta) / \|\nabla \mathcal{L}(\theta)\|, \quad (5)$$

where n is the unit vector in the direction of the likelihood gradient and thus the normal vector to an iso-likelihood surface. Note that, because we are only using the direction of the gradient, one can equivalently use the gradient of the log-likelihood, i.e. the score, which is more efficient to compute. We summarize the algorithm in §B.

As highlighted in Neal (2003), Eq. 5 is only exact when the point where the reflection of the trajectory of the particle takes place is exactly on the boundary $\mathcal{L}(\theta) = \mathcal{L}_*$. In practice, we can either use a small tolerance ϵ to define a neighborhood around the slice and reflect a trajectory whenever the particle is within this neighborhood, or reflect a trajectory whenever the particle lands at a point outside the boundary. The latter method has a theoretical risk of a particle getting “stuck” behind the boundary (in which case the trajectory would be rejected, and a new initial momentum would be chosen).

HSS (or GMC) has been used for nested sampling before (Betancourt, 2011; Feroz & Skilling, 2013; Speagle, 2020). However, the `dynesty` implementation and defaults of HSS lacks the efficiency and reliability of other sampling methods. In addition, in these public implementations, a score has to be manually provided since the package is not compatible with modern differentiable programming frameworks.

²dypolychord.readthedocs.io/en/latest/

³dynesty.readthedocs.io/en/stable/

3. Contributions

In this section, we outline the key combination of ingredients in GGNS we use to significantly improve its performance in high-dimensional settings in comparison with existing publicly available tools.

In brief: we introduce trimming & adaptive step size techniques to remove the hyperparameter tuning difficulties that have beset previous implementations, bring in the current state-of-the-art in parallelization and cluster recognition, and implement in differentiable programming which removes the requirement of providing a score function. With these innovations we find that one only needs $\sim O(1)$ bounces to have decorrelated the chain from the start point, allowing sublinear f_{sampler} scaling. Finally, for maximum posterior scaling, the fact that gradients guide the path means one no longer requires $n_{\text{live}} \sim O(d)$, giving an in-principle linear scaling with dimensionality for the purposes of posterior estimation.

This linear scaling has a theoretical basis. For methods such as slice sampling, taking n steps in a d dimensional space leads to sampling an n -dimensional subspace. Therefore, we need to reach $O(d)$ steps to explore the full space. For Hamiltonian slice sampling, on the other hand, every time we use the gradient for a reflection, we get information about the full d dimensional space. Therefore, each step is exploring the full volume, leading to the requirement of $O(1)$ reflections. This is a similar argument to the better scaling with the dimensionality of Hamiltonian Monte Carlo methods, compared to methods such as random walk Metropolis-Hastings.

In detail, our contributions are the following:

Adaptive Time Step Control We add an adaptive time-step control mechanism in the HSS algorithm. In HSS, particles move in straight lines and eventually reflect off the hard likelihood boundary. To ensure the trajectories between reflections strike a balance between efficiency and accuracy, we introduce the concept of a variable time step, denoted as dt . This time step is adjusted dynamically during the course of the algorithm. By monitoring the number of reflections, we increase or decrease dt to optimize the computational efficiency while maintaining trajectory integrity. This approach, inspired by Neal’s work in (Neal, 2003), enables us to employ larger time steps, thereby reducing the number of reflections without compromising trajectory quality.

Trajectory Preservation In our second enhancement, we introduce a novel approach to preserving and utilizing trajectory⁴ information during the HSS updates. Specifically,

⁴The term trajectory here refers to the states of the chain, not the intermediate states of a Hamiltonian trajectory, as in Nishimura & Dunson (2020).

we store all points along the trajectory after a designated number of reflections, where $\text{min_ref} < \text{max_ref}$. This archive of trajectory points allows us to efficiently select a new live point by uniformly sampling the stored trajectories in a fully parallel manner. We also perturb trajectories by adding some noise delta_p , to achieve faster decorrelation of the samples.

Pruning Mechanism To further enhance efficiency, we introduce a “pruning” mechanism during the HSS process. Points that have remained outside the slice for an extended duration are identified and removed from consideration. These pruned points are then reset to their initial positions, and new momenta are randomly assigned. This mechanism significantly improves the computational efficiency of the proposed method, as we do not waste computational resources evaluating the likelihood of points that have drifted far away from the slice.

Parallel Evolution of Live Points As in (Burkoff et al., 2012; Henderson & Goggans, 2014; Martiniani et al., 2014; Handley et al., 2015b), we implement a dynamic approach to live point management, whereby half of the live points are “killed” at each iteration and replaced with new points. The new set of live points evolves with our HSS algorithm entirely in parallel, given that the HSS algorithm boils down to simulating simple dynamics for all the live points. This parallelism dramatically accelerates the algorithm’s execution.

Mode Collapse Mitigation To address the issue of mode collapse, we incorporate a clustering recognition and evolution algorithm as developed and implemented in PolyChord (Handley et al., 2015b). During the execution of the nested sampling process, we identify clusters of points and keep track of the volume of each cluster. Then, we spawn points proportionally to this volume. This addition helps maintain diversity among live points, preventing them from converging prematurely to a single mode.

Robust Termination Criterion Our final contribution involves the introduction of an alternative termination criterion, which we find to be more robust. Unlike previous implementations of nested sampling that rely on the remaining prior volume X , we utilize the property that the quantity $X\mathcal{L}(\theta)$ follows a characteristic trajectory—initially increasing, reaching a peak, and then decreasing. We terminate the algorithm when $X\mathcal{L}(\theta)$ has decreased by a predetermined fraction from its maximum value. This termination is further explained in §G. This criterion proves to be more resilient to variations in hyperparameters, including the number of live points.

Differentiable Programming Whilst nested sampling algorithms written in differential programming languages exist in `jax` (Albert, 2020) and `torch` (Paszke et al., 2019;

Anau Montel et al., 2023), these do not make use of gradients in guiding the choice of a new live point. Therefore, their choice of using a differentiable programming language is motivated mainly by the advantages of GPU interoperability. To our knowledge, ours is the first algorithm utilizing the gradients derived by differentiable programming to guide the choice of a new live point. Furthermore, this adaptation of nested sampling to hardware intended for modern machine learning workflows, featuring massive parallelization on GPUs; is particularly important in data processing settings that combine nested sampling with deep learning, such as when the prior or likelihood models are given by deep neural networks. We show an example of this when we combine nested sampling and generative flow networks in §5.

We summarize the hyperparameters in §C, and provide extended ablation studies in §E. We also show a summarized version of our algorithm in, including all details, in §F.

Algorithm 1 A summarized version of the GGNS algorithm. (Full version in §F.)

- 1: Initialise n_{live} live points from the prior $\pi(\theta)$.
- 2: Initialise an empty set of dead points.
- 3: Evaluate the likelihood $\mathcal{L}_i = \mathcal{L}(\theta_i)$ for each live point.
- 4: Set $\Delta(X\mathcal{L}) = 0$.
- 5: Set $(X\mathcal{L})_{\text{max}} = 0$.
- 6: Set $dt = dt_{\text{ini}}$.
- 7: **while** $\Delta(X\mathcal{L}) > \text{tol}$ **do**
- 8: Update number of clusters with Alg. 5
- 9: Starting from the remaining $n_{\text{live}}/2$ live points, use HSS to generate $n_{\text{live}}/2$ new points θ_{new} , under the condition $\mathcal{L}(\theta_{\text{new}}) > \mathcal{L}_j \forall \theta_{\text{new}}$.
- 10: Update dt based on the number of trajectory steps
- 11: Add θ_{new} to the set of live points.
- 12: Set $(X\mathcal{L})_{\text{max}} = \max(X\mathcal{L}_{\text{max}}, (X\mathcal{L})_{\text{max}})$, where \mathcal{L}_{max} is the maximum likelihood amongst the live points.
- 13: Set $(\Delta X\mathcal{L}) = X\mathcal{L}_{\text{max}} / (X\mathcal{L})_{\text{max}}$
- 14: **end while**
- 15: **for** $i = 1, \dots, n_{\text{live}}$ **do**
- 16: Select the live point with the lowest likelihood \mathcal{L}_j and move it from the set of live points to the set of dead points.
- 17: Set $Z = Z + \frac{1}{n_{\text{live}}+1} X\mathcal{L}_j$.
- 18: Set $X = X \frac{n_{\text{live}}}{n_{\text{live}}+1}$.
- 19: **end for**

4. Experiments

4.1. Comparison with Other Nested Sampling Methods

We compare the performance of gradient-guided nested sampling with two popular nested sampling algorithms, already

Table 1. Log-evidence function estimation bias (mean and standard deviation over 10 runs).

Method	Gaussian mixture	Funnel	Many well
SMC	-0.569 ± 0.010	0.561 ± 0.801	14.990 ± 1.078
PIS	1.274 ± 0.218	0.262 ± 0.008	2.69 ± 0.04
FAB	0.0027 ± 0.0012	0.0022 ± 0.0011	0.0315 ± 0.0094
Ours	0.029 ± 0.132	-0.051 ± 0.353	0.10 ± 0.60

introduced in §2: PolyChord and dynesty. We use the same likelihood function for all algorithms, which is a Gaussian likelihood with a diagonal covariance matrix, and therefore has $\mathcal{D}_{\text{KL}}(\mathcal{P}|\Pi) \propto d$.

For PolyChord, since $f_{\text{sampler}} \propto n_{\text{repeats}} = 5d$, from Eq. 2 we therefore expect $n_{\text{like}} \propto n_{\text{live}} \times 5d^2$. For dynesty, its default mode swaps between a region sampler with $n_{\text{like}} \propto n_{\text{live}} \times e^{d/d_0} d$ in low dimensions to a slice sampler with $n_{\text{repeats}} = d$, giving $n_{\text{like}} \propto n_{\text{live}} \times d^2$. For GGNS, since $f_{\text{sampler}} \sim \text{max_ref} \sim \mathcal{O}(1)$, we instead expect

$$n_{\text{like}} \propto n_{\text{live}} \times d. \tag{6}$$

For demonstrating the various competing effects discussed in §2.1, §2.2, §3, we set $n_{\text{live}} = 200$, independent of dimensionality. Note that constant n_{live} mode is not usually recommended for these samplers, since, as discussed in §2.1, we need a minimum number of live points to tune the live point generation hyperparameters. Since GGNS uses gradients to guide the choice of live points, it is not restricted in this way.

The results are shown in Fig. 1. We observe the scaling expected from the discussion above. At constant n_{live} , PolyChord has quadratic scaling with dimensionality, providing good evidence estimates until the dimensionality becomes similar to the $n_{\text{live}} = 200$. dynesty is most efficient but exponentially scaling in low dimensions, and swaps to quadratic scaling in higher dimensions when it moves over to slice sampling, at a lower constant than PolyChord due to its default $n_{\text{repeats}} = d$ in comparison with $5d$. Note, however, that this factor of 5 default efficiency is traded off against poor evidence estimates, even in low dimensions, once it is in slice sampling mode.

GGNS, as predicted, has by far the best (linear) scaling and performs evidence estimation accurately even as the dimensionality approaches the number of live points since its live point generation is guided by gradients rather than the other live points. Note, however, that as expected from Eq. 3, the error increases with the square root of the dimensionality at fixed n_{live} .

4.2. Synthetic Datasets

We compare our algorithm with other methods for sampling posterior distributions in a series of commonly used




Figure 1. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log Z$ for different dimensionalities between this work (blue), and other nested sampling algorithms (PolyChord in green and `dynesty` in orange; showing more efficient log-log linear scaling while achieving a higher-fidelity estimate. All comparisons are done for a Gaussian likelihood with a diagonal covariance matrix. The error bars show the standard deviation over 10 runs. Error bars for PolyChord and `dynesty` are also present but barely visible. Note that the last point for `dynesty` is not shown, as it is too large to fit in the plot.




Figure 2. Contour lines for the target distribution, and samples, for the first four dimensions of the 32-dimensional “Many Wells” problem. On the left, we show results for Flow Annealed Importance Sampling Bootstrap (FAB) with a replay buffer, and on the right for GGNS.

synthetic tasks. We compare with the following methods: Sequential Monte Carlo (SMC, Halton, 1962; Gordon et al., 1993; Chopin, 2002; Del Moral et al., 2006), Path Integral Sampler (PIS, Zhang & Chen, 2022), generative flow networks (GFlowNet Bengio et al., 2021; 2023; Lahlou et al., 2023), and Flow Annealed Importance Sampling Bootstrap (FAB, Midgley et al., 2023). For SMC, we use the particles implementation⁵. For PIS, we compare with the on-policy version alone, as it obtains better results than the off-policy version, and we use the reweighted estimate of $\log Z$. Finally, for FAB, we use the publicly available PyTorch implementation⁶, and always use the FAB version with a replay buffer.

We compare these methods with GGNS in three tasks, already introduced in (Hoffman et al., 2014; 2019; Zhang

⁵<https://particles-sequential-monte-carlo-in-python.readthedocs.io/en/latest/>

⁶<https://github.com/lollcat/fab-torch>

& Chen, 2022; Lahlou et al., 2023; Midgley et al., 2023):

The first one is a **Gaussian mixture** in 2-dimension, which consists of a mixture of 9 mode-separated Gaussians. The second one is known as Neal’s **funnel distribution** (Neal, 2003), which is a 10D distribution with a funnel shape. Finally, the third one is the 32-dimensional **Many Well** problem. The Many Well problem is a particularly challenging one, due to the high its high multimodality.

As our benchmark, we use the accuracy of the estimate in the log-evidence, or log-partition function. The calculation of the Bayesian evidence is a good way to evaluate the performance of inference algorithms. We report the mean and standard error of the estimation bias over 10 independent runs in Table 1. We observe that gradient-guided nested sampling obtains unbiased estimates in all three tasks, something that does not happen for any of the other methods studied in this work. While our standard deviation is higher than that of other methods, these can be reduced by adjusting the hyperparameters of our method. However, Eq. 3 shows that the nested sampling log-evidence error can only be reduced sublinearly by increasing the number of live points n_{live} , which increases the computational cost. We cannot, therefore, expect substantial improvements in GGNS log-evidence error bars without innovations in the nested sampling algorithm itself.

When focusing on the quality of generated samples, Fig. 2 shows a comparison between our samples, and those generated using the state of the art method (FAB), on the Many Well task. We see how our methods generates samples of similar quality to FAB, despite a much quicker running time (our method takes only a few minutes in a CPU for this task), and less need for hyperparameter tuning.




Figure 3. **First row:** The true image and noise that we aim to reconstruct. **Second row:** The mean from our gradient-guided nested sampling and the standard deviation. We see how the GGNS posterior matches the expected one.

4.3. Image Generation

We also tested the performance of GGNS on a high-dimensional problem, sampling the posterior distribution over image pixels. To do this, we chose the problem of inferring the pixel values of background galaxies in strong gravitational lensing systems (e.g., Adam et al., 2022). We assumed a correlated (and non-zero mean) normal prior distribution for the background source. A sample from the prior was generated (representing the background galaxy) and was distorted by the potential of a foreground lens. Gaussian noise was then added to produce a noisy simulated data. Given the data, the posterior of a model (a pixelated image of the undistorted background source) could be calculated by adding the likelihood and the prior terms. Furthermore since the model is perfectly linear (and known) and the noise and the prior are Gaussian, the posterior is a high-dimensional Gaussian posterior that could be calculated analytically, allowing us to compare the samples drawn with GGNS with the analytic solution.

Fig. 3 shows a comparison between the true image and its noise and the one recovered by GGNS. We see that we can recover both the correct image, and the noise distribution. We emphasize that this is a uni-modal problem and that the experiment’s goal is to demonstrate the capability of GGNS to sample in high dimensions (in this case, 256), such as images, and to test the agreement between the samples and a baseline analytic solution. We can also compare samples from our method, with samples obtained using Hamiltonian Monte Carlo. We use the `HamiltonTorch` implementation (Cobb, 2023). Using HMC, we get samples with a Wasserstein distance of 1.91 with respect to ground truth values, compared to 1.53 with GGNS. Importantly, needed a sweep of hyperparameters, as most settings just failed with-




Figure 4. Energy histograms for the Lennard-Jones potential with 13 particles (left) and 55 particles (right). Ground truth is shown in blue, and GGNS in orange.

out accepting any new samples, while our method works directly with the default hyperparameters

4.4. Lennard-Jones Potential

To further assess the performance of GGNS on challenging high-dimensional problems, we evaluated the algorithm on the Lennard-Jones potential, a widely used model in computational chemistry. The Lennard-Jones potential describes the interactions between pairs of neutral atoms or molecules. We considered the 13- and 55-dimensional Lennard-Jones potentials, representing a system with many interacting particles. This task has been used for benchmarking in previous work (see, for example, Köhler et al., 2020; Klein et al., 2024; Akhound-Sadegh et al., 2024). Sampling from the resulting high-dimensional (39-dimensional and 165-dimensional, respectively) probability distribution is a notoriously difficult task due to the complex energy landscape characterized by numerous local minima and barriers. It is increasingly difficult for conventional sampling methods to explore the distribution effectively with increasing dimensionality, and many state-of-the-art algorithms struggle to maintain efficiency and accuracy, often suffering from mode collapse or poor mixing.

We show energy histograms for LJ13 and LJ55 in Fig. 4. We can see how GGNS can obtain high-quality samples even for LJ55, for which other samplers such as DDS and PIS simply diverge, while FAB cannot reach low enough energies (see Figure 4 of Akhound-Sadegh et al. (2024)). Remarkably, we can run LJ13 GGNS in approximately 30 minutes in a single CPU and even LJ55 in 12 hours in one CPU (and much faster if we use parallelization).

5. Combination with Generative Flow Networks

We show how the samples obtained from the proposed nested sampling procedure can augment amortized sampling algorithms, such as the generative flow networks considered in §4.2. In Lahlou et al. (2023), it was shown that Euler-Maruyama integration of a stochastic differential equation (SDE) can be viewed as the generative process of a gen-

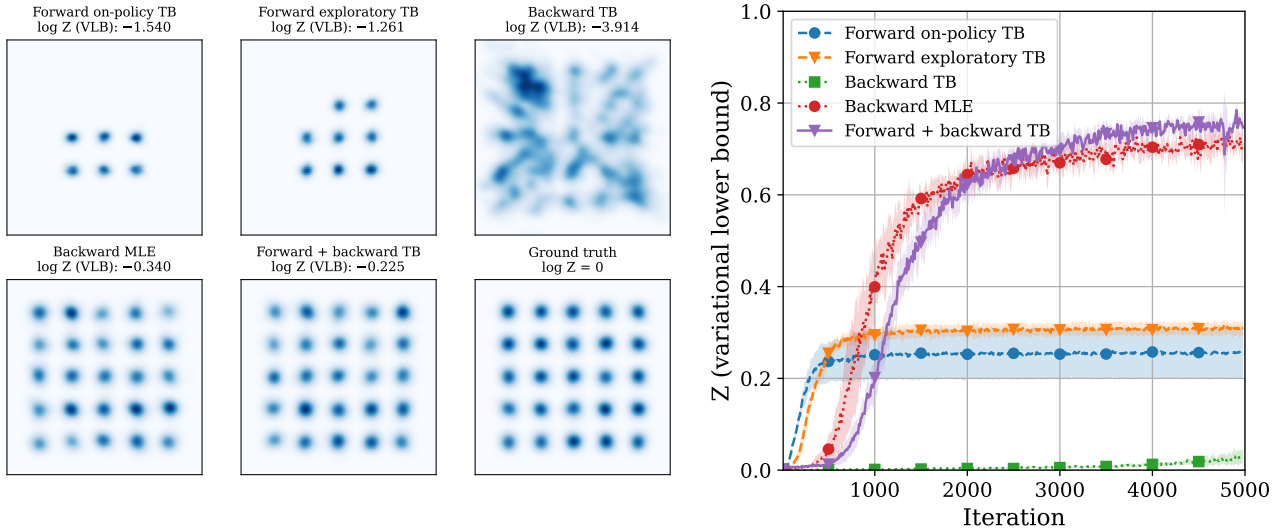


Figure 5. **Above:** KDE plots of samples from trained stochastic control models (ground truth distribution at lower right). Mixing forward sampling with noising trajectories initialized at nested sampling outputs results in all modes being modelled accurately. **Below:** Variational lower bound on the partition function for the five samplers. The theoretical maximum – achieved by the Schrödinger bridge between the Dirac distribution at the origin and the target distribution – is $Z = 1$.

erative flow network. The drift and diffusion terms of the SDE can be trained as the GFlowNet’s forward policy to sample from a target distribution given by an unnormalized density. In particular, GFlowNet objectives can be used to learn the reverse to a Brownian bridge between a target distribution and a point, amounting to approximating the reverse to particular kind of diffusion process. The trajectory balance objective – which directly optimizes for agreement of the forward and reverse path measures – is equivalent in expected gradient to the objective of the path integral sampler (Zhang & Chen, 2022) when trained using on-policy forward exploration, but can also be trained using off-policy trajectories to accelerate mode discovery, which was found to be beneficial in Lahlou et al. (2023); Malkin et al. (2023).

Extending the setup of Zhang & Chen (2022); Lahlou et al. (2023), we consider the problem of sampling from a mixture of 25 well-separated Gaussians (see Fig. 5), with horizontal and vertical spacing of 5 between the component means and each component having variance 0.3. The learned sampler integrates the SDE $dx_t = \mu(x, t) dt + 5 dw_t$, where μ is the output of a neural network (a small MLP) taking x and t as input, with initial condition $x_0 = \mathbf{0}$ up to time $t = 1$. The reward for x_1 is the density of the target distribution. The neural network architecture and training hyperparameters are the same as in Lahlou et al. (2023).

We generate a dataset \mathcal{D} of 2715 approximate samples from the target distribution first using GGNS, and then we use bootstrapping to generate equally weighted samples, using the bootstrapping algorithm in (Handley, 2019). We consider five algorithms for training the SDE drift μ :

- (1) **On-policy TB:** We train μ by optimizing the trajectory balance (TB) objective on trajectories obtained by integration of the SDE being trained (equivalent to the path integral sampler objective and to minimization of the KL divergence between forward and reverse path measures).
- (2) **Exploratory TB:** We optimize the trajectory balance objective on trajectories obtained from a noised version of the SDE, which adds Gaussian noise with standard deviation σ to the drift term at each step. Consistent with Lahlou et al. (2023), we linearly reduce σ from 0.1 to 0 over the first 2500 training iterations. Such exploration is expected to aid in discovery of modes.
- (3) **Backward TB:** We optimize the trajectory balance objective on trajectories sampled from the reverse (diffusion) process begun at samples from \mathcal{D} .
- (4) **Backward MLE:** We sample trajectories from the reverse process begun at samples from \mathcal{D} and train μ so as to maximize the log-likelihood of these trajectories under the forward process. This objective amounts to training a diffusion model or score-based generative model (Song & Ermon, 2019; Ho et al., 2020) on \mathcal{D} , as the optimal μ is the score of the target distribution convolved with a Gaussian and appropriately scaled.
- (5) **Forward + backward TB:** We optimize the trajectory balance objective both on trajectories obtained by integrating the SDE forward from samples from \mathcal{D} and on reverse trajectories begun at samples from \mathcal{D} . This method resembles the training policy used by Zhang et al. (2022).

KDE plots of samples from the trained models, as well as

training metric plots, are shown in Fig. 5. Training with on-policy TB alone (1) results in mode collapse, a typical effect of training with a reverse KL objective (Malkin et al., 2023). We see that while noise introduced in forward exploration (2) helps mode discovery, it is insufficient for all modes to be found. Training using trajectory balance on backward trajectories (3) results in spurious modes, as the model is unlikely to see states that are far from those seen along reverse trajectories from \mathcal{D} during training. Maximum (4) discovers all modes of the distribution, as they are represented in \mathcal{D} , but closer inspection reveals that they are not modeled as accurately; this effect is more pronounced when the dataset \mathcal{D} is small. The best sampling performance is reached by models that perform a mix of forward exploration and reverse trajectories from the dataset samples.

It is important to note that with well-tuned exploratory policies – as in (2) – it is possible to coax the model into discovering all of the modes and modeling them with high fidelity. However, the model is highly sensitive to the exploration parameters: if the exploration rate is too high or not reduced slowly enough, the model is slow to converge and blurs or ‘fattens’ the modes, while an exploration rate that is too low results in mode collapse. On the other hand, mixing forward exploration with backward trajectories from the approximate samples allows the sampler to model all of the modes accurately without such tuning. Notably, we found that the forward trajectories in (5) can be sampled either on-policy or from a tempered policy, with little difference in performance.

6. Limitations

While the proposed gradient-guided nested sampling (GGNS) algorithm offers significant improvements over existing nested sampling methods and provides a powerful tool for sampling problems, there are some limitations that should be acknowledged:

Gradient Computation The efficiency of GGNS relies heavily on the availability and ease of computing gradients of the likelihood function. For problems where the likelihood function is non-differentiable or has discontinuous gradients, the performance of GGNS may be hindered. In such cases, alternative approaches like finite differencing or subgradient methods could be explored, but these may come with additional computational costs.

Hyperparameter Tuning Despite the efforts to reduce hyperparameter tuning, GGNS still requires the specification of certain hyperparameters, such as the initial step size, trajectory preservation parameters, and termination criteria. The performance of the algorithm may be sensitive to these hyperparameter choices, and finding optimal settings can be challenging, especially for complex problems.

High-Dimensional Scaling While GGNS demonstrates improved scaling with dimensionality compared to existing nested sampling methods, its performance in extremely high-dimensional spaces (e.g., thousands of dimensions) remains uncertain. Additional investigations and potential algorithmic enhancements may be required to maintain efficiency in such extreme cases.

Trajectory Perturbations with non-uniform priors Our trajectory perturbation method does not preserve detailed balance in the case of non-uniform priors. There are some potential solutions to this problem: 1) We can “absorb” the prior into the likelihood, thus making the prior again uniform. 2) We can do trajectory perturbations only at the start of the trajectory to achieve decorrelation of the points, followed by some trajectory without noise perturbation. In doing that, we effectively use the part of the trajectory with the noise to find a new starting point and then do a trajectory without noise, and, therefore with the right stationary distribution. 3) We can abandon trajectory perturbations altogether. The goal of the noise is to achieve faster decorrelation of new points, but it is not central to the algorithm’s functioning. Therefore, we can replace it with longer trajectories without noise.

Fixed mass matrix Our algorithm works with a fixed mass matrix in HMC. A potential improvement for future work is to use an adaptive mass matrix, which might improve performance in certain situations.

Despite these limitations, GGNS represents a significant step forward in the field of nested sampling, offering a powerful and scalable approach for Bayesian inference in challenging high-dimensional problems.

7. Discussion and Conclusions

We have introduced a new nested sampling algorithm based on Hamiltonian Slice Sampling. Gradient-guided nested sampling improves on previous nested sampling algorithms by removing the linear dependence of the number of live points on dimensionality. It also makes use of the power of differentiable programming frameworks and parallelization for significant speed improvements. We have shown that the proposed method scales much better with dimensionality than other nested sampling algorithms, thanks to the use of gradient information. This better scaling allows us to apply nested sampling in high-dimensional problems that were too computationally expensive for previous methods. We have also shown that GGNS can be combined with generative flow networks to obtain large numbers of samples from complex posterior distributions. Applications of GGNS to difficult real-world inference problems, both on its own and in combination with amortized sampling methods, are left for future work.

Impact Statement

The problems of posterior sampling and normalizing constant estimation are of general interest in pure machine learning and statistical modelling, as well as in many scientific applications. We see no specific negative impacts or risks of misuse beyond those of other general-purpose probabilistic inference algorithms.

Reproducibility Statement

An implementation of GGNS in PyTorch (Paszke et al., 2019), along with notebooks to reproduce the results from the experiments, is available at <https://github.com/Pablo-Lemos/GGNS>.

ACKNOWLEDGMENTS

This research was supported by the Munich Institute for Astro-, Particle and BioPhysics (MIAPbP) which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC-2094 - 390783311.

References

- Alexandre Adam, Adam Coogan, Nikolay Malkin, Ronan Legin, Laurence Perreault-Levasseur, Yashar Hezaveh, and Yoshua Bengio. Posterior samples of source galaxies in strong gravitational lenses with score-based priors. *arXiv preprint arXiv:2211.03812*, 2022.
- Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated denoising energy matching for sampling from Boltzmann densities. *International Conference on Machine Learning (ICML)*, 2024.
- Joshua G. Albert. JAXNS: a high-performance nested sampling package based on JAX. *arXiv preprint arXiv:2012.15286*, 2020.
- Noemi Anau Montel, James Alvey, and Christoph Weniger. Scalable inference with Autoregressive Neural Ratio Estimation. *arXiv preprint arXiv:2308.08597*, 2023.
- Greg Ashton, Noam Bernstein, Johannes Buchner, Xi Chen, Gábor Csányi, Andrew Fowlie, Farhan Feroz, Matthew Griffiths, Will Handley, Michael Habeck, et al. Nested sampling for physical scientists. *Nature Reviews Methods Primers*, 2(1):39, 2022.
- Robert J. N. Baldock, Noam Bernstein, K. Michael Salerno, Lívía B. Pártay, and Gábor Csányi. Constant-pressure nested sampling with atomistic dynamics. *Physical Review E*, 96(4):043311, October 2017.
- Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Neural Information Processing Systems (NeurIPS)*, 2021.
- Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. GFlowNet foundations. *Journal of Machine Learning Research*, 24:1–76, 2023.
- Julian Besag. Comments on “representations of knowledge in complex systems” by u. grenander and mi miller. *J. Roy. Statist. Soc. Ser. B*, 56(591-592):4, 1994.
- Michael Betancourt. Nested sampling with constrained Hamiltonian Monte Carlo. In *AIP Conference Proceedings*, volume 1305, pp. 165–172. American Institute of Physics, 2011.
- Benjamin Bloem-Reddy and John Cunningham. Slice sampling on hamiltonian trajectories. *International Conference on Machine Learning (ICML)*, 2016.
- Brendon J. Brewer and Daniel Foreman-Mackey. DNest4: Diffusive Nested Sampling in C++ and Python. *arXiv preprint arXiv:1606.03757*, 2016.
- Johannes Buchner. UltraNest - a robust, general purpose Bayesian inference engine. *The Journal of Open Source Software*, 6(60):3001, April 2021.
- Johannes Buchner. Nested sampling methods. *arXiv preprint arXiv:2101.09675*, 2021.
- Nikolas S Burkoff, Csilla Várnai, Stephen A Wells, and David L Wild. Exploring the energy landscapes of protein folding simulations with Bayesian computation. *Biophysical journal*, 102(4):878–886, 2012.
- Xiaohao Cai, Jason D McEwen, and Marcelo Pereyra. Proximal nested sampling for high-dimensional bayesian model selection. *Statistics and Computing*, 32(5):87, 2022.
- Nicolas Chopin. A sequential particle filter method for static models. *Biometrika*, 89(3):539–552, 2002.
- Adam D Cobb. hamiltorch: A pytorch-based library for hamiltonian monte carlo. In *Proceedings of Cyber-Physical Systems and Internet of Things Week 2023*, pp. 114–115. 2023.
- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.

- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.
- F. Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses. *Monthly Notices of the Royal Astronomical Society*, 384(2):449–463, February 2008.
- Farhan Feroz and John Skilling. Exploring multi-modal distributions with nested sampling. In *AIP Conference Proceedings*, volume 1553, pp. 106–113. American Institute of Physics, 2013.
- Farhan Feroz, MP Hobson, and Michael Bridges. MultiNest: an efficient and robust bayesian inference tool for cosmology and particle physics. *Monthly Notices of the Royal Astronomical Society*, 398(4):1601–1614, 2009.
- Farhan Feroz, Michael P. Hobson, Ewan Cameron, and Anthony N. Pettitt. Importance Nested Sampling and the MultiNest Algorithm. *The Open Journal of Astrophysics*, 2(1):10, November 2019.
- Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pp. 107–113. IET, 1993.
- Maximilian N Günther and Tansu Daylan. allesfitter: Flexible star and exoplanet inference from photometry and radial velocity. *The Astrophysical Journal Supplement Series*, 254(1):13, 2021.
- Michael Habeck. Nested sampling with demons. In *Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MAXENT 2014)*, volume 1641 of *American Institute of Physics Conference Series*, pp. 121–129, January 2015.
- John H Halton. Sequential Monte Carlo. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 58, pp. 57–78. Cambridge University Press, 1962.
- Will Handley. anesthetic: nested sampling visualisation. *The Journal of Open Source Software*, 4(37):1414, Jun 2019.
- Will Handley. The scaling frontier of nested sampling. *MaxEnt proceedings (under review)*, October 2023.
- Will Handley and Pablo Lemos. Quantifying tensions in cosmological parameters: Interpreting the des evidence ratio. *Physical Review D*, 100(4):043504, 2019.
- WJ Handley, MP Hobson, and AN Lasenby. Polychord: nested sampling for cosmology. *Monthly Notices of the Royal Astronomical Society: Letters*, 450(1):L61–L65, 2015a.
- WJ Handley, MP Hobson, and AN Lasenby. POLYCHORD: next-generation nested sampling. *Monthly Notices of the Royal Astronomical Society*, 453(4):4384–4398, 2015b.
- R Wesley Henderson and Paul M Goggans. Parallelized nested sampling. In *AIP Conference Proceedings*, volume 1636, pp. 100–105. American Institute of Physics, 2014.
- Edward Higson, Will Handley, Michael Hobson, and Anthony Lasenby. Dynamic nested sampling: an improved algorithm for parameter estimation and evidence calculation. *Statistics and Computing*, 29:891–913, 2019.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Neural Information Processing Systems (NeurIPS)*, 2020.
- Matthew Hoffman, Pavel Sountsov, Joshua V Dillon, Ian Langmore, Dustin Tran, and Srinivas Vasudevan. NeuTralizing bad geometry in Hamiltonian Monte Carlo using neural transport. *arXiv preprint arXiv:1903.03704*, 2019.
- Matthew D Hoffman, Andrew Gelman, et al. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Neural Information Processing Systems (NeurIPS)*, 2024.
- Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. *International Conference on Machine Learning (ICML)*, 2020.
- Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghui Zhang, Alexandra Volokhova, Alex Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of continuous generative flow networks. *International Conference on Machine Learning (ICML)*, 2023.
- Baptiste Lavie, João M Mendonça, Christoph Mordasini, Matej Malik, Mickaël Bonnefoy, Brice-Olivier Demory, Maria Oreshenko, Simon L Grimm, David Ehrenreich, and Kevin Heng. Helios-retrieval: an open-source, nested sampling atmospheric retrieval code; application to the hr 8799 exoplanets and inferred constraints for planet formation. *The Astronomical Journal*, 154(3):91, 2017.
- David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.
- Nikolay Malkin, Salem Lahlou, Tristan Deleu, Xu Ji, Edward Hu, Katie Everett, Dinghui Zhang, and Yoshua

- Bengio. GFlowNets and variational inference. *International Conference on Learning Representations (ICLR)*, 2023.
- Phil Marshall, Nutan Rajguru, and Anže Slosar. Bayesian evidence as a tool for comparing datasets. *Physical Review D*, 73(6):067302, 2006.
- Stefano Martiniani, Jacob D Stevenson, David J Wales, and Daan Frenkel. Superposition enhanced nested sampling. *Physical Review X*, 4(3):031034, 2014.
- Jason D. McEwen, Tobías I. Liaudat, Matthew A. Price, Xiaohao Cai, and Marcelo Pereyra. Proximal nested sampling with data-driven priors for physical scientists. *arXiv preprint arXiv:2307.00056*, 2023.
- Laurence Illing Midgley, Vincent Stimper, Gregor NC Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. *International Conference on Learning Representations (ICLR)*, 2023.
- Hadi Mohasel Afshar and Justin Domke. Reflection, refraction, and Hamiltonian Monte Carlo. *Neural Information Processing Systems (NIPS)*, 2015.
- Adam Moss. Accelerated Bayesian inference using deep learning. *Monthly Notices of the Royal Astronomical Society*, 496(1):328–338, July 2020.
- Pia Mukherjee, David Parkinson, and Andrew R Liddle. A nested sampling algorithm for cosmological model selection. *The Astrophysical Journal*, 638(2):L51, 2006.
- Radford M Neal. Slice sampling. *The annals of statistics*, 31(3):705–767, 2003.
- Radford M Neal et al. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2, 2011.
- Steven O. Nielsen. Nested sampling in the canonical ensemble: Direct calculation of the partition function from NVT trajectories. *The Journal of Chemical Physics*, 139(12):124104, 09 2013.
- Akihiko Nishimura and David Dunson. Recycling intermediate steps to improve Hamiltonian Monte Carlo. *Bayesian Analysis*, 15(4):1087 – 1108, 2020.
- Livia B. Pártay, Gábor Csányi, and Noam Bernstein. Nested sampling for materials. *The European Physical Journal B*, 94(8):159, Aug 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Neural Information Processing Systems (NeurIPS)*, 2019.
- Aleksandr Petrosyan and William James Handley. SuperNest: accelerated nested sampling applied to astrophysics and cosmology. *arXiv preprint arXiv:2212.01760*, 2022.
- Gareth O Roberts and Jeffrey S Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.
- Gareth O Roberts and Richard L Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pp. 341–363, 1996.
- Patricio Maturana Russel, Brendon J Brewer, Steffen Klaere, and Remco R Bouckaert. Model Selection and Parameter Inference in Phylogenetics Using Nested Sampling. *Systematic Biology*, 68(2):219–233, 06 2018.
- John Skilling. Nested sampling for general Bayesian computation. *Bayesian Analysis*, 1(4):833 – 859, 2006.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Neural Information Processing Systems (NeurIPS)*, 2019.
- Joshua S Speagle. dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, 493(3):3132–3158, 2020.
- Michael J. Williams, John Veitch, and Chris Messenger. Nested sampling with normalizing flows for gravitational-wave inference. *Physical Review D*, 103(10):103006, May 2021.
- Michael J. Williams, John Veitch, and Chris Messenger. Importance nested sampling with normalising flows. *Machine Learning: Science and Technology*, 4(3):035011, September 2023.
- David Yallup, Timo Janßen, Steffen Schumann, and Will Handley. Exploring phase space with nested sampling. *The European Physical Journal C*, 82(8):678, 2022.
- Dinghuai Zhang, Nikolay Malkin, Zhen Liu, Alexandra Volokhova, Aaron Courville, and Yoshua Bengio. Generative flow networks for discrete probabilistic modeling. *International Conference on Machine Learning (ICML)*, 2022.
- Qinsheng Zhang and Yongxin Chen. Path integral sampler: a stochastic control approach for sampling. *International Conference on Learning Representations (ICLR)*, 2022.
- Yizhe Zhang, Xiangyu Wang, Changyou Chen, Ricardo Henao, Kai Fan, and Lawrence Carin. Towards unifying Hamiltonian Monte Carlo and slice sampling. *Neural Information Processing Systems (NeurIPS)*, 2016.

A. Nested Sampling Review

Algorithm 2 A simple nested sampling algorithm. Note that more sophisticated implementations include clustering of live points, a calculation of the error in the estimate of $\log Z$, and other techniques to improve performance.

-
- 1: Initialise n_{live} live points from the prior $\pi(\theta)$.
 - 2: Initialise an empty set of dead points.
 - 3: Evaluate the likelihood $\mathcal{L}_i = \mathcal{L}(\theta_i)$ for each live point.
 - 4: Set $X = 1$.
 - 5: Set $Z = 0$.
 - 6: **while** $X_N > \text{tol} \cdot Z$ **do**
 - 7: Select the live point with the lowest likelihood \mathcal{L}_j , and move it from the set of live points to the set of dead points.
 - 8: Sample a new point θ_{new} from the prior $\pi(\theta)$, under the condition $\mathcal{L}(\theta_{\text{new}}) > \mathcal{L}_j$.
 - 9: Set $Z \rightarrow Z + \frac{1}{n_{\text{live}}+1} X \mathcal{L}_j$.
 - 10: Set $X \rightarrow X \frac{n_{\text{live}}}{n_{\text{live}}+1}$.
 - 11: Add θ_{new} to the set of live points.
 - 12: **end while**
 - 13: **for** $i = 1, \dots, n_{\text{live}}$ **do**
 - 14: Select the live point with the lowest likelihood \mathcal{L}_j , and move it from the set of live points to the set of dead points.
 - 15: Set $Z \rightarrow Z + \frac{1}{n_{\text{live}}+1} X \mathcal{L}_j$.
 - 16: Set $X \rightarrow X \frac{n_{\text{live}}}{n_{\text{live}}+1}$.
 - 17: **end for**
-

Nested sampling was initially introduced as a method to calculate the Bayesian evidence or marginal likelihood:

$$\mathcal{Z} = \int \mathcal{L}(\theta) \pi(\theta) d\theta, \quad (7)$$

where $\mathcal{L}(\theta)$ is the likelihood function, and $\pi(\theta)$ is the prior distribution.

The key idea of nested sampling is to define a new variable called the *cumulative prior mass* or the *prior volume* as:

$$X(\theta) = \int_{\mathcal{L}(\theta') > \mathcal{L}(\theta)} \pi(\theta') d\theta', \quad (8)$$

which is the fraction of the prior mass that has a likelihood greater than the likelihood of the current point. This variable is bounded between 0 and 1, and can be used to rewrite the evidence as:

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X) dX, \quad (9)$$

which is a one-dimensional integral. Therefore, we can evaluate the likelihoods of a set of points $\{\theta_i\}$ sorted by their likelihood, and use them to estimate the evidence by approximating the integral in Eq. 9 as a sum:

$$\mathcal{Z} \approx \sum_i \mathcal{L}(X_i) \Delta X_i, \quad (10)$$

where $\Delta X_i = X_{i-1} - X_i$ is the difference in prior volume between the i -th and the $(i-1)$ -th point. This approximation is exact in the limit of an infinite number of points, and can be used to estimate the evidence to arbitrary precision.

The key idea of nested sampling is the following: We start by sampling a set of n_{live} *live points* from the prior distribution. We then find the point with the lowest likelihood, and remove it from the set, adding it to the set of *dead points*. We then replace it with a new point sampled from the prior, subject to the constraint that its likelihood is greater than the likelihood of the point that was removed. This means that, while it is unfeasible to calculate $X(\theta)$ for each of the new points exactly, we can approximate it by using the fact that, at each iteration, the prior volume is contracted by approximately:

$$\Delta X_i \approx \frac{n_{\text{live}}}{n_{\text{live}} + 1}. \quad (11)$$

This process is repeated until the remaining posterior mass is smaller than some fraction of the current estimate of \mathcal{Z} . The set of points that we have sampled can then be used to estimate the evidence using Eq. 10.

Furthermore, Nested Sampling can be used for parameter inference. To do that, we assign the following importance weight to each point:

$$p_i = \frac{\mathcal{L}(w_i)}{Z}, \tag{12}$$

where w_i is the prior volume of the shell that was used to sample the i -th point:

$$w_i = X_{i-1} - X_i. \tag{13}$$

An example implementation of a nested sampling algorithm is shown in Alg. 2.

B. Hamiltonian Slice Sampling Algorithm

Algorithm 3 Hamiltonian or Reflective Slice Sampling

- 1: Choose a point θ from the existing set of live points.
 - 2: Choose a direction d .
 - 3: Choose an initial momentum $p_{\text{ini}} \sim \mathcal{N}(0, 1)$.
 - 4: Set $p = p_{\text{ini}}$.
 - 5: Set $x = \theta$.
 - 6: Set $t = 0$.
 - 7: **while** $t < T$ **do**
 - 8: Set $x = x + p dt$.
 - 9: **if** x is outside the slice **then**
 - 10: Take $n = \nabla \mathcal{L}(\theta) / \|\nabla \mathcal{L}(\theta)\|$
 - 11: Set $p = p - 2(p \cdot n)n$.
 - 12: **end if**
 - 13: Set $t = t + dt$.
 - 14: **end while**
 - 15: Set $\theta' = x$.
-

We show an example implementation of the Hamiltonian or Reflective Slice Sampling algorithm in Alg. 3.

C. Hyperparameters of GGNS

Table 2.

Parameter	Default Value	Description
<code>n_{live}</code>	Number of live points. A higher number leads to better mode coverage.	200 ⁷
<code>tol</code>	Tolerance. The stopping criterion. GGNS terminates when $\mathcal{L}_i X_i / \max(\mathcal{L}_i X_i) < \text{tol}$.	0.01
<code>min_ref</code>	The minimum number of reflections. We sample points after they have reflected of the boundary at least <code>min_ref</code> .	1
<code>max_ref</code>	The maximum number of reflections. We stop each HSS iteration after the point has reflected <code>max_ref</code> times off the boundary.	3
<code>delta_p</code>	The number of noise added to the momentum at each HSS step, to decorrelate samples faster.	0.05

Table Table 2, shows the different hyperparameters of GGNS. This shows the little tuning required for GGNS to perform unbiased sampling.

D. Sampling Complex Distributions

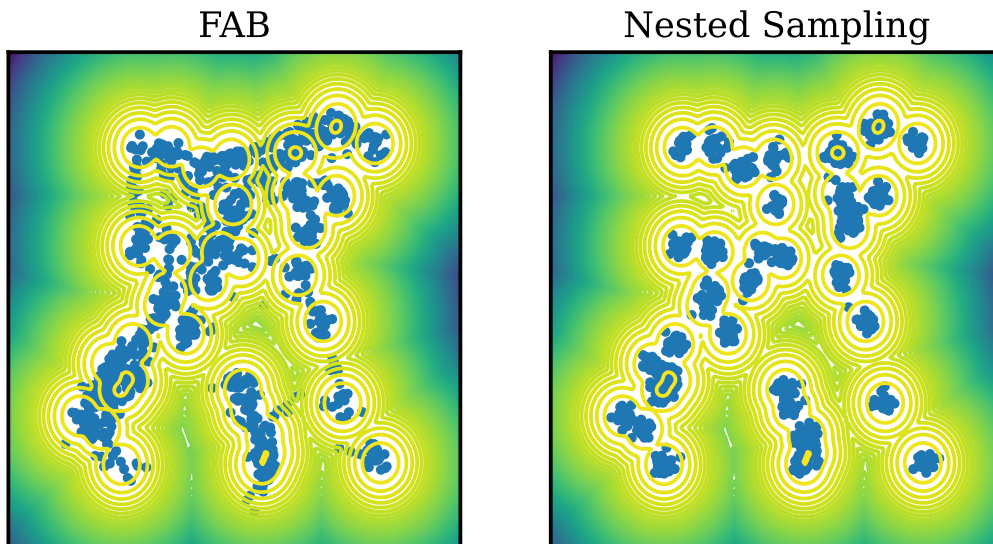


Figure 6. Contour lines for the target distribution, and samples, for the mixture of 40 Gaussians from (Midgley et al., 2023). On the left, we show results for Flow Annealed Importance Sampling Bootstrap (FAB), and on the right for GGNS.

We further test the capacity of GGNS to model several complex distributions that are for different reasons challenging for inference algorithms. For these examples, we use a visual comparison with samples from the true distribution. We increase the number of live points to 2000 for these examples, to ensure that we have enough samples to compare with the true distribution. Because nested sampling produces weighted samples, all the figures use an alpha blending of the samples, with the alpha value proportional to the importance weight.

Firstly, we re-use the Gaussian mixture distribution from the previous example, but we increase the number of modes to 81. This distribution is difficult due to its very high multimodality. The results are shown in the top panel of Fig. 7, where we see that our method successfully recovers all modes.

Secondly, the "five Swiss rolls" example consists of five copies of the "Swiss roll" distribution in two dimensions. It combines multimodality with the difficulty of sampling the complex structure of each mode of the distribution. As shown in the bottom panel of Fig. 7, the proposed method successfully samples the distribution.

Finally, we show our performance on the mixture of 40 Gaussians used in (Midgley et al., 2023), in Fig. 6. The image shows GGNS samples all the modes of the distribution.

E. Ablation Study

Given the multiple different components included in our algorithm, described in §3, it is important to perform an ablation study to fully understand how the different parts of the algorithm contribute to the observed improvements in performance. Therefore, this section removes each of the improvements introduced in §3 one by one and analyses the effect of these changes on the sampling performance. We use the task introduced in §4.1 to perform this comparison unless otherwise specified.

E.1. Adaptive Time Step Control

We repeat the analysis of §4.1 with a fixed time step in the Hamiltonian slice sampling steps instead of using adaptive time step control. We attempt three different time steps: $dt = 0.5$ and $dt = 0.1$. Note that, in §4.1, we start with $dt = 0.1$ but adapt it as the algorithm progresses.

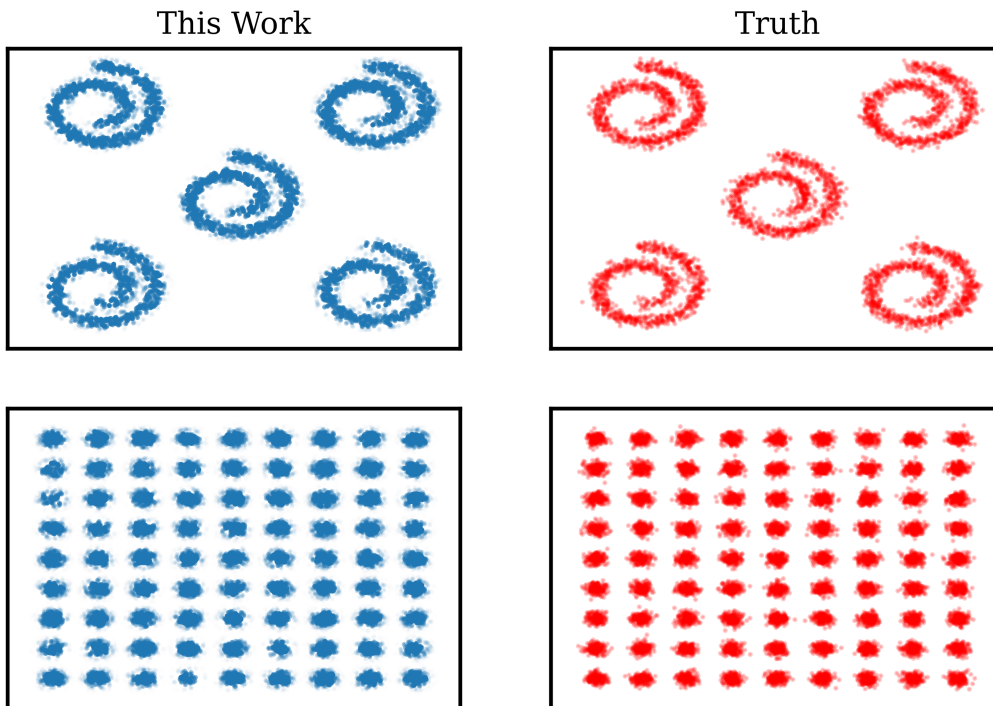


Figure 7. Comparison between the proposed method (left) and the truth (right) on the "five Swiss rolls" distribution. We show our method successfully recovers every model of this highly multi-modal distribution.

The results, shown in Fig. 8, show that when we use a large dt we can reduce the number of likelihood evaluations, as we achieve the minimum number of reflections faster, but we get a biased estimate of $\log Z$, as we fail to appropriately sample each slice. For small dt , on the other hand, we get less bias in $\log Z$, but the number of likelihood evaluations goes up.

In general, the advantage of the adaptive step is that as the algorithm progresses, the volume of the region defined by Eq. 4 we are exploring decreases. Therefore, a step size that is appropriate at a given point will become too large eventually as the algorithm progresses.

E.2. Trajectory preservation

GGNS uses a novel approach to sample the trajectories and to ensure that samples are correlated, where we ensure a certain number of boundary reflections. We compare what happens when we use the simpler approach of integrating our trajectory for a fixed number of steps n_{steps} and simply keeping the last sample. We repeat the analysis for $n_{\text{steps}} = 20$ and $n_{\text{steps}} = 200$.

Fig. 9 shows the results. We see that a fixed number of steps leads to a biased estimate of $\log Z$. The argument for this, similarly to what it was for trajectory preservation, is that the volume of the region Eq. 4 decreases as the algorithm progresses.

We also study what happens when we use trajectory preservation but do not add noise δ_p to achieve a faster decorrelation of the samples. The results, shown in Fig. 10, are intuitive: No noise in the trajectories reduces the likelihood evaluations, as trajectories are less noisy but lead to biased evidence estimates, as the samples are not fully decorrelated.

E.3. Mode Collapse Mitigation

To study the effect of this, we need a multimodal distribution. We use a mixture of nine Gaussians, shown in Fig. 11. To




Figure 8. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log Z$ for different dimensionalities between baseline GGNS (blue), and GGNS without adaptive step control ($dt = 0.5$ in green and $dt = 0.1$ in orange)




Figure 9. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log Z$ for different dimensionalities between baseline GGNS (blue), and GGNS without trajectory preservation ($N_{\text{steps}} = 200$ in green and $N_{\text{steps}} = 20$ in orange). Both of these lead to correlated estimates of $\log Z$, as shown by the right panel.

study the effect of our mode collapse mitigation, we run nested sampling on this problem, with and without this setting. The main hyperparameter that affects the number of modes found, is the number of live points n_{live} . Therefore, we run nested sampling on this problem for different values of this hyperparameter. For each configuration, we run the algorithm 10 times, and count the number of modes found. We define a mode as being found, if at least one of the samples is within a distance σ of the center of the mode.

The results are shown in Table 3. We see how, generally, mode collapse mitigation helps us find a higher number of modes. Although the number of live points is the most important hyperparameter when it comes to mode finding, the ability to find all modes for a fixed n_{live} is higher when using mode collapse mitigation.

E.4. Termination Criterion

Finally, we repeat the analysis of §4.1 using the termination criterion used by other nested sampling algorithms such as DyNesty and PolyChord, in which we terminate the algorithm when $\mathcal{L}_{\text{max}} X_i < \text{tol}$, for some tolerance hyperparameter. We used a tolerance 0.01, a value often used by nested sampling practitioners.

We show the results in Fig. 12. The number of likelihood evaluations appears similar, but at high dimensions, the previous termination leads to a biased estimate of $\log Z$. Indeed, while the number of likelihood estimations is of the same order of magnitude, the termination used by GGNS leads to slightly more evaluations ($\sim 5.8 \cdot 10^6$) than the previous one ($\sim 5.4 \cdot 10^6$) for $d = 128$. These 400,000 evaluations are likely to drive the underestimation of the evidence by the previous method.




Figure 10. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log Z$ for different dimensionalities between baseline GGNS (blue), and GGNS without adding noise to the trajectories (orange). Less noise decreases the number of evaluations but leads to a biased $\log Z$ estimate.

Table 3. The average number of modes found over 10 nested sampling runs, sampling the distribution shown in Fig. 11, with and without mode collapse mitigation, for varying number of live points.

Method	$n_{\text{live}} = 20$	$n_{\text{live}} = 50$	$n_{\text{live}} = 100$	$n_{\text{live}} = 200$
Without Mode Collapse Mitigation	3.6	6.4	8.2	8.9
With Mode Collapse Mitigation	4.1	6.4	8.4	9

E.5. Comparison with DyNesty hslice

The DyNesty algorithm also offers the option of using gradients to perform Hamiltonian slice sampling, as discussed in the main text. However, we have found the performance to be inefficient, something that (Speagle, 2020) also alludes to. Fig. 13 compares GGNS with DyNesty’s version of HSS. We see how, while the evidence estimation is more robust (right panel), the number of likelihood evaluations is extremely large (left), making it practically unusable in high-dimensional settings.

E.6. Other changes

We could not study other changes, such as the pruning mechanism or parallel Evolution of live points, as this would have led to a full rewrite of the algorithm. We leave this study for future work.

E.7. Conclusions

The main conclusion of this ablation study is that naive changes to GGNS quickly lead to biased sampling. It is the combination of the contributions introduced in §3 that leads to robust evidence on high dimensions. Of course, it is true that for each of these parts, there are settings that will work for any problem, i.e. we can always make the step size small enough, the number of steps high enough, etc. However, the main advantage of our algorithm is that it works without the need for fine-tuning all these parameters.

F. GGNS algorithm

We show the full GGNS algorithm in Alg. 4. Our cluster-finding algorithm and our Hamiltonian slice sampling algorithm are shown in Alg. 5 and Alg. 6 respectively and are both used in Alg. 4. The cluster statistics formalism presented in this section follows (Handley et al., 2015b). We refer the reader to the original paper for derivations on where these formulas come from.

The summary statistics are initiated using the following equation:

Algorithm 4 The GGNS algorithm

```

1: Initialise  $n_{\text{live}}$  live points from the prior  $\pi(\theta)$ .
2: Initialise an empty set of dead points.
3: Evaluate the likelihood  $\mathcal{L}_i = \mathcal{L}(\theta_i)$  for each live point.
4: Initiate summary statistics  $\{\bar{Z}, \bar{X}, \bar{Z}_p, \dots\}$ , using Eq. 14 to Eq. 25.
5: Set  $\Delta(X\mathcal{L}) = 0$ .
6: Set  $(X\mathcal{L})_{\text{max}} = 0$ .
7: Set  $n_{\text{clusters}} = 1$ .
8: Set  $dt = dt_{\text{ini}}$ .
9: while  $\Delta(X\mathcal{L}) > \text{tol}$  do
10:  for  $i = 1, \dots, n_{\text{clusters}}$  do
11:    Use cluster finding Alg. 5
12:    If new clusters are found, initiate them by splitting the cluster  $i$ , using Eq. 36 to Eq. 42
13:  end for
14:  for  $j = 1, \dots, n_{\text{live}}//2$  do
15:    Select the point with the lowest likelihood  $\mathcal{L}_j$  and remove them from the set of live points to the set of dead points.
16:    Update the summary statistics, using Eq. 26 to Eq. 35.
17:    If a cluster has no points, remove it, and set  $n_{\text{clusters}} - = 1$ 
18:  end for
19:  Generate the cluster labels for the next points  $n_{\text{live}}//2$ , proportionally to  $\bar{X}_p$ 
20:  Sample  $x \sim \theta_{\text{live}}$ , from the appropriate clusters
21:  Use Alg. 6, to get  $\theta_{\text{new}}$ , and out_frac, under the condition  $\mathcal{L}(\theta_{\text{new}}) > \mathcal{L}_j \forall \theta_{\text{new}}$ .
22:  if out_frac  $> 0.15$  then
23:    Set  $dt = dt * 0.9$ 
24:  else if out_frac  $< 0.05$  then
25:    Set  $dt = dt * 1.1$ 
26:  end if
27:  Add  $\theta_{\text{new}}$  to the set of live points.
28:  Set  $(X\mathcal{L})_{\text{max}} = \max(X\mathcal{L}_{\text{max}}, (X\mathcal{L})_{\text{max}})$ , where  $\mathcal{L}_{\text{max}}$  is the maximum likelihood amongst the live points.
29:  Set  $(\Delta X\mathcal{L}) = X\mathcal{L}_{\text{max}}/(X\mathcal{L})_{\text{max}}$ 
30: end while
31: for  $i = 1, \dots, n_{\text{live}}$  do
32:  Select the live point with the lowest likelihood  $\mathcal{L}_j$  and move it from the set of live points to the set of dead points.
33:  Set  $Z = Z + \frac{1}{n_{\text{live}}+1} X\mathcal{L}_j$ .
34:  Set  $X = X + \frac{n_{\text{live}}}{n_{\text{live}}+1}$ .
35: end for

```

Algorithm 5 The cluster finding algorithm used in Alg. 4, for a cluster containing n_{points} points.

```

1: Initialise prev_sizes = None.
2: for  $k = 2, \dots, n_{\text{points}}$  do
3:  Run k-nearest-neighbours (KNN) on the cluster points, with value  $k$ 
4:  Set cluster_sizes as the number of points in each KNN cluster
5:  if cluster_sizes = prev_sizes then
6:    Break
7:  else
8:    Set prev_sizes = cluster_sizes
9:  end if
10: end for
11: return The number of KNN clusters.

```




Figure 11. The distribution used for our ablation study on mode collapse mitigation, true samples on the right, and samples with GGNS ($n_{\text{live}} = 100$) on the left.

$$\bar{Z} = 0, \quad (14)$$

$$\bar{Z}_p = \{\bar{Z}\}, \quad (15)$$

$$\bar{Z}^2 = 0, \quad (16)$$

$$\bar{Z}_p^2 = \{\bar{Z}^2\}, \quad (17)$$

$$\bar{Z}\bar{X} = 0, \quad (18)$$

$$\bar{Z}\bar{X}_p = \{\bar{Z}\bar{X}\}, \quad (19)$$

$$\bar{Z}_p\bar{X}_p = \{\bar{Z}\bar{X}\}, \quad (20)$$

$$\bar{X} = 1, \quad (21)$$

$$\bar{X}_p = \{\bar{X}\}, \quad (22)$$

$$\bar{X}_p^2 = \{\bar{X}^2\}, \quad (23)$$

$$\bar{X}_p\bar{X}_q = 0 \quad (q \neq p), \quad (24)$$

$$(25)$$

where p and q refer to the cluster numbers, initially 1. To update the summary statistics, we use:




Figure 12. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log \mathcal{Z}$ for different dimensionalities between baseline GGNS (blue), and GGNS using the termination criterion used by MultiNest, PolyChord and DyNesty.

Algorithm 6 The Hamiltonian slice sampling algorithm used in Alg. 4, starting from n points with position x , and with step size dt ; and with a likelihood barrier \mathcal{L}_{\min}

```

1: Set num_out_steps = 0, num_in_steps = 0
2: Set  $p \sim \mathcal{N}(0, 1)$ .
3: Set num_reflections[1, ..., n] < -0
4: Set x_saved = {}
5: while min(num_reflections) < max_reflections do
6:   Set  $x_+ = p * dt$ 
7:   Call the likelihood function, to get  $\mathcal{L}$  and  $\nabla \mathcal{L}$ 
8:   Set outside[1, ..., n] =  $\mathcal{L} < \mathcal{L}_{\min}$ 
9:   Take  $n = \nabla \mathcal{L} / \|\nabla \mathcal{L}\|$ 
10:  Set  $p[\text{outside}] = p[\text{outside}] - 2(p \cdot n)n[\text{outside}]$ .
11:  Set  $\epsilon \sim \mathcal{N}(0, 1)$ 
12:  Set  $p = p * (1 + \epsilon * \text{delta}_p)$ 
13:  Set num_reflections+ = outside
14:  if min(num_reflections) < min_reflections then
15:    Add  $x[\sim \text{outside}]$  to x_saved
16:  end if
17:  Set num_out_steps+ =  $\sum(\text{outside})$ 
18:  Set num_in_steps+ =  $\sum(\sim \text{outside})$ 
19: end while
20: Set out_frac = num_out_steps / (num_out_steps + num_in_steps)
21: Samples  $\theta \sim x\_saved$ 
22: return  $\theta$ , out_frac
    
```

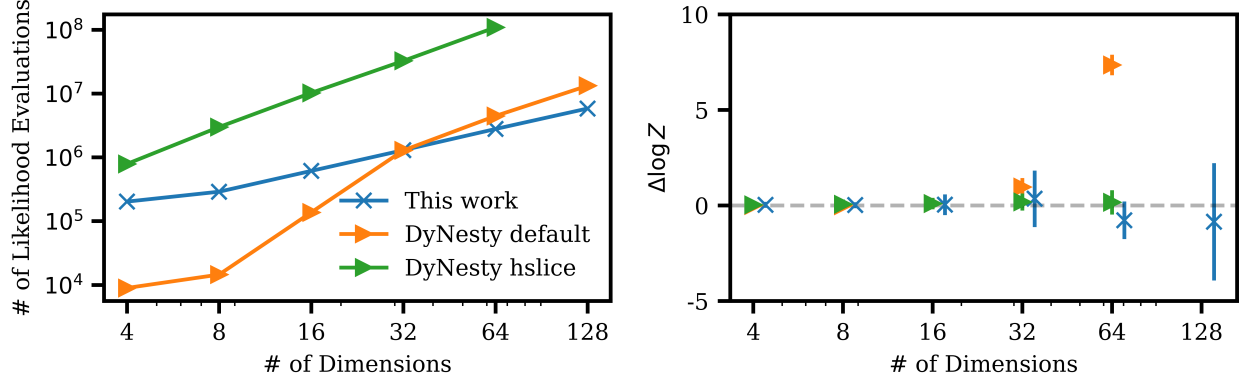


Figure 13. Comparison of likelihood evaluations (left panel) and error in the estimation of $\log Z$ for different dimensionalities between GGNS (blue), DyNesty default (orange) and DyNesty using HSS (green).

$$\bar{Z} \rightarrow \bar{Z} + \frac{\bar{X}_p \mathcal{L}}{n_p + 1}, \quad (26)$$

$$\bar{Z}_p \rightarrow \bar{Z}_p + \frac{\bar{X}_p \mathcal{L}}{n_p + 1}, \quad (27)$$

$$\bar{X}_p \rightarrow \frac{n_p \bar{X}_p}{n_p + 1}, \quad (28)$$

$$\bar{Z}^2 \rightarrow \bar{Z}^2 + \frac{2\bar{Z}\bar{X}_p \mathcal{L}}{n_p + 1} + \frac{2\bar{X}_p^2 \mathcal{L}^2}{(n_p + 1)(n_p + 2)}, \quad (29)$$

$$\bar{Z}_p^2 \rightarrow \bar{Z}_p^2 + \frac{2\bar{Z}_p \bar{X}_p \mathcal{L}}{n_p + 1} + \frac{2\bar{X}_p^2 \mathcal{L}^2}{(n_p + 1)(n_p + 2)}, \quad (30)$$

$$\overline{\mathcal{Z}X_p} \rightarrow \frac{n_p \overline{\mathcal{Z}X_p}}{n_p + 1} + \frac{n_p \bar{X}_p^2 \mathcal{L}}{(n_p + 1)(n_p + 2)}, \quad (31)$$

$$\overline{\mathcal{Z}X_q} \rightarrow \overline{\mathcal{Z}X_p} + \frac{\bar{X}_p \bar{X}_q \mathcal{L}}{(n_p + 1)} \quad (q \neq p), \quad (32)$$


$$\overline{\mathcal{Z}_p X_p} \rightarrow \frac{n_p \overline{\mathcal{Z}_p X_p}}{n_p + 1} + \frac{n_p \bar{X}_p^2 \mathcal{L}}{(n_p + 1)(n_p + 2)}, \quad (33)$$

$$\bar{X}_p^2 \rightarrow \frac{n_p \bar{X}_p^2}{n_p + 2}, \quad (34)$$

$$\overline{X_p X_q} \rightarrow \frac{n_p \overline{X_p X_q}}{n_p + 1} \quad (q \neq p). \quad (35)$$

When we need to split a cluster p into multiple clusters i , we use:

Figure 14. In nested sampling, the likelihood \mathcal{L} (dotted line) goes up, as the algorithm progresses from $X = 1$ to $X = 0$ (right to left in the plot). However, the product $X\mathcal{L}(X)$ starts low, peaks, end goes down again. Therefore, our termination criterion checks when $X\mathcal{L}(X)$ has peaked, and again gone close to zero. Image credit (Handley et al., 2015b).



$$\overline{X}_i = \frac{n_i}{n} \overline{X}_p, \quad (36)$$

$$\overline{X}_i^2 = \frac{n_i(n_i + 1)}{n(n + 1)} \overline{X}_p^2, \quad (37)$$

$$\overline{X}_i X_j = \frac{n_i n_j}{n(n + 1)} \overline{X}_p^2, \quad (38)$$

$$\overline{X}_i Y = \frac{n_i}{n} \overline{X}_p Y \quad Y \in \{Z, Z_p, X_q\}, \quad (39)$$

$$\overline{Z}_i = \frac{n_i}{n} \overline{Z}_p, \quad (40)$$

$$\overline{Z}_i X_i = \frac{n_i(n_i + 1)}{n(n + 1)} \overline{Z}_p \overline{X}_p, \quad (41)$$

$$\overline{Z}_i^2 = \frac{n_i(n_i + 1)}{n(n + 1)} \overline{Z}_p^2. \quad (42)$$

$$(43)$$

F.1. Complexity

- The HSS subroutine has complexity $\mathcal{O}(T)$ where T is the number of time steps we integrate. Our algorithm works until reaching a number of reflections, instead of a number of time steps. However, because we have an adaptive time step that enforces one reflection every N steps (by default $N = 10$), the algorithm's complexity does not scale with latent dimensionality. With respect to number of live points, the algorithm does scale linearly, as each live point evolves independently ($\mathcal{O}(n_{\text{live}})$), however, thanks to our parallelization, if we have n_{live} CPUs/GPUs available, the time complexity will still be $\mathcal{O}(1)$. This is why, as we discuss in the paragraph starting in line 162, right side, the complexity of the subroutine is $\mathcal{O}(1)$, leading to the improved scalability of our algorithm.

- The current clustering subroutine has worse scalability: Because we run k nearest neighbours for the n_{live} points, up to n_{live} times, leading to a complexity $\mathcal{O}(n_{\text{live}}^2)$. This is why we run the cluster-finding algorithm very sporadically. We will explore replacing this clustering algorithm with one with better complexity in future work. Because our algorithm does not need to scale n_{live} with latent dimensionality, however, the scaling with dimensionality is the same one as KNNs $\mathcal{O}(d)$.

G. Termination Criterion Explained

Fig. 14 explains the intuition behind the new termination criterion introduced in §3. In nested sampling, the likelihood increases as the algorithm progresses, corresponding to going from right ($X = 1$) to left ($X = 0$) in the figure. However, the product of the prior volume X and the likelihood \mathcal{L} always follow the same pattern: It starts low, reaches a peak, and then decreases towards zero again. Therefore, our termination criterion consists of checking the ratio of the current value of $X\mathcal{L}$ to the maximum value that $X\mathcal{L}$ has reached throughout the algorithm. When that fraction is smaller than some threshold,




Figure 15. The bias in the estimate of $\log Z$ as a function of number of likelihood evaluations, for GGNS (blue), and other nested sampling algorithms (PolyChord in green and `dynesty` in orange). We achieve different numbers of like evaluations by changing the number of live points. We run each algorithm with $n_{\text{live}} = \{50, 100, 200, 400\}$. Note that the `dynesty` runs with $n_{\text{live}} = 50$ and $n_{\text{live}} = 100$ are not in the plot, as they are too far up in the y-axis.

we stop the algorithm.

H. Evidence Estimation as a Function of Likelihood Evaluations

In this appendix, we study the relationship between the number of likelihood evaluations, and the estimate of $\Delta \log Z$. The easiest way to vary the number of likelihood evaluations, is by varying the number of live points used. We repeat the analysis of §4.1, for two values of the number of dimensions $d = 32$ and $d = 64$, for each of the algorithms; varying the number of live points in the range $n_{\text{live}} = \{50, 100, 200, 400\}$.

We see how, in both cases, `dynesty` can lead to very biased inference, if the number of live points is low (Fig. 15). On the other hand, `PolyChord` reliably achieves unbiased inference, at the expense of a much higher number of likelihood evaluations. GGNS gets the best of each, by achieving unbiased inference with less likelihood evaluations. We also see how our algorithm scales better with dimensionality, when we compare its performance between the left and the right plots, with the other algorithms.