




# Exploring phase space with nested sampling

David Yallup<sup>1,a</sup> , Timo Janßen<sup>2</sup>, Steffen Schumann<sup>2</sup>, Will Handley<sup>1</sup>

<sup>1</sup> Cavendish Laboratory and Kavli Institute for Cosmology, University of Cambridge, JJ Thomson Avenue, Cambridge CB3 0HE, UK

<sup>2</sup> Institut für Theoretische Physik, Georg-August-Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany

Received: 17 May 2022 / Accepted: 23 July 2022 / Published online: 5 August 2022  
© The Author(s) 2022

**Abstract** We present the first application of a Nested Sampling algorithm to explore the high-dimensional phase space of particle collision events. We describe the adaptation of the algorithm, designed to perform Bayesian inference computations, to the integration of partonic scattering cross sections and the generation of individual events distributed according to the corresponding squared matrix element. As a first concrete example we consider gluon scattering processes into 3-, 4- and 5-gluon final states and compare the performance with established sampling techniques. Starting from a flat prior distribution Nested Sampling outperforms the VEGAS algorithm and achieves results comparable to a dedicated multi-channel importance sampler. We outline possible approaches to combine Nested Sampling with non-flat prior distributions to further reduce the variance of integral estimates and to increase unweighting efficiencies.

## 1 Introduction

Realistic simulations of scattering events at particle collider experiments play an indispensable role in the analysis and interpretation of actual measurement data for example at the Large Hadron Collider (LHC) [1, 2]. A central component of such event simulations is the generation of hard scattering configurations according to a density given by the squared transition matrix element of the concrete process under consideration. This is needed both for the evaluation of corresponding cross sections, as well as the explicit generation of individual events that potentially get further processed, e.g. by attaching parton showers, invoking phenomenological models to account for the parton-to-hadron transition, and eventually, a detector simulation. To adequately address the physics needs of the LHC experiments requires the evaluation of a wide range of high-multiplicity hard processes that feature a highly non-trivial multimodal target density

that is rather costly to evaluate. The structure of the target is thereby affected by the appearance of intermediate resonances, quantum interferences, the emission of soft and/or collinear massless gauge bosons, or non-trivial phase space constraints, due to kinematic cuts on the final state particles. Dimensionality and complexity of the phase space sampling problem make the usage of numerical methods, and in particular Monte Carlo techniques, for its solution indispensable.

The most widely used approach relies on adaptive multi-channel importance sampling, see for example [3–7]. However, to achieve good performance detailed knowledge of the target distribution, i.e. the squared matrix element, is needed. To this end information about the topology of scattering amplitudes contributing to the considered process is employed in the construction of individual channels. Alternatively, and also used in combination with importance sampling phase space maps, variants of the self-adaptive VEGAS algorithm [8] are routinely applied [9–12].

An alternative approach for sampling according to a desired probability density is offered by Markov Chain Monte Carlo (MCMC) algorithms. However, in the context of phase space sampling in high energy physics these techniques attracted rather limited attention, see in particular [13, 14]. More recently a mixed kernel method combining multi-channel sampling and MCMC, dubbed (MC)<sup>3</sup>, has been presented [15]. A typical feature of such MCMC based algorithms is the potential autocorrelation of events that can affect their direct applicability in typical use case scenarios of event generators.

To meet the computing challenges posed by the upcoming and future LHC collider runs and the corresponding event simulation campaigns, improvements of the existing phase space sampling and event unweighting techniques will be crucial [16, 17]. This has sparked renewed interest in the subject, largely driven by applications of machine learning techniques, see for instance [18–36].

In this article we explore an alternative direction. We here study the application of Nested Sampling [37] as imple-

<sup>a</sup> e-mail: dy297@cam.ac.uk (corresponding author)

mented in POLYCHORD [38] to phase space integration and event generation for high energy particle collisions. We here assume no prior knowledge about the target and investigate the ability of the algorithm to adapt to the problem. Nested Sampling has originally been proposed to perform Bayesian inference computations for high dimensional parameter spaces, providing also the evidence integral, i.e. the integral of the likelihood over the prior density. This makes it ideally suited for our purpose. In Sect. 2 we will introduce Nested Sampling as a method to perform cross section integrals and event generation, including a reliable uncertainty estimation. In Sect. 3 we will apply the method to gluon scattering to 3-, 4-, and 5-gluon final states as a benchmark for jet production at hadron colliders, thereby comparing results for total cross sections and differential distributions with established standard techniques. Evaluation of the important features of the algorithm when applied in the particle physics context is also discussed in this section. In Sect. 4 we illustrate several avenues for future research, extending the work presented here. Finally, we present our conclusions in Sect. 5.

## 2 Nested sampling for event generation

The central task when exploring the phase space of scattering processes in particle physics is to compute the cross section integral,  $\sigma$ . This requires the evaluation of the transition squared matrix element,  $|\mathcal{M}|^2$ , integrated over the phase space volume,  $\Omega$ , where  $\Omega$  is composed of all possible kinematic configurations,  $\Phi$ , of the external particles. Up to some constant phase space factors this amounts to performing the integral,

$$\sigma = \int_{\Omega} d\Phi |\mathcal{M}|^2(\Phi). \quad (1)$$

In practice rather than sampling the physical phase space variables, i.e. the particles' four-momenta, it is typical to integrate over configurations,  $\theta \in [0, 1]^D$ , from the  $D$ -dimensional unit hypercube. Some mapping,  $\Pi : [0, 1]^D \rightarrow \Omega$ , is then employed to translate the sampled variables to the physical momenta. The mapping is defined as,  $\Phi = \Pi(\theta)$ , and the integral in Eq. (1) is written,

$$\sigma = \int_{[0,1]^D} d\theta |\mathcal{M}|^2(\Pi(\theta)) \mathcal{J}(\theta) = \int_{[0,1]^D} d\theta \mathcal{L}(\theta). \quad (2)$$

A Jacobian associated with the change of coordinates between  $\theta$  and  $\Phi$  has been introduced,  $\mathcal{J}$ , and then absorbed into the definition of  $\mathcal{L}(\theta) = |\mathcal{M}|^2(\Pi(\theta)) \mathcal{J}(\theta)$ . With no general analytic solution to the sorts of scatterings considered

at the high energy frontier, this integral must be estimated with numerical techniques. Numerical integration involves sampling from the  $|\mathcal{M}|^2$  distribution in a manner that gives a convergent estimate of the true integral when the samples are summed. As a byproduct this set of samples can be used to estimate integrals of arbitrary sub-selections of the integrated phase space volume, decomposing the total cross section into differential cross section elements,  $d\sigma$ . Additionally these samples can be unweighted and used as pseudo-data to emulate the experimental observations of the collisions. The current state of the art techniques for performing these tasks were briefly reviewed in Sect. 1.

Importance Sampling (IS) is a Monte Carlo technique used extensively in particle physics when one needs to draw samples from a distribution with an unknown *target* probability density function,  $P(\Phi)$ . Importance Sampling approaches this problem by instead drawing from a known *sampling* distribution,  $Q(\Phi)$  (A number of standard texts for inference give more thorough exposition of the general sampling theory used in this paper, see e.g. [39]). Samples drawn from  $Q$  are assigned a weight,  $w = P(\Phi)/Q(\Phi)$ , adjusting the importance of each sampled point. The performance of IS rests heavily on how well the sampling distribution can be chosen to match the target, and adaptive schemes like VEGAS are employed to refine initial proposals. It is well established that as the dimensionality and complexity of the target increase, the task of constructing a viable sampling distribution becomes increasingly challenging.

Markov Chain based approaches fundamentally differ in that they employ a local sampling distribution and define an acceptance probability with which to accept new samples. Markov Chain Monte Carlo (MCMC) algorithms are widely used in Bayesian inference. Numerical Bayesian methods have to be able to iteratively refine the prior distribution to the posterior, even in cases where the two distributions are largely disparate, making stochastic MCMC refinement an indispensable tool in many cases. This is an important conceptual point; in the particle physics problems presented in this work we are sampling from exact theoretically derived distributions. The lack of noise and a priori well known structure make methods with deterministic proposal distributions such as IS more initially appealing, however at some point increasing the complexity and dimensionality of the problem forces one to use stochastic methods. Lattice QCD calculations are a prominent example set of adjacent problems sampling from theoretical distributions that make extensive use of MCMC approaches [40]. MCMC algorithms introduce an orthogonal set of challenges to IS; a local proposal is inherently simpler to construct, however issues with exploration of multimodal target distributions and autocorrelation of samples become new challenges to address.


Nested Sampling (NS) is a well established algorithm for numerical evaluation of high dimensional integrals [37]. NS

differs from typical MCMC samplers as it is primarily an integration algorithm, hence by definition has to overcome a lot of the difficulties MCMC samplers face in multimodal problems. A recent community review of its various applications in the physical sciences, and various implementations of the algorithm has been presented in [41].

At its core NS operates by maintaining a number,  $n_{\text{live}}$ , of *live point* samples. This ensemble of live points is initially uniformly sampled from  $\theta \in [0, 1]^D$  – distributed in the physical volume  $\Omega$  according to the shape of the mapping  $\Pi$ . These live points are sorted in order of  $\mathcal{L}(\theta)$  evaluated at the phase space point, and the point with the lowest  $\mathcal{L}$ ,  $\mathcal{L}_{\text{min}}$ , in the population is identified. A replacement for this point is found by sampling uniformly under a hard constraint requiring,  $\mathcal{L} > \mathcal{L}_{\text{min}}$ . The volume enclosed by this next iteration of live points has contracted and the procedure of identifying the lowest  $\mathcal{L}$  point and replacing it is repeated. An illustration of three different stages of this iterative compression on an example two-dimensional function are shown in Fig. 1. The example function used in this case has four identical local maxima to find, practical exploration and discovery of the modes is achieved by having a sufficient ( $\mathcal{O}(10)$ ) initial samples in the basis of attraction of each mode. This can either be achieved by brute force sampling a large number of initial samples, or by picking an initial mapping distribution that better reflects the multi-modal structure. By continually uniformly sampling from a steadily compressing volume, NS can estimate the density of points which is necessary for computing an integral as given by Eq. (1). Once the iterative procedure reaches a point where the live point ensemble occupies a predefined small fraction of the initial volume,  $T_C$ , the algorithm terminates. The fraction  $T_C$  can be characterised as the *termination criterion*. The discarded points throughout the evolution are termed *dead points* which can be joined with the remaining live points to form a representative sample of the function, that can be used to estimate the integral or to provide a random sample of events.

To estimate the integral and generate (weighted) random samples, Nested Sampling achieves this by probabilistically estimating the volume of the shell between the two outermost points as approximately  $\frac{1}{n_{\text{live}}}$  of the current live volume. The volume  $X_j$  within the contour  $\mathcal{L}_j$  – defined by the point with  $\mathcal{L}_{\text{min}}$  – at iteration  $j$  may therefore be estimated as,

$$\begin{aligned}
 X_j &= \int_{\mathcal{L}(\theta) > \mathcal{L}_j} d\theta \\
 \Rightarrow X_0 &= 1, \\
 P(X_j | X_{j-1}) &= \frac{X_j^{n_{\text{live}}-1}}{n_{\text{live}} X_{j-1}^{n_{\text{live}}}} \\
 \Rightarrow \log X_j &\approx \frac{-j \pm \sqrt{j}}{n_{\text{live}}}.
 \end{aligned}$$



**Fig. 1** Schematic of live point evolution (blue dots) in Nested Sampling, over a two-dimensional function whose logarithm is the negative Himmelblau function (contours). Points are initially drawn from the unit hypercube (top panel). The points on the lowest contours are successively deleted, causing the live points to contract around the peak(s) of the function. After sufficient compression is achieved, the dead points (orange) may be weighted to compute the volume under the surface and samples from probability distributions derived from the function

The cross section and probability weights can therefore be estimated as,

$$\begin{aligned}\sigma &= \int d\theta \mathcal{L}(\theta) = \int dX \mathcal{L}(X) \\ &\approx \sum_j \mathcal{L}_j \Delta X_j, \quad w_j \approx \frac{\Delta X_j \mathcal{L}_j}{\sigma}.\end{aligned}\quad (3)$$

Importantly, for all of the above the approximation signs indicate errors in the procedure of probabilistic volume estimation, which are fully quantifiable.

The method to sample new live points under a hard constraint can be realised in multiple ways, and this is one of the key differences in the various implementations of NS. In this work we employ the POLYCHORD implementation of Nested Sampling [38], which uses slice sampling [42] MCMC steps to evolve the live points. NS can be viewed as being an ensemble of many short Markov Chains.

Much of the development and usage of NS has focused on the problem of calculation of marginal likelihoods (or evidences) in Bayesian inference, particularly within the field of Cosmology [43–48]. We can define the Bayesian evidence,  $\mathcal{Z}$ , analogously to the particle physics cross section,  $\sigma$ . NS in this context evaluates the integral,

$$\mathcal{Z} = \int d\theta \mathcal{L}(\theta) \pi(\theta), \quad (4)$$

where the likelihood function,  $\mathcal{L}$ , plays a similar role to  $|\mathcal{M}|^2$ . In the Bayesian inference context, the phase space over which we are integrating,  $\theta$ , has a measure defined by the prior distribution,  $\pi(\theta)$ , which without loss of generality under a suitable coordinate transformation can be taken to be uniform over the unit hypercube. Making the analogy between the evidence and the cross section explicit will allow us to apply some of the information theoretic metrics commonly used in Bayesian inference to the particle physics context [49], and provide terminology used throughout this work. Among a wide array of sampling methods for Bayesian inference, NS possesses some unique properties that enable it to successfully compute the high dimensional integral associated with Eq. (4). These properties also bear a striking similarity to the requirements one would like to have to explore particle physics phase spaces. These are briefly qualitatively described as follows:

- NS is primarily a *numerical integration method that produces posterior samples as a by product*. In this respect it is comfortably similar to Importance Sampling as the established tool in particle physics event generation. It might initially be tempting to approach the particle physics event generation task purely as a posterior sampling problem. Standard Markov Chain based sampling

tools cannot generically give good estimates of the integral, so are not suited to compute the cross section. Additionally issues with coverage of the full phase space from the resulting event samples are accounted for by default by obtaining a convergent estimate of the integral over all of the phase space.

- NS naturally *handles multimodal problems* [45,46]. The iterative compression can be augmented by inserting steps that cluster the live points periodically throughout the run. Defining subsets of live points and evolving them separately allows NS to naturally tune itself to the modality of unseen problems.
- NS requires a construction that can handle *sampling under a hard likelihood constraint* in order to perform the compression of the volume throughout the run. Hard boundaries in the physics problem, such as un-physical or deliberately cut phase space regions, manifest themselves in the sampling space as a natural extension of these constraints.
- NS is *largely self-tuning*. Usage in Bayesian inference has found that NS can be applied to a broad range of problems with little optimisation of hyper-parameters necessary [50–52]. NS can adapt to different processes in particle physics *without any prior knowledge of the underlying process needed*.

The challenge to present NS in this new context is to find an even comparison of sampling performance between NS and IS. It is typical in phase space sampling to compare the difference between the target and the sampling distribution as reducing the variation between these two distributions gives a clear metric of performance for IS. For NS there is no such global sampling distribution; the closest analogue being the prior which is then iteratively refined with local proposals to an estimate of the target. In Sect. 2.1 we attempt to compare the sampling distribution between NS and IS using a toy problem, however in the full physical gluon scattering example presented in Sect. 3 we instead focus directly on the properties of the estimated target distribution as this is the most direct equitable point of comparison.

## 2.1 Illustrative example

To demonstrate the capabilities of NS we apply the algorithm to an illustrative sampling problem in two dimensions. Further examples validating POLYCHORD on a number of challenging sampling toy problems are included in the original paper [38], here we present a modified version of the *Gaussian Shells* scenario. An important distinction of the phase space use case not present in typical examples is the emphasis on calculating finely binned *differential* histograms of the total integral. As a comparison to NS, we sample the same problem with a method that is well-known in high energy


physics – adaptive Importance Sampling (IS), realised using the VEGAS algorithm.

For our toy example we introduce a “stop sign” target density, whose unnormalised distribution is defined by


$$\begin{aligned}
 f(x, y) = & \frac{1}{2\pi^2} \frac{\Delta r}{\left(\sqrt{(x - x_0)^2 + (y - y_0)^2} - r_0\right)^2 + (\Delta r)^2} \\
 & \cdot \frac{1}{\sqrt{(x - x_0)^2 + (y - y_0)^2}} \\
 & + \frac{1}{2\pi r_0} \frac{\Delta r}{((y - y_0) - (x - x_0))^2 + (\Delta r)^2} \\
 & \cdot \Theta\left(r_0 - \sqrt{(x - x_0)^2 + (y - y_0)^2}\right), \quad (5)
 \end{aligned}$$

where  $\Theta(x)$  is the Heaviside function. It is the sum of a ring and a line segment, both with a (truncated) Cauchy profile. The ring is centred at  $(x_0, y_0) = (0.5, 0.5)$  and has a radius of  $r_0 = 0.4$ . The line segment is located in the inner part of the ring and runs through the entire diameter. We set the width of the Cauchy profile to  $\Delta r = 0.002$ . This distribution can be seen as an example of a target where it makes sense to tackle the sampling problem with a multi-channel distribution. One channel could be chosen to sample the ring in polar coordinates and one to sample the line segment in Cartesian coordinates. However, here we deliberately use VEGAS as a single channel in order to highlight the limitations of the algorithm. From the perspective of a single channel, there is no coordinate system to factorise the target distribution. That poses a serious problem for VEGAS, as it uses a factorised sampling distribution where the variables are sampled individually. Both algorithms are given zero prior knowledge of the target, thus starting with a uniform prior distribution.

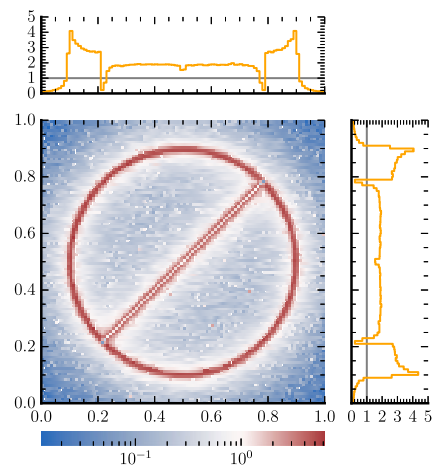
Our VEGAS grid has 200 bins per dimension. We train it over 10 iterations where we draw 30k points from the current VEGAS mapping and adapt the grid to the data. The distribution defined by the resulting grid is then used for IS without further adaptation. This corresponds to the typical use in an event generator, where there is first an integration phase in which, among other things, VEGAS is adapted, followed by a non-adaptive event generation phase. We note that VEGAS gets an advantage in this example comparison as we do not include the target evaluations from the training into the counting. However, it should be borne in mind that in a realistic application with a large number of events to be generated, the costs for training are comparatively low. For NS we use POLYCHORD with a number of live points  $n_{\text{live}} = 1000$  and a chain length  $n_{\text{repeats}} = 4$ , more complete detail of POLYCHORD settings and their implication are given in Sect. 3.1. Figure 2a shows the bivariate target distribution along with the marginal  $x$  and  $y$  distributions of the target, VEGAS and POLYCHORD. For this plot (as well as for Fig. 2c) we merged 70 independent runs of POLYCHORD to get a better visual



(a) target



(b) ratio target/VEGAS sampling density




(c) ratio target/POLYCHORD sampling density

**Fig. 2** A two-dimensional toy example: **a** Histogram of the target function along with the marginal sampling distributions of VEGAS and POLYCHORD. **b** Ratio of the target function and the probability density function of VEGAS. **c** Ratio of the target density to the sampling density of POLYCHORD

representation due to the larger sample size. It can be seen that both algorithms reproduce the marginal distributions reasonably well. There is some mismatch at the boundaries for VEGAS. This can be explained by the fact that VEGAS, as a variance-reduction method, focuses on the high-probability regions, where it puts many bins, and uses only few bins for the comparably flat low-probability regions. As a result, the bins next to the boundaries are very wide and overestimate the tails. POLYCHORD also oversamples the tails, reflecting the fact that in this example the prior is drastically different from the posterior, meaning the initial phase of prior sampling in POLYCHORD is very inefficient. In addition it puts too many points where the ring and the line segment join, which is where we find the highest values of the target function. This is not a generic feature of NS at the termination of the algorithm, rather it reflects the nature of having two intersecting sweeping degenerate modes in the problem, a rather unlikely scenario in any physical integral.

Figure 2b shows the ratio between the target distribution and the sampling distribution of VEGAS, representing the IS weights. It can be seen that the marginals of the ratio are relatively flat, with values between 0.1 and 5.7. However, in two dimensions the ratio reaches values up to  $1 \times 10^{-2}$ . By comparing Fig. 2a and b, paying particular attention to the very similar ranges of function values, it can be deduced that VEGAS almost completely misses to learn the structure of the target. It tries to represent the peak structure from the ring and the line segment by an enclosing square with nearly uniform probability distribution.

The same kind of plot is shown in Fig. 2c for the POLYCHORD data. NS does not strictly define a sampling distribution, however a proxy for this can be visualised by plotting the density of posterior samples. Here the values of the ratio are much smaller, between  $1 \times 10^{-2}$  and 7. POLYCHORD produces a flatter ratio function than VEGAS while not introducing additional artifacts that are not present in the original function. The smallest/largest values of the ratio are found in the same regions as the smallest/largest values of the target function, implying that POLYCHORD tends to overestimate the tails and to underestimate the peaks. This can be most clearly explained by examining the profile of where posterior mass is distributed throughout a run, an important diagnostic tool for NS runs [53]. It is shown in Fig. 3, where the algorithm runs from left to right; starting with the entire prior volume remaining enclosed by the live points,  $\log X = 0$ , and running to termination, when the live points contain a vanishingly small remaining prior volume. The posterior mass profile, shown in blue, is the analogue to the sampling density in VEGAS. To contextualise this against the target function, a profile of the log-likelihood of the lowest live point in the live point ensemble is similarly shown as a function of the



**Fig. 3** Likelihood ( $\log \mathcal{L}$ ) and posterior mass ( $\mathcal{L}X$ ) profiles for a run of POLYCHORD on the example target density. The  $x$ -axis tracks the prior volume remaining as the run progresses, with  $\log X = 0$  corresponding to the start of the run, with the algorithm compressing the volume from left to right, where the run terminates

remaining prior volume,  $X$ . Nested Sampling can be motivated as a *likelihood scanner*, sampling from monotonically increasing likelihood shells. These two profiles indicate some features of this problem, firstly a *phase transition* is visible in the posterior mass profile. This occurs when the degenerate peak of the ring structure is reached, the likelihood profile reaches a plateau where the iterations kill off the degenerate points at the peak of the ring, before proceeding to scan up the remaining line segment feature. An effective second plateau is found when the peak of the line segment is reached, with a final small detail being the superposition of the ring likelihood on the line segment. Once the live points are all occupying the extrema of the line segment, there is a sufficiently small prior volume remaining that the algorithm terminates. The majority of the posterior mass, and hence sampling density is distributed around the points where the two peaks are ascended. This reflects the stark contrast between the prior initial sampling density and the target, the samples are naturally distributed where the most information is needed to effectively compress the prior to the posterior.

We compare the efficiencies of the two algorithms for the generation of equal-weight events in Table 1. It shows that POLYCHORD achieves an overall efficiency of  $\epsilon = 0.0113(90)$  which is almost three times as high as the efficiency of VEGAS. While for VEGAS the overall efficiency  $\epsilon$  is identical to the unweighting efficiency  $\epsilon_{uw}$ , determined by the ratio of the average event weight over the maximal weight in the sample, for POLYCHORD we also have to take the slice sampling efficiency  $\epsilon_{ss}$  into account, which results from the thinning of the Markov Chain in the slice sampling step. Here, the total efficiency  $\epsilon = \epsilon_{ss}\epsilon_{uw}$  is dominated by the slice sampling efficiency. We point out that it is in the nature of the NS algorithm that the sample size is not deterministic. However, the variance is not very large and it is easily possible to merge several NS runs to obtain a larger sample.

**Table 1** Comparison of VEGAS and NS for the toy example in terms of size of event samples produced.  $N_{\mathcal{L}}$  gives the number of target evaluations,  $N_W$  the number of weighted events and  $N_{\text{equal}}$  the derived number of equal weight events. A MC slice sampling efficiency,  $\epsilon_{\text{ss}}$ , is listed for

| Algorithm | $N_{\mathcal{L}}$ | $\epsilon_{\text{ss}}$ | $N_W$        | $\epsilon_{\text{uw}}$ | $N_{\text{equal}}$ | $\epsilon$ |
|-----------|-------------------|------------------------|--------------|------------------------|--------------------|------------|
| VEGAS     | 300,000           |                        | 300,000      | 0.004(2)               | 1267 (460)         | 0.004 (2)  |
| NS        | 308,755 (17505)   | 0.041 (3)              | 12,669 (147) | 0.273 (7)              | 3462 (96)          | 0.0113 (9) |

**Table 2** Comparison of integrals calculated in the toy example with VEGAS and NS, along with the respective uncertainties

| Algorithm | I    | $\Delta\sigma_{\text{tot}}$ | $\Delta_w$ | $\Delta_{\text{MC}}$ |
|-----------|------|-----------------------------|------------|----------------------|
| VEGAS     | 1.71 | 0.02                        |            | 0.02                 |
| NS        | 1.65 | 0.05                        | 0.04       | 0.02                 |

Table 2 shows the integral estimates along with the corresponding uncertainty measures. While the pure Monte Carlo errors are of the same size for both algorithms, there is an additional uncertainty for NS. It carries an uncertainty on the weights of the sampled points, listed as  $\Delta_w$ . This arises due to the nature of NS using the volume enclosed by the live points at each iteration to estimate the volume of the likelihood shell. The variance in this volume estimate can be sampled, which is reflected as a sample of alternative weights for each dead point in the sample. Summing up these alternative weight samples gives a spread of predictions for the total integral estimate, and the standard deviation of these is quoted as  $\Delta_w$ . This additional uncertainty compounds the familiar statistical uncertainty, listed as  $\Delta_{\text{MC}}$  for all calculations. In Appendix A, we present the procedure needed to combine the two NS uncertainties to quote a total uncertainty,  $\Delta\sigma_{\text{tot}}$ , as naively adding in quadrature will overestimate the true error.

### 3 Application to gluon scattering

As a first application and benchmark for the Nested Sampling algorithm, we consider partonic gluon scattering processes into three-, four- and five-gluon final states at fixed centre-of-mass energies of  $\sqrt{s} = 1$  TeV. These channels have a complicated phase space structure that is similar to processes with quarks or jets, while the corresponding amplitude expressions are rather straightforward to generate. The fixed initial and final states allow us to focus on the underlying sampling problem. For regularisation we apply cuts to the invariant masses of all pairs of final state gluons such that  $m_{ij} > 30$  GeV and on the transverse momenta of all final state gluons such that  $p_{T,i} > 30$  GeV. The renormalisation scale is fixed to  $\mu_R = \sqrt{s}$ . The matrix elements are calculated

NS. A total,  $\epsilon$ , and unweighting,  $\epsilon_{\text{uw}}$ , efficiency are listed for both algorithms. We report the mean and standard deviation of ten independent runs of the respective algorithm

using a custom interface between POLYCHORD and the matrix element generator AMEGIC [5] within the SHERPA event generator framework [54]. Three established methods are used to provide benchmarks to compare NS to. Principle comparison is drawn to the HAAG sampler, optimised for QCD antenna structures [55], illustrating the exploration of phase space with the best a priori knowledge of the underlying physics included. It uses a cut-off parameter of  $s_0 = 900$  GeV<sup>2</sup>. Alongside this, two algorithms that will input no prior knowledge of the phase space, i.e. the integrand, are used; adaptive importance sampling as realised in the VEGAS algorithm [8] and a flat uniform sampler realised using the RAMBO algorithm [56,57]. VEGAS remaps the variables of the RAMBO parametrisation using 50, 70, 200 bins per dimension for the three-, four-, and five-gluon case, respectively. The grid is trained in 10 iterations using 100k training points each. Note, the dimensionality of the phase space for  $n$ -gluon production is  $D = 3n - 4$ , where total four-momentum conservation and on-shell conditions for the external particles are implicit.

As a first attempt to establish NS in this context, we treat the task of estimating the total and differential cross sections of the three processes *starting with no prior knowledge* of the underlying phase space distribution. For the purposes of running POLYCHORD we provide the flat RAMBO sampler as the prior, and the likelihood function provided is the squared matrix element. In contrast to HAAG, POLYCHORD performs the integration without any decomposition into channels, removing the need for any multichannel mapping. NS is a flexible procedure, and the objective of the algorithm can be modified to perform a variety of tasks, a recent example has presented NS for computation of small  $p$ -values in the particle physics context [58]. To establish NS for the task of phase space integration in this study, a standard usage of POLYCHORD is employed, mostly following default values used commonly in Bayesian inference problems.

The discussion of the application of NS to gluon-scattering processes is split into four parts. Firstly, the hyperparameters and general setup of POLYCHORD are explained in Sect. 3.1. In Sect. 3.2 a first validation of NS performing the core tasks of (differential) cross-section estimation from weighted events – against the HAAG algorithm – is presented. In Sect. 3.3 further information is given to contextualise the computational efficiency of NS against the alter-

**Table 3** POLYCHORD hyperparameters used for this analysis, parameters not listed follow the POLYCHORD defaults

| Parameter               | POLYCHORD name     | Value                     | Description                                   |
|-------------------------|--------------------|---------------------------|-----------------------------------------------|
| Number of dimensions    | $n_{\text{dim}}$   | [5, 8, 11]                | Dimension of sampling space                   |
| Number of live points   | $n_{\text{live}}$  | 10,000                    | Resolution of the algorithm                   |
| Number of repeats       | $n_{\text{rep}}$   | $n_{\text{dim}} \times 2$ | Length of Markov chains                       |
| Number of prior samples | $n_{\text{prior}}$ | $n_{\text{live}}$         | Number of initial samples from prior          |
| Boost posterior         |                    | $n_{\text{rep}}$          | Write out maximum number of posterior samples |

native established tools for these tasks. Finally a consideration of unweighted event generation with NS is presented in Sect. 3.4.

### 3.1 POLYCHORD hyperparameters

The hyperparameters chosen to steer POLYCHORD are listed in Table 3. These represent a typical set of choices for a high resolution run with the objective of producing a large number of posterior samples. The number of live points is one of the parameters that is most free to tune, being effectively the resolution of the algorithm. Typically  $n_{\text{live}}$  larger than  $\mathcal{O}(1000)$  gives diminishing returns on accuracy, Bayesian inference usage in particle physics has previously employed  $n_{\text{live}} = 4000$  [59] to provide some context for the choice made in this work. The particular event generation use case, partitioning the integral into arbitrarily small divisions (differential cross sections), logically favours a large  $n_{\text{live}}$  (resolution). The number of repeats is a parameter that controls the length of the slice sampling chains, the value chosen is the recommended default for reliable posterior sampling, whereas  $n_{\text{rep}} = n_{\text{dim}} \times 5$  is recommended for evidence (total integral) estimation. As this study aims to cover both differential and total cross sections, the smaller value is favoured as there is a strong limit on the overall efficiency imposed by how many samples are needed to decorrelate the Markov Chains.

An important point to note is in how POLYCHORD treats unphysical values of the phase space variables, e.g. if they fall outside the fiducial phase space defined by cuts on the particle momenta. This is not an explicit hyperparameter of POLYCHORD, rather how the algorithm treats points with zero likelihood. In both the established approaches and in POLYCHORD the sampling is performed in the unit hypercube, which is then translated to the physical variables which can be evaluated for consistency and rejected if they are not phys-

ically valid. One of the strengths of NS is that the default behavior is to consider points which return zero likelihood<sup>1</sup> as being excluded at the prior level. During the initial prior sampling phase, unphysical points are set to log-zero and the sampling proceeds until  $n_{\text{prior}}$  initial physical samples have been obtained. Provided each connected physical region contains some live points after this initial phase, the iterative phase of MCMC sampling will explore up to the unphysical boundary. This effect necessitates a correction factor to be applied to the integral, derived as the ratio of total initial prior samples to the physically valid prior samples. In practice the correction factor is found in the `prior_info` file written out by POLYCHORD. An uncertainty on this correction can be derived from order statistics [60], however it was found to be negligibly small for the purposes of this study so is not included.

Another standout choice of hyperparameter is the chosen value of  $n_{\text{prior}}$ . The number of prior samples is an important hyperparameter that would typically be set to some larger multiple of  $n_{\text{live}}$  in a Bayesian inference context,  $n_{\text{prior}} = 10 \times n_{\text{live}}$  would be considered sensible for a broad range of tasks. For the purpose of generating weighted events, using a larger value would generally be advantageous, however increasing  $n_{\text{prior}}$  will strongly decrease the efficiency in generating *unweighted* events. As the goal is to construct a generator taking an uninformed prior all the way through to unweighted events, the default value listed is used. However it is notable that this is a particular feature of starting from an uninformed prior, if more knowledge were to be included in the prior then a longer phase of prior sampling becomes advantageous. The final parameter noted, the factor by which to boost posterior samples, has no effect on POLYCHORD at runtime. Setting this to be equal to the number of repeats simply writes out the maximum number of dead points, hence is needed in this scenario. All plots and tables in the remainder of this section are composed of one single run of POLYCHORD with these settings, with the additional entries in Table 4 demonstrating a join of ten such runs.

### 3.2 Exploration and integrals

Before examining the performance of NS in detail, it is first important to validate that the technique is capable of fully exploring particle physics phase spaces in these chosen examples. The key test to validate this is to compare if various differential cross sections calculated with NS are statistically consistent with the established techniques. To do this, a single NS and HAAG sample of weighted events is produced, using approximately similar levels of computational

<sup>1</sup> Since POLYCHORD operates in log space, to avoid the infinity associated with  $\log(0)$ , log-zero is defined as a settable parameter. By default this is chosen to  $-1 \times 10^{-25}$ .



**Table 4** Comparison of integrals calculated for the three-, four- and five-gluon processes using RAMBO, VEGAS, NS and HAAG, along with the respective uncertainties

| Process | Algorithm          | $\sigma$ | $\Delta\sigma_{\text{tot}}$ | $\Delta_w$ | $\Delta_{\text{MC}}$ |
|---------|--------------------|----------|-----------------------------|------------|----------------------|
| 3-jet   | RAMBO              | 24.580   | 0.191                       |            | 0.191                |
|         | VEGAS              | 24.807   | 0.017                       |            | 0.017                |
|         | NS                 | 24.669   | 0.467                       | 0.484      | 0.100                |
|         | NS ( $\times 10$ ) | 24.888   | 0.145                       | 0.150      | 0.030                |
|         | HAAG               | 24.840   | 0.017                       |            | 0.017                |
| 4-jet   | RAMBO              | 9.876    | 0.107                       |            | 0.107                |
|         | VEGAS              | 9.849    | 0.009                       |            | 0.009                |
|         | NS                 | 9.837    | 0.194                       | 0.196      | 0.036                |
|         | NS ( $\times 10$ ) | 9.778    | 0.064                       | 0.066      | 0.011                |
|         | HAAG               | 9.853    | 0.006                       |            | 0.006                |
| 5-jet   | RAMBO              | 2.644    | 0.024                       |            | 0.024                |
|         | VEGAS              | 2.680    | 0.003                       |            | 0.003                |
|         | NS                 | 2.612    | 0.051                       | 0.048      | 0.009                |
|         | NS ( $\times 10$ ) | 2.667    | 0.017                       | 0.017      | 0.003                |
|         | HAAG               | 2.685    | 0.001                       |            | 0.001                |


overhead (more detail on this is given in Sect. 3.3). Both sets of weighted events are analysed using the default MC\_JETS RIVET routine [61]. RIVET produces binned differential cross sections as functions of various physical observables of the outgoing gluons. For each process, the total cross section for the NS sample is normalised to the HAAG sample, and a range of fine grained differential cross sections is calculated using both algorithms covering the following observables;  $\eta_i, y_i, p_{T,i}, \Delta\phi_{ij}, m_{ij}, \Delta R_{ij}, \Delta\eta_{ij}$ , where  $i \neq j$  label the final state jets, reconstructed using the anti- $k_T$  algorithm [62] with a radius parameter of  $R = 0.4$  and  $p_T > 30$  GeV. The normalised difference between the NS and HAAG differential cross section in each bin can be computed as,

$$\chi = \frac{d\sigma_{\text{HAAG}} - d\sigma_{\text{NS}}}{\sqrt{\Delta_{\text{HAAG}}^2 + \Delta_{\text{NS}}^2}}, \tag{6}$$

in effect this is the differences between the two algorithms normalised by the combined standard deviation. By summing up this  $\chi$  deviation across all the available bins in each process, a test to see if the two algorithms are convergent within their quoted uncertainties can be performed. Since over 500 bins are populated and considered in each process, it is expected that the rate of these  $\chi$  deviations should be approximately normally distributed. This indeed appears to hold, and these summed density estimates across all observables are shown in Fig. 4, alongside an overlaid normal distribution with mean zero and variance one,  $\mathcal{N}(0, 1)$ , to illustrate the expected outcome. Two example variables that were used to build this global deviation are also shown; the leading jet

$p_T$  in Fig. 5a and  $\Delta R_{12}$ , the distance of the two leading jets in the  $(\eta, \phi)$  plane, in Fig. 5b.

The composition of the quoted uncertainty for the two algorithms differs, demonstrating an important feature of an NS calculation. For HAAG, and IS in general, it is conventional to quote the uncertainty as the standard error from the effective number of fills in a bin. Nested Sampling on the other hand introduces an uncertainty on the weights used to fill the histograms themselves, effectively giving rise to multiple weight histories that must be sampled to derive the correct uncertainty on the NS calculation. Details on this calculation are supplied in Appendix A. In summary the alternative weight histories give an overlapping measure of the statistical uncertainty, so this effect must be accounted for in situ alongside taking the standard deviation of the weight histories. To contextualise this, the middle panels in Fig. 5 show the correct combined uncertainty (using the recipe from Appendix A) as a grey band, against the bands derived from the standard error of each individual algorithm (henceforth  $\Delta_{\text{MC}}$ ) as dashed lines, and the complete NS error treatment as a dotted line. The standard error (dashed) NS band in these panels is a naive estimation of the full NS uncertainty (dotted), however this illustrates an important point; at the level of fine grained differential observables the NS uncertainty is dominated by statistics and is hence reducible as one would expect by repeated runs. Based on the example observables we can initially conclude that whilst both algorithms appear compatible, when using weighted events NS generally has a larger uncertainty than HAAG across most of the range (given a roughly equivalent computational overhead). However, further inspection of the resulting unweighted event samples



**Fig. 4** Global rate of occurrence of per bin deviation,  $\chi$ , between HAAG and NS, for each considered scattering process. A normally distributed equivalent deviation rate is shown for comparison

derived from these weighted samples in the remaining sections reveals a more competitive picture between the two algorithms.

The estimates of the total cross sections, derived from the sum of weighted samples, provided in Table 4, give an alternative validation that NS is sufficiently exploring the phase space by ensuring that compatible estimates of the cross sections are produced between all the methods reviewed in this study. The central estimates of the total cross sections are generally consistent within the established error sources for all calculations considered. In this table the components of the error calculation for NS are listed separately;  $\Delta_w$  being the standard deviation resulting from the alternative weight histories and  $\Delta_{MC}$  being the standard error naively taken from the mean of the alternative NS weights. In contrast to the differential observables, the naive counting uncertainty is small so has negligible effect at the level of total cross sections. In summary, for a total cross section the spread of alternative weight histories gives a rough estimate of the total error, whereas for a fine grained differential cross section the standard error dominates. The way to correctly account for the effect of counting statistics within the weight histories is given in Appendix A.


Repeated runs of NS will reduce these uncertainties. The anesthetic package [63] is used to analyse the NS runs throughout this paper, and contains a utility to join samples. Once samples are joined consistently into a larger sample, the uncertainties can be derived as already detailed. The result of joining 10 equivalent NS runs with the previously motivated hyperparameters is also listed in Table 4. Joining 10 runs affects the  $\Delta\sigma_{tot}$  for NS in two ways; reducing the spread of weighted sums composing  $\Delta_w$  (i.e. reducing  $\Delta_{MC}$ ), and reducing the variance of distribution for each weight itself

(i.e. the part of  $\Delta_w$  that does not overlap with  $\Delta_{MC}$ ). The former is reduced by simply having an increased size of samples produced, increasing the number of effective fills by a factor of  $\sim 10$  in this case, with the latter reduced due to the increased effective number of live points used for the volume estimation.


### 3.3 Efficiency of event generation

An example particle physics workflow on this gluon scattering problem would be to take HAAG as an initial mapping of the phase space (effectively representing the best prior knowledge of the problem), and using VEGAS to refine the proposal distribution to optimally efficiently generate weighted events. Of the three existing tools presented in this study for comparison (HAAG, RAMBO, and VEGAS), NS bears most similarity to VEGAS, in that both algorithms learn the structure of the target integrand. To this end an atypical usage of VEGAS is employed, testing how well VEGAS could learn a proposal distribution from an uninformed starting point (RAMBO). This is equivalent to how NS was employed, starting from an uninformed prior (RAMBO) and generating posterior samples via Nested Sampling. It was motivated so far that roughly similar computational cost was used for the previous convergence checks, and that the hyperparameters of POLYCHORD were chosen to emphasise efficient generation of unweighted events. In what follows, we analyse more precisely this key issue of computational efficiency.

The statistics from a single run of the four algorithms for the three selected processes is listed in Table 5. NS is non deterministic in terms of number of matrix element evaluations ( $N_{\mathcal{L}}$ ), instead terminating from a pre determined convergence criterion of the integral. HAAG, VEGAS, RAMBO



(a) Differential cross section binned as a function of the leading jet transverse momentum, for the three considered processes.



(b) Differential cross section binned as a function of the separation of the leading two jets in the  $(\eta, \phi)$  plane, *i.e.*  $\Delta R_{12} = \sqrt{\Delta\eta_{12}^2 + \Delta\phi_{12}^2}$ , for the three considered processes.

**Fig. 5** Two example physical differential observables computed with weighted events using the HAAG and NS algorithms. The top panels show the physical distributions, the middle panels display the relative component error sources, and the bottom panel displays the normalised

deviation. The deviation plot has been normalised such that  $\chi = 1$  corresponds to an expected  $1\sigma$  deviation of a Gaussian distribution. Note that for illustrative purposes the cross sections for the four- and five-gluon processes have been scaled by global factors

**Table 5** Comparison of the four algorithms for the three processes in terms of size of event samples produces.  $N_{\mathcal{L}}$  gives the number of matrix element evaluations,  $N_W$  the number of weighted events,  $N_{W,\text{eff}}$  the effective number of weighted events and  $N_{\text{equal}}$  the derived number

of equal-weight events. A MC slice sampling efficiency,  $\epsilon_{\text{ss}}$ , is listed for NS. A total,  $\epsilon$ , and an unweighting,  $\epsilon_{\text{uw}}$ , efficiency are listed for all algorithms

| Process | Algorithm | $N_{\mathcal{L}} (\times 10^6)$ | $\epsilon_{\text{ss}}$ | $N_W (\times 10^6)$ | $\epsilon_{\text{uw}}$ | $N_{\text{equal}} (\times 10^6)$ | $\epsilon$ |
|---------|-----------|---------------------------------|------------------------|---------------------|------------------------|----------------------------------|------------|
| 3-jet   | RAMBO     | 10.00                           |                        | 10.00               | 0.0001                 | 0.001                            | 0.0001     |
|         | VEGAS     | 10.00                           |                        | 10.00               | 0.02                   | 0.20                             | 0.02       |
|         | NS        | 6.43                            | 0.03                   | 0.17                | 0.37                   | 0.06                             | 0.01       |
|         | HAAG      | 10.00                           |                        | 10.00               | 0.03                   | 0.29                             | 0.03       |
| 4-jet   | RAMBO     | 10.00                           |                        | 10.00               | 0.00003                | 0.0003                           | 0.00003    |
|         | VEGAS     | 10.00                           |                        | 10.00               | 0.005                  | 0.049                            | 0.005      |
|         | NS        | 7.94                            | 0.02                   | 0.19                | 0.43                   | 0.08                             | 0.01       |
|         | HAAG      | 10.00                           |                        | 10.00               | 0.02                   | 0.23                             | 0.02       |
| 5-jet   | RAMBO     | 10.00                           |                        | 10.00               | 0.00004                | 0.0004                           | 0.00004    |
|         | VEGAS     | 10.00                           |                        | 10.00               | 0.001                  | 0.013                            | 0.001      |
|         | NS        | 9.17                            | 0.02                   | 0.19                | 0.44                   | 0.08                             | 0.01       |
|         | HAAG      | 10.00                           |                        | 10.00               | 0.03                   | 0.25                             | 0.03       |

are all used to generate exactly 10M weighted events. The chosen POLYCHORD hyperparameters roughly align the NS method with the other three in terms of computational cost. One striking difference comes from the Markov Chain nature of NS. Default usage only retains a fraction of the total  $\mathcal{L}$  evaluations, inversely proportional to  $n_{\text{rep}}$ . This results in a smaller number of retained weighted events,  $N_W$ , than the number of  $\mathcal{L}$  evaluations,  $N_{\mathcal{L}}$ , for NS. However the retained weighted events by construction match the underlying distribution much closer than the other methods, resulting in a higher unweighting efficiency,  $\epsilon_{\text{uw}}$ , for the NS sample. Exact equal-weight unweighting can be achieved by accepting events with a probability proportional to the share of the sample weight they carry, this operation is performed for all samples of weighted events and the number of retained events is quoted as  $N_{\text{equal}}$ . NS as an unweighted event generator has some additional complexity due to the uncertainty in the weights themselves, this is given more attention in Sect. 3.4.

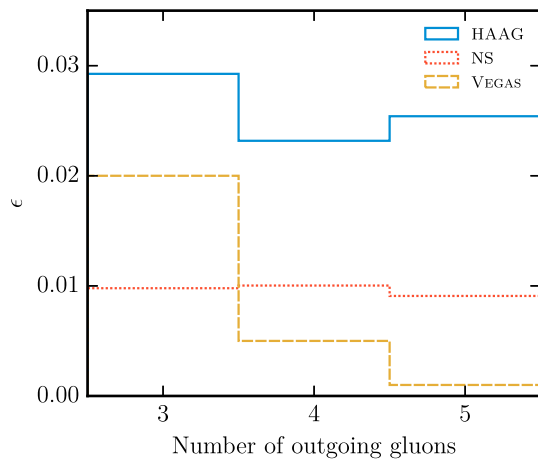
Due to differences in  $N_{\mathcal{L}}$  between NS and the other methods, it is most effective to compare the total efficiency in producing unweighted events,  $\epsilon = N_{\text{equal}}/N_{\mathcal{L}}$ . RAMBO as the baseline illustrates the performance one would expect, inputting no prior knowledge and not adapting to any acquired knowledge. As such RAMBO yields a tiny  $\epsilon$ . HAAG represents the performance using the best state of prior knowledge but without any adaptation, in these tests this represents the best attainable  $\epsilon$ . VEGAS and NS start from a similar point, both using RAMBO as an uninformed state of prior knowledge, but adapting to better approximate the phase space distribution as information is acquired. VEGAS starts with a higher efficiency than NS for the 3-gluon process, but the VEGAS efficiency drops by approximately an

order of magnitude as the dimensionality of phase space is increased to the 5-gluon process. NS maintains a consistent efficiency of approximately a percent, competitive with the consistent approximately three percent efficiency obtained by HAAG.

As the key point of comparison for this issue is the efficiency,  $\epsilon$ , this is highlighted with an additional visualisation in Fig. 6. The scaling behavior of the efficiency of each algorithm as a function of the number of outgoing gluons (corresponding to an increase in phase space dimensionality) is plotted for NS, HAAG and VEGAS. From the same starting point, NS and VEGAS can both learn a representation of the phase space, and do so in a way that yields a comparable efficiency to the static best available prior knowledge in HAAG. As the dimensionality of the space increases it appears that VEGAS starts to suffer in how accurately it can learn the mapping, however NS is still able to learn the mapping in a consistently efficient manner.

### 3.4 Unweighted event generation

The fact that NS leads to a set of alternative weight histories poses a technical challenge in operating as a generator of unweighted events in the expected manner. Exact unweighting, compressing the weighted sample to strictly equally weighted events leads to a different set of events being accepted for each weight history. Representative yields of unweighted events can be calculated as shown in Table 5 using the mean weight for each event, but the resulting differential distributions will underestimate the uncertainty if this is quoted simply as the standard error in the bin, as described in Appendix A. The correct uncertainty recipe can be propagated through naively, by separately unweight-



**Fig. 6** Visualisation of the efficiencies listed in Table 5

ing each weight history, however this requires saving as many event samples as required weight variations. Partial unweighting is commonly used in HEP event generation to allow a slight deviation from strict unit weights, to increase efficiency in practical settings. A modification to the partial unweighting procedure could be used to propagate the spread of weights to variations around accepted, approximate unit weight, events.

To conclude the exploration of the properties of NS as a generator for particle physics, a representative physical distribution calculated from a sample of exact unit-weight events is shown in Fig. 7. This sample is derived from the same weighted sample described in Table 5 and previously presented as a weighted event sample in Fig. 5a. The full set of NS variation weights is used to calculate the mean weight for each event, which is used to unweight the sample, for the chosen observable this is a very reasonable approximation as the fine binning means the standard error is the dominant uncertainty. The range of the leading jet transverse momenta has been extended into the tail of this distribution by modifying the default RIVET routine. This distribution largely reflects the information about the total efficiency previously illustrated in Fig. 6, projected onto a familiar differential observable. The total efficiency,  $\epsilon$ , was noted as being approximately one percent from NS, compared to approximately three percent from HAAG across all processes. If the total number of matrix element evaluations,  $N_{\mathcal{L}}$ , were to be made equal across all algorithms and processes, the performance would be further consistent.

#### 4 Future research directions


Throughout Sect. 3, the performance of Nested Sampling in the context of particle physics phase space sampling and event generation was presented. A single choice of hyperpa-

rameters was made, effectively performing a single NS run as an entire end-to-end event generator; starting from zero knowledge of the phase space all the way through to generating unweighted events. Simplifying the potential options of NS to a single version of the algorithm was a deliberate choice to more clearly illustrate the performance of NS in this new context, using the same settings for multiple tasks gives multiple orthogonal views on how the algorithm performs. However this was a limiting choice, NS has a number of variants and applications that could more effectively be tuned to a subset of the tasks presented. Some of the possible simple alterations – such as increasing  $n_{\text{prior}}$  to improve weighted event generation at the expense of unweighting efficiency – were motivated already in this paper. In this section we outline four broad topics that extend the workflow presented here, bringing together further ideas from the worlds of Nested Sampling and phase space exploration.

##### 4.1 Physics challenges in event generation

The physical processes studied in this work, up to 5-gluon scattering problems, are representative of the complexity of phase space calculation needed for the current precision demands of the LHC experiment collaborations [64]. However part of the motivation for this work, and indeed the broader increased interest in phase space integration methods, is due to the impending breaking point current pipelines face under the increased precision challenges of the HL-LHC programme. Firstly we observe that the phase space dimensionality of the highest multiplicity process studied here is 11. In broader Bayesian inference terms this is rather small, with NS being typically used for problems  $\mathcal{O}(10)$  to  $\mathcal{O}(100)$  dimensions, where it is uniquely able to perform numerical integration without approximation or strictly matching prior knowledge. The POLYCHORD implementation is styled as *next-generation* Nested Sampling, designed to have polynomial scaling with dimensionality aiming for robust performance as inference is extended to  $\mathcal{O}(100)$  dimensions. Earlier implementations of NS, such as MULTINEST [46], whilst having worse dimensional scaling properties, may be a useful avenue of investigation for the lower dimensional problems considered in this paper.

This work validated NS in a context where current tools still can perform the required tasks, albeit at times at immense computational costs. Requirements from the HL-LHC strain the existing LHC event generation pipeline in many ways and pushing the sampling problem to higher dimensions is no exception [2]. Importance Sampling becomes exponentially more sensitive to how close the proposal distribution matches the target in higher dimensions, a clear challenge for particle physics in two directions; multileg processes rapidly increasing the sampling dimension [65] and corresponding radiative corrections (real and virtual) make it increasingly



**Fig. 7** The equivalent leading jet transverse momentum observable as calculated in Fig. 5a, using an exact unit weight compression of the same samples. A modified version of the default MC\_JETS routine has been used to extend the  $p_T$  range shown

hard to provide an accurate proposal, e.g. through the sheer number of phase space channels needed and by having to probe deep into singular phase space regions [66]. We propose that NS is an excellent complement to further investigation on both these fronts. The robust dimensional scaling of NS illustrated against VEGAS in Fig. 6 encapsulates both solid performance with increasing dimension, and the adherence to an uninformed prior whilst still attaining this scaling is promising for scenarios where accurate proposals are harder to construct.

#### 4.2 Using prior knowledge

Perhaps the most obvious choice that makes the application here stylised is in always starting from an uninformed prior state of knowledge. Using Equations (2) and (4), the cross section integral with a phase space mapping was motivated as being exactly the Bayesian evidence integral with a choice of prior. To that end there is no real distinction between taking the non-uniform HAAG distribution as the prior instead of the flat RAMBO density that was used in this study. In this respect NS could be styled as learning an additional compression to the posterior distribution, refining the static proposal distributions typically employed to initiate the generation of a phase space mapping (noting that this is precisely what VEGAS aims to do in this context).

Naively applying a non-flat mapping exposes the conflicting aims at play in this set of problems however; efficiently generating events from a strongly peaked distribution, and generating high statistics estimates of the tails of the same distribution. Taking a flat RAMBO prior is well suited to the latter problem, whereas taking a HAAG prior is better suited to the former. One particular hyperparameter of POLYCHORD that was introduced can be tuned to this purpose; the number of prior samples,  $n_{\text{prior}}$ . If future work is to use a non flat, partially informed starting point, increasing  $n_{\text{prior}}$  well above the minimum (equal to the number of live points required) used in this study would be needed. A more complete direction for further work would be to investigate the possibility of mixing multiple proposal distributions [67,68].


As a demonstration, we again apply NS to the toy example of Sect. 2.1 but this time using a non-uniform prior distribution. While a good prior would be an approximation of the target distribution, we choose to purposely miss an important feature of the target, the straight line segment, that the sampler still has to explore. Considering that in HEP applications the prior knowledge may be encoded in the mixture distributions of a multi-channel importance sampler, this is an extreme version of a realistic situation. As typically the number of channels grows dramatically with increasing final-state particle multiplicity, e.g. factorially when channels correspond to the topologies of contributing Feynman diagrams, one might choose to disable some sub-dominant channels in

order to avoid a prohibitively large set of channels. However, this would lead to a mis-modelling of the target in certain phase-space regions.

Here we use only the ring part of the target, truncated on a circle that covers the unit hypercube, as our prior. Without an additional coordinate transformation this prior would not be of much use for VEGAS as the line part remains on the diagonal. To sample from the prior, we first transform to polar coordinates. Then we sample the angle uniformly and the radial coordinate using a Cauchy distribution truncated to the interval  $(0, 1/\sqrt{2}]$ . In order to have good coverage of the tails, despite the strongly peaked prior, we increase  $n_{\text{prior}}$  to  $50 \times n_{\text{live}}$ . This results in a total efficiency of  $\epsilon = 0.037(4)$ , more than three times the value obtained with a uniform prior, cf. Table 1. While the unweighting efficiency reduces to  $\epsilon_{\text{uw}} = 0.17(2)$ , the slice sampling efficiency increases to  $\epsilon_{\text{ss}} = 0.216(7)$ . In Fig. 8 we show the ratio between the target function and the POLYCHORD sampling distribution. Compared to Fig. 2c, the ratio has a smaller range of values. Along the peak of the ring part of the target function, the ratio is approximately one. The largest values can be found around the line segment with POLYCHORD generating up to ten times less samples than required by the target distribution. It can be concluded that even with an intentionally poor prior distribution, POLYCHORD benefits from the prior knowledge in terms of efficiency and still correctly samples the target distribution including the features absent from the prior.

### 4.3 Dynamic nested sampling

In addition to using a more informed prior to initiate the Nested Sampling process, a previous NS run can be used to further tune the algorithm itself to a particular problem. This is an existing idea in the literature known as dynamic Nested Sampling [69]. Dynamic NS uses information acquired about the likelihood shells in a previous NS run to varying the number of live points *dynamically* throughout the run. This results in a more efficient allocation of the computation towards the core aim of compressing the prior to the posterior. We expect that this would only increase the efficiency of the unweighting process, as the density of weighted events would be trimmed to even more closely match the underlying phase space density. Dynamic Nested Sampling naturally combines with the proposal of *using prior knowledge* to make a more familiar generator chain, however one that is driven primarily by NS. This mirrors the current established usage of VEGAS in this context; using VEGAS to refine the initial mapping by a redistribution of the input variables, to more efficiently generate from the acquired mapping.



**Fig. 8** The ratio of the target function of the two-dimensional toy example and the probability density function of POLYCHORD using a non-uniform prior distribution. Black histogram bins have not been filled by any data due to limited sample size

### 4.4 Connection to modern machine learning techniques

There has been a great deal of recent activity coincident to this work, approaching similar sets of problems in particle physics event generation using modern Machine Learning (ML) techniques [70]. Much of this work is still exploratory in nature, and covers such a broad range of activity that comprehensively reviewing the potential for combining ML and NS is beyond the scope of this work. It is however clear that there is strong potential to include NS into a pipeline that modern ML is already aiming to optimise. To that aim, we identify a particular technique that has been studied previously in the particle physics context; using Normalising Flows to train phase space mappings [29–31]. In spirit a flow based approach, training an invertible probabilistic mapping between prior and posterior, bears a great deal of similarity to the core compression idea behind Nested Sampling. The potential in dovetailing Nested Sampling with a flow based approach has been noted in the NS literature [71], further motivating the potential for synergy here.

The ability of NS to construct mappings of high dimensional phase spaces without needing any strong prior knowledge, can be motivated as being an ideal forward model with

which to train a Normalising Flow. In effect this replaces the generator part of the process with an importance sampler, whilst still using NS to generate the mappings. This is particularly ideal in this context, as the computational overhead required to decorrelate the Markov Chains imposes a harsh limit on the efficiency of a pure NS based approach. Combining these techniques in this way could retain the desirable features of both and serve to mitigate the ever increasing computational demands of energy frontier particle physics.

We close by noting that also in the area of lattice field theory Normalising Flows have recently attracted attention, see e.g. [72, 73], to address the sampling of multimodal target function. We envisage that also in these applications Nested Sampling could be applied.

## 5 Conclusions

The establishing study presented here had two main aims. Firstly to introduce the technique of Nested Sampling, applied to a realistic problem, to researchers in the particle physics community. Secondly to provide a translation back to researchers working on Bayesian inference techniques, presenting an important and active set of problems in particle physics that Nested Sampling could provide a valuable contribution to. The physical example presented used POLY-CHORD to perform an end-to-end generation of events without any input prior knowledge. This is a stylised version of the event generator problem, intended to validate Nested Sampling in this new context and demonstrate some key features. For the considered multi-gluon production processes Nested Sampling was able to learn a mapping in an efficient manner that exhibits promising scaling properties with phase space dimension. We have outlined some potential future research directions; highlighting where the strengths of this approach could be most effective, and how to embed Nested Sampling in a more complete event generator workflow. Along these lines, we envisage an implementation of the Nested Sampling technique for the SHERPA event generator framework [54], possibly also supporting operation on GPUs [74]. This will provide additional means to address the computing challenges for event generation posed by the upcoming LHC runs.

**Acknowledgements** This work has received funding from the European Union's Horizon 2020 research and innovation programme as part of the Marie Skłodowska-Curie Innovative Training Network MCnetITN3 (Grant agreement no. 722104). SS and TJ acknowledge support from BMBF (contract 05H21MGCAB). SS acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - project number 456104544. WH and DY are funded by the Royal Society. This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service ([www.csd3.cam.ac.uk](http://www.csd3.cam.ac.uk)), provided by Dell EMC and Intel using Tier-2

funding from the Engineering and Physical Sciences Research Council (capital Grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council ([www.dirac.ac.uk](http://www.dirac.ac.uk)).

**Data Availability Statement** This manuscript has no associated data or the data will not be deposited. [Authors' comment: The results involve purely theoretical comparisons, publicly available software packages to recreate the calculations are listed and referenced where appropriate.]

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP<sup>3</sup>. SCOAP<sup>3</sup> supports the goals of the International Year of Basic Sciences for Sustainable Development.

## Appendix: Uncertainties in nested sampling

Typical Nested Sampling literature focuses on two main sources of uncertainty; an uncertainty on the weights of the dead points due to the uncertainty in the volume contraction at each iteration, and an uncertainty on the overall volume arising from the path the Markov Chain takes through the space to perform each iteration. The former source is what we consider in this work, and can be calculated as a sample of weights for each dead point using `anesthetic`. The latter source can be estimated using the `nestcheck` package [75], the method presented here uses combinations of multiple runs to form integral estimates meaning the best strategy to minimise this effect is already baked in. Further use cases would benefit from more thorough cross checks using `nestcheck`.

The usual source of uncertainty in a binned histogram in particle physics comes from the standard error. Importance Sampling draws sample events with associated weights  $w_i$ , with the sum of these sample weights giving the estimated cross section in a bin. The effective number of fills in a bin using weighted samples is,

$$N = \frac{(\sum_i w_i)^2}{\sum_i w_i^2}. \quad (7)$$

The inverse square root of  $N$  then constitutes the standard error on the cross section in the bin. In practice this means that the standard deviation of an integral estimated with Importance Sampling can be quoted as  $\Delta_{MC} = \sqrt{\sum_i (w_i^2)}$ . In typical NS applications this is significantly smaller than the pre-



viously mentioned sources, and thus often not considered. However, when using NS as a phase space event generator for finely binned differential observables, the statistical uncertainty can become a significant effect so must be taken into account. Adding the standard error to the weight uncertainty in quadrature is a suitable upper bound for the NS uncertainty but is found to overestimate the uncertainty in some bins. While the standard error gives a measure of the spread of weights around the mean weight in a bin, alternative weights from the sampling history in NS also give an overlapping measure of this.

To correctly account for the statistical error in this context a revised recipe is needed. The following proposed procedure reweights the alternative weight samples to account for the spread of the resulting effective fills in each bin. The effective number of entries in a bin arising from a NS run can be written as,

$$N_j = \frac{(\sum_i w_{j,i})^2}{\sum_i w_{j,i}^2}, \quad (8)$$

where  $i$  indexes the number of weighted samples in each bin, and  $j$  indexes the alternative weights. The result of the  $j$  sampled weight variations is a set of  $j$  different effective counts in each bin. These counts can be modelled as  $j$  trials of a multinomial distribution with  $j$  categories, written as,

$$P(N | \alpha) = \frac{j!}{\prod_j N_j!} \prod_j \alpha_j^{N_j}, \quad (9)$$

where a probability of sampling each category,  $\alpha_j$ , has been introduced. The desired unknown distribution of  $\alpha_j$  can be found using Bayes theorem to invert the arguments. If an uninformative conjugate prior to the multinomial distribution is used, the Dirichlet distribution, the desired inverted probability can also be written in the form of a Dirichlet distribution,

$$P(\alpha | N) = \Gamma \left( \sum_j N_j \right) \prod_j \frac{\alpha_j^{N_j-1}}{\Gamma(N_j)}. \quad (10)$$

A sample vector of  $\alpha_j$  from this Dirichlet distribution, will give a probability of observing each category  $N_j$ . This probability can be used to weight the categories giving a weighted set of effective number of fills,  $\{\alpha_j N_j\}$ . This considers each alternative weight sample as a discrete sample from an underlying continuous distribution  $N_j$  is sampled from. The set of weighted effective fills can be used to quote a weighted set of samples of the bin cross section by multiplying by the square of the sum of the weights,  $\{\sigma_j\} = \{\alpha_j N_j \sum_i (w_{j,i}^2)\}$ . The estimated cross section in the bin is then the expected value of this

set,  $\sigma = E[\sigma_j]$ , and the total standard deviation on this cross section is derived from the variance,  $\Delta\sigma_{\text{tot}} = (\text{Var}[\sigma_j])^2$ .

## References

1. A. Buckley et al., General-purpose event generators for LHC physics. *Phys. Rep.* **504**, 145 (2011). <https://doi.org/10.1016/j.physrep.2011.03.005>. arXiv:1101.2599
2. J.M. Campbell et al., Event Generators for High-Energy Physics Experiments, in 2022 Snowmass Summer Study, 3, (2022). arxiv:2203.11110
3. R. Kleiss, R. Pittau, Weight optimization in multichannel Monte Carlo. *Comput. Phys. Commun.* **83**, 141 (1994). [https://doi.org/10.1016/0010-4655\(94\)90043-4](https://doi.org/10.1016/0010-4655(94)90043-4). arXiv:hep-ph/9405257
4. C.G. Papadopoulos, PHEGAS: a phase space generator for automatic cross-section computation. *Comput. Phys. Commun.* **137**, 247 (2001). [https://doi.org/10.1016/S0010-4655\(01\)00163-1](https://doi.org/10.1016/S0010-4655(01)00163-1). arXiv:hep-ph/0007335
5. F. Krauss, R. Kuhn, G. Soff, AMEGIC++ 1.0: a matrix element generator in C++. *JHEP* **02**, 044 (2002). <https://doi.org/10.1088/1126-6708/2002/02/044>. arXiv:hep-ph/0109036
6. F. Maltoni, T. Stelzer, MadEvent: automatic event generation with MadGraph. *JHEP* **02**, 027 (2003). <https://doi.org/10.1088/1126-6708/2003/02/027>. arXiv:hep-ph/0208156
7. T. Gleisberg, S. Hoeche, Comix, a new matrix element generator. *JHEP* **12**, 039 (2008). <https://doi.org/10.1088/1126-6708/2008/12/039>. arXiv:0808.3674
8. G.P. Lepage, A new algorithm for adaptive multidimensional integration. *J. Comput. Phys.* **27**, 192 (1978). [https://doi.org/10.1016/0021-9991\(78\)90004-9](https://doi.org/10.1016/0021-9991(78)90004-9)
9. T. Ohl, Vegas revisited: adaptive Monte Carlo integration beyond factorization. *Comput. Phys. Commun.* **120**, 13 (1999). [https://doi.org/10.1016/S0010-4655\(99\)00209-X](https://doi.org/10.1016/S0010-4655(99)00209-X). arXiv:hep-ph/9806432
10. S. Jadach, Foam: multidimensional general purpose Monte Carlo generator with selfadapting symplectic grid. *Comput. Phys. Commun.* **130**, 244 (2000). [https://doi.org/10.1016/S0010-4655\(00\)00047-3](https://doi.org/10.1016/S0010-4655(00)00047-3). arXiv:physics/9910004
11. T. Hahn, CUBA: a Library for multidimensional numerical integration. *Comput. Phys. Commun.* **168**, 78 (2005). <https://doi.org/10.1016/j.cpc.2005.01.010>. arXiv:hep-ph/0404043
12. A. van Hameren, PARNI for importance sampling and density estimation. *Acta Phys. Pol. B Ser* **40**, 259 (2009). arXiv:0710.2448
13. H. Kharraziha, S. Moretti, The Metropolis algorithm for on-shell four momentum phase space. *Comput. Phys. Commun.* **127**, 242 (2000). [https://doi.org/10.1016/S0010-4655\(99\)00504-4](https://doi.org/10.1016/S0010-4655(99)00504-4). arXiv:hep-ph/9909313
14. S. Weinzierl, A general algorithm to generate unweighted events for next-to-leading order calculations in electron positron annihilation. *JHEP* **08**, 028 (2001). <https://doi.org/10.1088/1126-6708/2001/08/028>. arXiv:hep-ph/0106146
15. K. Kröniger, S. Schumann, B. Willenberg, (MC)\*\*3: a multi-channel Markov chain Monte Carlo algorithm for phase-space sampling. *Comput. Phys. Commun.* **186**, 1 (2015). <https://doi.org/10.1016/j.cpc.2014.08.024>. arXiv:1404.4328
16. H.S.F. Physics Event Generator, WG collaboration, Challenges in Monte Carlo event generator software for high-luminosity LHC. *Comput. Softw. Big Sci.* **5**, 12 (2021). <https://doi.org/10.1007/s41781-021-00055-1>. arxiv:2004.13687
17. HSF Physics Event Generator WG collaboration, E. Yazgan et al., HL-LHC Computing Review Stage-2, Common Software Projects: Event Generators, 9 (2021)
18. J. Bendavid, Efficient Monte Carlo Integration Using Boosted Decision Trees and Generative Deep Neural Networks, 6 (2017)

19. M.D. Klimek, M. Perelstein, Neural network-based approach to phase space integration. *SciPost Phys.* **9**, 053 (2020). <https://doi.org/10.21468/SciPostPhys.9.4.053>. arXiv:1810.11509
20. S. Otten, S. Caron, W. de Swart, M. van Beekveld, L. Hendriks, C. van Leeuwen et al., Event generation and statistical sampling for physics with deep generative models and a density information buffer. *Nat. Commun.* **12**, 2985 (2021). <https://doi.org/10.1038/s41467-021-22616-z>. arXiv:1901.00875
21. R. Di Sipio, M.F. Giannelli, S.K. Haghighat, S. Palazzo, DijetGAN: a generative-adversarial network approach for the simulation of QCD Dijet Events at the LHC. *JHEP* **08**, 110 (2019). [https://doi.org/10.1007/JHEP08\(2019\)110](https://doi.org/10.1007/JHEP08(2019)110). arxiv:1903.02433
22. A. Butter, T. Plehn, R. Winterhalder, How to GAN LHC Events. *SciPost Phys.* **7**, 075 (2019). <https://doi.org/10.21468/SciPostPhys.7.6.075>. arXiv:1907.03764
23. Y. Alanazi et al., Simulation of electron–proton scattering events by a Feature-Augmented and Transformed Generative Adversarial Network (FAT-GAN), 1 (2020)
24. Y. Alanazi et al., AI-based Monte Carlo event generator for electron–proton scattering, 8. (2020)
25. S. Diefenbacher, E. Eren, G. Kasieczka, A. Korol, B. Nachman, D. Shih, DCTRGAN: Improving the Precision of Generative Models with Reweighting. *JINST* **15** (2020). <https://doi.org/10.1088/1748-0221/15/11/P11004>. arXiv:2009.03796
26. A. Butter, S. Diefenbacher, G. Kasieczka, B. Nachman, T. Plehn, GANplifying event samples. *SciPost Phys.* **10**, 139 (2021). <https://doi.org/10.21468/SciPostPhys.10.6.139>. arXiv:2008.06545
27. A. Butter, S. Diefenbacher, G. Kasieczka, B. Nachman, T. Plehn, GANplifying Event Samples. *SciPost Phys.* **10**, 139 (2021). <https://doi.org/10.21468/SciPostPhys.10.6.139>. arXiv:2008.06545
28. K.T. Matchev, A. Roman, P. Shyamsundar, Uncertainties associated with GAN-generated datasets in high energy physics. *SciPost Phys.* **12**, 104 (2022). <https://doi.org/10.21468/SciPostPhys.12.3.104>. arXiv:2002.06307
29. E. Bothmann, T. Janßen, M. Knobbe, T. Schmale, S. Schumann, Exploring phase space with neural importance sampling. *SciPost Phys.* **8**, 069 (2020). <https://doi.org/10.21468/SciPostPhys.8.4.069>. arXiv:2001.05478
30. C. Gao, J. Isaacson, C. Krause, i-flow: high-dimensional integration and sampling with normalizing flows. *Mach. Learn. Sci. Tech.* **1**, 045023 (2020). <https://doi.org/10.1088/2632-2153/abab62>. arXiv:2001.05486
31. C. Gao, S. Höche, J. Isaacson, C. Krause, H. Schulz, Event generation with normalizing flows. *Phys. Rev. D* **101**, 076002 (2020). <https://doi.org/10.1103/PhysRevD.101.076002>. arXiv:2001.10028
32. B. Stienen, R. Verheyen, Phase space sampling and inference from weighted events with autoregressive flows. *SciPost Phys.* **10**, 038 (2021). <https://doi.org/10.21468/SciPostPhys.10.2.038>. arXiv:2011.13445
33. K. Danziger, T. Janßen, S. Schumann, F. Siegert, Accelerating Monte Carlo event generation—rejection sampling using neural network event-weight estimates, 9 (2021)
34. M. Backes, A. Butter, T. Plehn, R. Winterhalder, How to GAN event unweighting. *SciPost Phys.* **10**, 089 (2021). <https://doi.org/10.21468/SciPostPhys.10.4.089>. arXiv:2012.07873
35. M. Bellagente, M. Haußmann, M. Luchmann, T. Plehn, Understanding event-generation networks via uncertainties, 4 (2021)
36. A. Butter, T. Heimel, S. Hummerich, T. Krebs, T. Plehn, A. Rousset et al., Generative networks for precision enthusiasts, 10 (2021)
37. J. Skilling, Nested sampling for general Bayesian computation. *Bayesian Anal.* **1**, 833 (2006). <https://doi.org/10.1214/06-BA127>
38. W.J. Handley, M.P. Hobson, A.N. Lasenby, polychord: next-generation nested sampling. *Mon. Not. R. Astron. Soc.* **453**, 4385 (2015). <https://doi.org/10.1093/mnras/stv1911>. arXiv:1506.00171
39. D.J.C. MacKay, *Information Theory. Inference & Learning Algorithms* (Cambridge University Press, Cambridge, 2002)
40. Particle Data Group collaboration, Review of Particle Physics. *PTEP* **2020**, 083C01 (2020). <https://doi.org/10.1093/ptep/ptaa104>
41. G. Ashton et al., Nested sampling for physical scientists. *Nature* **2** (2022). <https://doi.org/10.1038/s43586-022-00121-x>. arXiv:2205.15570
42. R.M. Neal, Slice sampling. *Ann. Stat.* **31**, 705 (2003). <https://doi.org/10.1214/aos/1056562461>
43. P. Mukherjee, D. Parkinson, A.R. Liddle, A nested sampling algorithm for cosmological model selection. *Astrophys. J. Lett.* **638**, L51 (2006). <https://doi.org/10.1086/501068>. arXiv:astro-ph/0508461
44. R. Shaw, M. Bridges, M.P. Hobson, Clustered nested sampling: efficient Bayesian inference for cosmology. *Mon. Not. R. Astron. Soc.* **378**, 1365 (2007). <https://doi.org/10.1111/j.1365-2966.2007.11871.x>. arXiv:astro-ph/0701867
45. F. Feroz, M.P. Hobson, Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis. *Mon. Not. R. Astron. Soc.* **384**, 449 (2008). <https://doi.org/10.1111/j.1365-2966.2007.12353.x>. arXiv:0704.3704
46. F. Feroz, M.P. Hobson, M. Bridges, MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009). <https://doi.org/10.1111/j.1365-2966.2009.14548.x>. arXiv:0809.3437
47. F. Feroz, M.P. Hobson, E. Cameron, A.N. Pettitt, Importance nested sampling and the MultiNest algorithm. *Open J. Astrophys.* **2**, 10 (2019). <https://doi.org/10.21105/astro.1306.2144>. arXiv:1306.2144
48. W.J. Handley, M.P. Hobson, A.N. Lasenby, PolyChord: nested sampling for cosmology. *Mon. Not. R. Astron. Soc.* **450**, L61 (2015). <https://doi.org/10.1093/mnras/rlv047>. arXiv:1502.01856
49. W. Handley, P. Lemos, Quantifying dimensionality: Bayesian cosmological model complexities. *Phys. Rev. D* **100**, 023512 (2019). <https://doi.org/10.1103/PhysRevD.100.023512>. arXiv:1903.06682
50. S.S. AbdusSalam et al., Simple and statistically sound recommendations for analysing physical theories, 12 (2020)
51. GAMBIT collaboration, Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module. *Eur. Phys. J. C* **77**, 761 (2017). <https://doi.org/10.1140/epjc/s10052-017-5274-y>. arXiv:1705.07959
52. A. Fowlie, W. Handley, L. Su, Nested sampling with plateaus. *Mon. Not. R. Astron. Soc.* **503**, 1199 (2021). <https://doi.org/10.1093/mnras/stab590>. arXiv:2010.13884
53. E. Higson, W. Handley, M. Hobson, A. Lasenby, Sampling errors in nested sampling parameter estimation. *Bayesian Analysis series* **13** (2018). <https://doi.org/10.1214/17-ba1075>
54. Sherpa collaboration, Event Generation with Sherpa 2.2. *SciPost Phys.* **7**, 034 (2019). <https://doi.org/10.21468/SciPostPhys.7.3.034>. arXiv:1905.09127
55. A. van Hameren, C.G. Papadopoulos, A hierarchical phase space generator for QCD antenna structures. *Eur. Phys. J. C* **25**, 563 (2002). <https://doi.org/10.1007/s10052-002-1000-4>. arXiv:hep-ph/0204055
56. R. Kleiss, W.J. Stirling, S.D. Ellis, A new Monte Carlo treatment of multiparticle phase space at high-energies. *Comput. Phys. Commun.* **40**, 359 (1986). [https://doi.org/10.1016/0010-4655\(86\)90119-0](https://doi.org/10.1016/0010-4655(86)90119-0)
57. S. Plätzer, RAMBO on diet, 8 (2013). arXiv:1308.2922
58. A. Fowlie, S. Hoof, W. Handley, Nested sampling for frequentist computation: fast estimation of small p values. *Phys. Rev. Lett.* **128**, 021801 (2022). <https://doi.org/10.1103/PhysRevLett.128.021801>. arXiv:2105.13923
59. E. Carragher, W. Handley, D. Murnane, P. Stangl, W. Su, M. White et al., Convergent Bayesian global fits of 4D compos-

- ite Higgs models. *JHEP* **05**, 237 (2021). [https://doi.org/10.1007/JHEP05\(2021\)237](https://doi.org/10.1007/JHEP05(2021)237). arXiv: 2101.00428
60. A. Fowlie, W. Handley, L. Su, Nested sampling cross-checks using order statistics. *Mon. Not. R. Astron. Soc.* **497**, 5256 (2020). <https://doi.org/10.1093/mnras/staa2345>. arXiv:2006.03371
  61. C. Bierlich et al., Robust independent validation of experiment and theory: rivet version 3. *SciPost Phys.* **8**, 026 (2020). <https://doi.org/10.21468/SciPostPhys.8.2.026>. arXiv:1912.05451
  62. M. Cacciari, G.P. Salam, G. Soyez, The anti- $k_t$  jet clustering algorithm. *JHEP* **04**, 063 (2008). <https://doi.org/10.1088/1126-6708/2008/04/063>. arXiv:0802.1189
  63. W. Handley, anesthetic: nested sampling visualisation. *J. Open Sour. Softw.* **4**, 1414 (2019). <https://doi.org/10.21105/joss.01414>. arXiv:1905.04768
  64. ATLAS collaboration, Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the ATLAS experiment, 12 (2021). arXiv:2112.09588
  65. S. Höche, S. Prestel, H. Schulz, Simulation of vector boson plus many jet final states at the high luminosity LHC. *Phys. Rev. D* **100**, 014024 (2019). <https://doi.org/10.1103/PhysRevD.100.014024>. arXiv:1905.05120
  66. T. Gleisberg, F. Krauss, Automating dipole subtraction for QCD NLO calculations. *Eur. Phys. J. C* **53**, 501 (2008). <https://doi.org/10.1140/epjc/s10052-007-0495-0>. arXiv: 0709.2881
  67. A. Petrosyan, W. Handley, SuperNest: accelerated nested sampling applied to astrophysics and cosmology, Maximum Entropy (accepted for Oral presentation) (2022)
  68. “SuperNest.” <https://gitlab.com/a-p-petrosyan/sspr>. <https://pypi.org/project/supernest/>
  69. E. Higson, W. Handley, M. Hobson, A. Lasenby, Dynamic nested sampling: an improved algorithm for parameter estimation and evidence calculation. *Stat. Comput.* **29**, 891–913 (2018). <https://doi.org/10.1007/s11222-018-9844-0>
  70. S. Badger et al., Machine Learning and LHC Event Generation, in 2022 Snowmass Summer Study, A. Butter, T. Plehn and S. Schumann, eds., 3 (2022). arXiv:2203.07460
  71. J. Alsing, W. Handley, Nested sampling with any prior you like. *Mon. Not. R. Astron. Soc.* **505**, L95 (2021). <https://doi.org/10.1093/mnras/slab057>. arXiv:2102.12478
  72. L. Del Debbio, J.M. Rossney, M. Wilson, Efficient modeling of trivializing maps for lattice  $\phi^4$  theory using normalizing flows: a first look at scalability. *Phys. Rev. D* **104**, 094507 (2021). <https://doi.org/10.1103/PhysRevD.104.094507>. arXiv:2105.12481
  73. D.C. Hackett, C.-C. Hsieh, M.S. Albergo, D. Boyda, J.-W. Chen, K.-F. Chen et al., Flow-based sampling for multimodal distributions in lattice field theory, 7 (2021)
  74. E. Bothmann, W. Giele, S. Höche, J. Isaacson, M. Knobbe, Many-gluon tree amplitudes on modern GPUs: a case study for novel event generators, 6 (2021)
  75. E. Higson, W. Handley, M. Hobson, A. Lasenby, Nestcheck: diagnostic tests for nested sampling calculations. *Mon. Not. R. Astron. Soc.* **483**, 2044 (2019). <https://doi.org/10.1093/mnras/sty3090>. arXiv:1804.06406