



A Positive Resampler for Monte Carlo events with negative weights

Jeppe R. Andersen^{1,a}, Christian Gütschow^{2,b}, Andreas Maier^{3,c}, Stefan Prestel^{4,d}

¹ Institute for Particle Physics Phenomenology, University of Durham, Durham DH1 3LE, UK

² Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

³ Deutsches Elektronen-Synchrotron, DESY, Platanenallee 6, 15738 Zeuthen, Germany

⁴ Theoretical Particle Physics, Department of Astronomy and Theoretical Physics, Lund University, Sölvegatan 14 A, 223 62 Lund, Sweden

Received: 20 July 2020 / Accepted: 13 October 2020 / Published online: 31 October 2020
© The Author(s) 2020

Abstract We propose the *Positive Resampler* to solve the problem associated with event samples from state-of-the-art predictions for scattering processes at hadron colliders typically involving a sizeable number of events contributing with negative weight. The proposed method guarantees positive weights for all physical distributions, and a correct description of all observables. A desirable side product of the method is the possibility to reduce the size of event samples produced by General Purpose Event Generators, thus lowering the resource demands for subsequent computing-intensive event processing steps. We demonstrate the viability and efficiency of our approach by considering its application to a next-to-leading order + parton shower merged prediction for the production of a W boson in association with multiple jets.

1 Introduction

General Purpose Event Generators [1–3] form a crucial component of studies in high-energy physics, since they produce detailed predictions used for the design and calibration of detectors, interpretations of the measurements as well as the investigations of theoretical models. More often than not it is necessary to take into account the effects from the perturbative showering and the hadronisation models implemented in these generators, in order to achieve an accurate prediction for the cuts and observables chosen for experimental measurements.

High-accuracy perturbative event generator predictions can be obtained by first matching each jet multiplicity to next-

to-leading order (NLO) using the methods of e.g. MC@NLO [4] or POWHEG [5], followed by a merging of these exclusive samples using approaches such as MEPS@NLO [6] or UNLOPS [7,8]. The increased accuracy comes at a significant cost in additional computing resources, and these calculations increasingly contribute to the LHC computing footprint. The result of these merged NLO-accurate event generator simulations are event samples containing events of both positive and negative weights, meaning that the correspondence between the number of events in a bin of a distribution and the cross section in that bin is lost.¹ Even when the event samples are unweighted to constitute events with weights of ± 1 , the number of negative weighted events can be significant. This reduces the statistical significance of the sample compared to one with the same number of all positive weight events.

The outcome of event generation is often processed through time-consuming detector simulation – which currently constitutes the major part of the LHC computing budget. Since both positive and negative weight events are afterwards processed at a significant cost, it is beneficial to reduce the cancellation of events of negative and positive weight. This can be done by reducing the occurrence of negative weight events in the event generation, see e.g. [9] for an approach in MC@NLO.

We report here on an alternative and NLO-matching-independent approach to completely remove all negative weights from any already generated sample, and re-introducing the correspondence between the number of events in a bin and the local contribution to the cross section. This *positive resampling* will be achieved in a two-stage process: (1)

^a e-mail: jeppe.andersen@durham.ac.uk

^b e-mail: chris.g@cern.ch

^c e-mail: andreas.martin.maier@desy.de

^d e-mail: stefan.prestel@thep.lu.se (corresponding author)

¹ The fraction of negative weights can vary wildly between different matching or merging schemes, but in general will be non-negligible for processes containing multiple light jets. This is illustrated on Figs. 1 and 2.

modify the weights of events to be all positive and possibly smaller in magnitude, and (2) apply a standard unweighting. The second step should be taken only if the number of events is sought to be reduced. Reducing the event sample can significantly lower the computing budget for the steps in the analysis following the event generation, both in terms of CPU and disk.²

It may seem counter intuitive to allow for or even seek a reduction in the number of events, since traditionally the statistical significance of a sample, or the variance, is linked to the number of events, and reducing the number of events would therefore reduce the statistical significance of the sample. However, when the sample contains events with both positive and negative weights, the number of events can indeed be reduced without impacting the statistical significance: The effective cross section is given by $\sigma = \sigma_p - \sigma_n$, where σ_p is the contribution from events with positive weights, and σ_n the absolute value of the negative-weight contribution. The Monte Carlo variance associated with the sample is $s_p^2/N_p + s_n^2/N_n$, where s_p^2, s_n^2 is the variance of the integrand. If we replace the sample with one that has N_s positively weighted events, the variance of this new sample is s_s^2/N_s . If therefore $N_s/(N_p + N_n)$ is similar to s_s/s_p or s_s/s_n , the variance can be unchanged. This is achieved with the *Positive Resampler*, a simplified description of which transforms the weight of each event to its absolute value and multiplies by $(\sigma_p - \sigma_n)/(\sigma_p + \sigma_n)$. If there is a large cancellation, then the weight of each event is much reduced, leading also to a reduction in the Monte Carlo estimate of the variance. Hence fewer events are needed for the same statistical certainty. The *Positive Resampler* is introduced in Sects. 3 and 4 showcases results obtained based on samples of $W+0, 1, 2$ jets at NLO fixed-order accuracy merged with UNLOPS.

2 Weights in UNLOPS merging

High-precision predictions are important ingredients to LHC data analysis. If the analysis is sensitive to the effect of yet higher jet multiplicities, high precision is obtained by “merging” several distinct calculations. NLO merged calculations provide the state-of-the-art for LHC phenomenology, and contribute significantly to the overall computing resource usage. At the same time, increased precision almost always comes at the cost of having to rely more heavily on weighted event generation. Typically, issues due to a reduction in sta-

tistical convergence worsen with every additional NLO calculation included in the merging.

NLO merging schemes aim to produce inclusive event samples that both comply with NLO fixed-order accuracy for several multi-jet processes, and ensure that the accuracy of parton showering is preserved. Various methods have been developed to this effect [6, 7, 12–16], all differing slightly in the concrete goals and the definition of the target accuracy. Each NLO merging method suffers from various sources of negative weights. Predictions based on unitarised merging schemes [7] are the only predictions that not only comply with both of the above criteria, but also guarantee that, for an arbitrary base process and *arbitrary* parton multiplicity, inclusive n -jet cross sections are preserved exactly, without introducing sub-dominant contributions due to the merging.³ This desirable feature comes at the price of introducing new sources of counter-events and/or event weights compared to other methods such as [6].

Hence, unitarised NLO merging provides a very non-trivial test of unweighting methods. In the following, we will employ the NLO merging as implemented in the DIRE plugin to the PYTHIA event generator as a test case. This NLO merging implementation is based on UNLOPS [7], includes QCD, QED and electroweak vector boson emissions, and thus allows merging of calculations with multiple hard jets, photons/leptons or electroweak vector bosons. In particular, the “EW-improved” merging of [19] is extended to NLO QCD. Unordered configurations (i.e. for which no history of ordered emissions can be reconstructed for an input event) are treated according to the MOPS+unordered prescription of [20]. The latter means that the merging also heavily relies on matrix elements extracted from MadGraph5_aMC@NLO [21]. Within this NLO merging framework, several sources of event weights arise:

- The input NLO QCD short-distance cross sections (generated with MadGraph5_aMC@NLO for this study) can contain positive and negative weights. The input events are typically unweighted to $\pm w$, where w is a unit weight.
- The procedure of assigning a parton shower history to a high-multiplicity input event can necessitate weights. The history is chosen among all possible shower histories according to the shower probability, which may contain non-positive definite splitting functions in DIRE. This leads to a corrective (positive or negative) weight.
- The merging procedure enforces a consistent renormalisation and factorisation scale setting by introducing weights. These weights are almost entirely positive and tend to fluctuate only mildly. The merging scheme fur-

² Note that within the specific field of sequential Monte Carlo methods [10], the term “resampling” has a more constrained meaning than in statistics more widely. “Resampling” in this more constrained meaning has been employed for parton-shower simulations in [11]. The latter has very little overlap with our method, or the goals of the current article.

³ It should however be noted that calculations that are valid for specific processes [14, 17], or up to a maximal multiplicity [13, 18] also fulfill similar consistency criteria.

ther removes overlap between different input samples by including no-emission probabilities. These factors are essentially in $[0, 1]$, but can, in rare cases (e.g. due to negative NLO parton distribution functions or splitting functions), lead to negative weights.

- The $\mathcal{O}(\alpha_s)$ expansion of the weight discussed in the previous point – which is necessary to guarantee the NLO accuracy of the method – is often negative, thus introducing a non-negligible source of weights.
- The accuracy of inclusive cross sections is enforced by explicit unitarisation [7, 8]. This means that a large fraction of events will be employed as “counter-events” with negative weight.
- Independently of NLO merging, the subsequent showering may also generate event positive or negative weights if the splitting functions are not positive-definite.

The regions of phase space where negatively weighted events contribute most depends on the source of weights, and thus also on the merging setup. An overall larger fraction of negative weights can be expected when merging more NLO calculations. This is best illustrated with an example. Consider a precise background prediction for a vector-boson + jets measurement. If only the inclusive zero-jet prediction is NLO accurate, then the predominant source of negative weights is the unitarisation of reweighted one-jet LO configurations. This will lead to a moderate amount of negatively weighted events, which produce a relatively flat negative contribution to the vector boson rapidity spectrum. The same events will induce larger negative weight fraction at small boson $p_{\perp V}$ than at high $p_{\perp V}$, since showering from counter-events is constrained to the soft/collinear regions. Negative weights will have a negligible impact on observables that require one, two or more jets. If the calculation is extended by also including an NLO calculation for the inclusive V +one-jet rate, then $p_{\perp V}$ will also exhibit negative-weight contributions at high values, since new sources of negative weights ($\mathcal{O}(\alpha_s)$ expansions, unitarisation of the two-jet LO sample) arise. Observables that depend on two jets will now acquire a negative component at small jet separation due to the mechanism of unitarisation. This effect is illustrated in Fig. 1.

3 A positive resampler

Our main goal is to restore the connection seen when generating all positive weight events between the number of events in the neighbourhood of any phase space point (e.g. the bin in any histogram) and the local contribution to the cross section. This will require a modification of the weight of all events in the neighbourhood of events with negative weights. The idea is simply stated to (1) calculate the cross section σ_i from the

N_i positive and negative weight events in a neighbourhood, (2) change each weight to be positive and rescale all weights to preserve σ_i . One can then proceed with a unweighting procedure over all the neighbourhoods i to restore the connection between the number of events and the cross section from each neighbourhood.

In Monte Carlo event generation, the cross sections are generated exclusively in all momenta. We will demonstrate first the method for the idealised situation where the cross section is stored differentially in all momenta relevant for the later event analysis – this could be e.g. momenta of the jets, leptons etc. We will then demonstrate how it works also in the case of using simple binned observables for the unweighting.

3.1 Multi-dimensional resampling

We begin by considering the idealistic case (i.e. the limit of infinitely narrow bin widths) where a differential distribution in the observable $\mathcal{O}_1$ is calculated with both positive and negative weight events in a f -body phase space region Ω_f . The differential distribution in the observable $\mathcal{O}_1$ is then constructed as

$$\frac{d\sigma}{d\mathcal{O}_1} = \frac{d\sigma}{d\Omega_f} \frac{d\Omega_f}{d\Omega_n} \frac{d\Omega_n}{d\mathcal{O}_1}, \quad (1)$$

where

$\frac{d\sigma}{d\Omega_f}$ signifies the cross section calculated in terms of the final state momenta.
 $\frac{d\Omega_f}{d\Omega_n}$ encodes e.g. the jet clustering and is the Jacobian for the f -body phase space into the n -body phase space ($n < f$) that the observable depends on.
 $\frac{d\Omega_n}{d\mathcal{O}_1}$ is traditionally included in the calculations by the binning of the n -body phase space in terms of $\mathcal{O}_1$.

Equation (1) just constitutes the chain rule of differentiation. Histogrammed distributions are obtained by simple integrations of this relation. It is therefore not surprising that if the events of positive and negative weights arising in the generation (represented by $\frac{d\sigma}{d\Omega_f}$) can be turned into events of all positive weights, then the distribution in any $\mathcal{O}_1$ can be calculated with these events (since the calculation of the other factors in Eq. (1) are unchanged). Of course the f -body phase space can have a very high dimension and it may seem impractical to perform the clustering of events in bins in all the dimensions of Ω_f followed by the reweighting procedure outlined above. One could of course perform the reweighting in the lower-dimensional Ω_n , the phase space of all objects (jets etc.) entering observables. It is clear here that the reweighting works for any observable $\mathcal{O}_1$. What is perhaps surprising is that if the reweighting is performed in neighbourhoods (bins after integration) in $\mathcal{O}_1$ then $\frac{d\sigma}{d\mathcal{O}_2}$ for

any other observable $\mathcal{O}_2$ can still be constructed: We have

$$\frac{d\sigma}{d\mathcal{O}_2} = \frac{d\Omega_f}{d\mathcal{O}_2} \left\{ \left(\frac{d\sigma}{d\Omega_f} \frac{d\Omega_f}{d\Omega_n} \frac{d\Omega_n}{d\mathcal{O}_1} \right) \frac{d\mathcal{O}_1}{d\Omega_f} \right\}. \quad (2)$$

All reference to the cross section σ is within the brackets ($\cdots$), so the cancellation of positive and negative weight events can be implemented here in terms of reweighting as above to the distribution in $\mathcal{O}_1$. The effect of $d\mathcal{O}_1/d\Omega_f$ is taken into account by calculating the value (or bin) of $\mathcal{O}_1$ starting from the phase space points in Ω_f . The spectrum for the observable $\mathcal{O}_2$ is then calculated by constructing the quantities

$\frac{d\Omega_f}{d\mathcal{O}_2}$ by binning the contribution from the phase space points in $\mathcal{O}_f$ in terms of the observable $\mathcal{O}_2$
 $\{\cdots\}$ by finding all Ω_f resulting in $\mathcal{O}_1$, and multiply by the differential distribution in $\mathcal{O}_1$.

So, in terms of truly differential distributions, it would not matter which observable $\mathcal{O}_i$ one would start from. The distribution in $\mathcal{O}_j$ can be obtained by the procedure above. This is correct up to effects in the binning in $\mathcal{O}_1$, which we will later show are modest indeed, and can be reduced further by decreasing the bin widths used in the Positive Resampler. Indeed, in the situation where the impact of the negative weight events in phase space Ω_f mapped into the equivalent bins in $\mathcal{O}_1$ and $\mathcal{O}_2$ is identical, then the result for $\mathcal{O}_2$ using the Positive Resampler in $\mathcal{O}_1$ is exact. We will see in Sect. 4.1 how convergence can be achieved by resampling in multiple dimensions.

3.2 Resampling the total cross section

The extreme opposite of resampling in all n dimensions of the momenta of constructs used in the analysis, such as jets and leptons, is resampling just the cross section – effectively using just one bin. This approach may be considered too coarse in practice, but we include the discussion here since it provides a simple example of the algorithms used. For clarity, we choose to illustrate the method in terms of bins and weights, obtained as integrals and MC sampling of the distributions discussed in Sect. 3.1.

Let us start by considering the total cross section

$$\sigma = \sum_{i=1}^N w_i, \quad (3)$$

obtained from N events with weights w_i . Introducing a convenient factor of one, we can write

$$\sigma = \sum_{i=1}^N w_i = \frac{\left(\sum_{i=1}^N |w_i| \right) \left(\sum_{i=1}^N w_i \right)}{\sum_{i=1}^N |w_i|} \equiv P \sum_{i=1}^N |w_i|, \quad (4)$$

where $0 \leq P = \frac{\sum_i w_i}{\sum_i |w_i|} \leq 1$. Effectively, this amounts to replacing each event weight w_i by $P|w_i|$. The total cross section is preserved by construction, but what is the effect on binned distributions?

To answer this question, let us select an arbitrary distribution $\frac{d\sigma}{d\mathcal{O}}$ and an arbitrary bin B ranging from $\mathcal{O}_B$ to $\mathcal{O}_{B+1}$ and containing $M \gg 1$ events. Without loss of generality we can assume that the bin contains the events $i = 1, \dots, M$. The height of the bin is given by

$$\sigma_{\mathcal{O},B} \equiv \frac{d\sigma}{d\mathcal{O}}(\mathcal{O}_{B+1} - \mathcal{O}_B) = \sum_{i=1}^M w_i = \frac{\sum_{i=1}^M |w_i|}{\sum_{i=1}^M |w_i|} \sum_{i=1}^M w_i, \quad (5)$$

where we have introduced the same factor of one as previously. In the simplified case with a uniform distribution of negative weights,⁴

$$\frac{\sum_{i=1}^M w_i}{\sum_{i=1}^M |w_i|} \approx \frac{\sum_{i=1}^M w_i}{\sum_{i=1}^M |w_i|}, \quad (6)$$

we can effectively swap the summation limits in the numerator of the Eq. (5) for the bin height:

$$\begin{aligned} \sigma_{\mathcal{O},B} &= \sum_{i=1}^M w_i = \frac{\sum_{i=1}^M |w_i|}{\sum_{i=1}^M |w_i|} \sum_{i=1}^M w_i \\ &\approx \frac{\sum_{i=1}^M w_i}{\sum_{i=1}^M |w_i|} \sum_{i=1}^M |w_i| = P \sum_{i=1}^M |w_i|. \end{aligned} \quad (7)$$

This tells us that simply replacing each event weight w_i by $P|w_i|$ preserves the total cross section *exactly* and also all bin heights in each distribution in the limit of large statistics, i.e. within statistical fluctuations for finite statistics. In practice, negative weights will not necessarily be distributed uniformly. In such a case, we expect to introduce a systematic error that grows with the gradient of the negative-weight contribution. We will analyse the real-world performance for the example of $W + \text{jets}$ production in Sect. 4.

⁴ This follows directly from the law of large numbers $\frac{1}{N} \sum_{i=1}^N x_i \approx \frac{1}{M} \sum_{i=1}^M x_i$ for random variables $x_i \in \{w_i, |w_i|\}$, but also holds more generally. For example, generating larger event samples in specific regions of phase space introduces a bias in the event weights but does not affect the following conclusions.

3.3 Resampling a distribution

For one-dimensional sampling in terms of an observable $\mathcal{O}$, the method can be adapted to preserve exactly the distribution in $\mathcal{O}$. The accuracy in other distribution will be determined by the variation in the impact of negative weight events for the phase space Ω_f mapped into each bin in the two distributions. In analogy to Eq. (4) we introduce separate rescaling factors $P_{\mathcal{O},B}$ for each bin B containing the events $i = i_B, \dots, i_{B+1}$:

$$\sigma_{\mathcal{O},B} = \sum_{i=i_B}^{i_{B+1}} w_i = \frac{\sum_{i=i_B}^{i_{B+1}} |w_i|}{\sum_{i=i_B}^{i_{B+1}} |w_i|} \sum_{i=i_B}^{i_{B+1}} w_i \equiv P_{\mathcal{O},B} \sum_{i=i_B}^{i_{B+1}} |w_i|. \quad (8)$$

This preserves all bin heights exactly and therefore the full distribution in $\mathcal{O}$. The total cross section is just the sum over all bins, which also remains unchanged. In cases where the systematic variation in the distribution of negative weights is negligible, Eq. (6) guarantees that all rescaling factors $P_{\mathcal{O}',B'}$ for all observables $\mathcal{O}'$ and bins B' converge to the same value. This implies that all other distributions remain correct in the limit of large statistics.

It is straightforward to generalise the argument to multi-differential distributions. For instance, if we resample bins in a double-differential distribution in $\mathcal{O}$ and $\mathcal{O}'$ then also the respective single-differential distributions in $\mathcal{O}$ and $\mathcal{O}'$ as well as the total cross section are preserved exactly. The most extreme case is a differential distribution in all final-state momenta. In the limit of infinitesimal bin widths we recover the idealised scenario already described in Sect. 3.1. The convergence to the correct result can be achieved by appropriately increasing the dimensions used in the Positive Resampler.

Consider now the situation where two observables $\mathcal{O}$ and $\mathcal{O}'$ should be described simultaneously. The correct result for both observables would be obtained if the P -factors were calculated separately for each set of events in the overlaps of the bins of $\mathcal{O}$ and $\mathcal{O}'$. The i 'th bin in the distribution of $\mathcal{O}'$ would then be found as

$$\sigma_{\mathcal{O}',i} = \sum_j P_{\mathcal{O},j \cap \mathcal{O}',i} \sum_{k \in \mathcal{O}, j \cap \mathcal{O}',i} |w_k|. \quad (9)$$

This method can be straightforwardly generalised to more observables, and will be simpler than reverting to binning the full phase space Ω_n .

Consider now the situation where a one-dimensional observable $\mathcal{O}$ and distribution is used for the resampling, and the result is then binned in another observable $\mathcal{O}'$. The dependence of P on the bins i in $\mathcal{O}'$ in each bin j of $\mathcal{O}$ is then not known, and instead of Eq. (9) one can at best obtain

the approximation

$$\sigma_{\mathcal{O}',i} \approx \sum_j P_{\mathcal{O},j} \sum_{k \in \mathcal{O}, j \cap \mathcal{O}',i} |w_k|. \quad (10)$$

The error made in going from Eq. (9) to Eq. (10) is proportional to the variation in P for sub-division of the bin j . Narrower bin widths in j will improve the situation, especially if there is any correlation between the observables $\mathcal{O}$ and $\mathcal{O}'$. We will later (see Fig. 3) investigate the outcome of resampling in the transverse momentum of the leading jet, and then bin in the rapidity of this jet. This is a “worst-case” scenario, where there is little correlation between the observables.

4 Results from the positive resampler

To demonstrate the practical performance of the algorithm outlined in Sect. 3 we consider an α_s -driven NLO-merged description of the W -boson + jets process, as an example of a resource-intensive calculation. We focus on results for proton-proton collisions at 14 TeV centre-of-mass energy. The inclusive event sample merging NLO QCD calculations for 0, 1 and 2 additional jets is generated by using MadGraph5_aMC@NLO⁵ merged and showered by DIRE. Multiparton interactions, hadronisation and beam remnants are handled by PYTHIA. The outputs are analysed using RIVET [22], which provides the input distributions to the Positive Resampler, and performs the analysis and plotting. We present results for the standard Rivet analyses MC_WINC and MC_WJETS in order to illustrate the performance of the Positive Resampler on analyses already in use by the community for testing the description of the production of a W -boson.

We then apply a series of resampling steps to each generated weighted event sample. Since the events in each sample are based on fixed-order input with a uniform number of observable final state particles, the resampling will be automatically differential in this quantity. In addition, we choose a further target (binned) distribution which we aim to preserve exactly during the procedure. We are free to pick an arbitrary distribution for this purpose. In the present work, we consider first the following one-dimensional examples.

⁵ NLO inputs are generated with aMC@NLO, with loose cuts and employing PYTHIA shower subtractions, and post-processed by performing a single PYTHIA evolution step using the settings recommended in the MadGraph5_aMC@NLO documentation. The results after post-processing are stored and used as input for DIRE. LO events are also generated with loose cuts. Tighter merging scale cuts on LO and NLO samples are applied by DIRE. Table 1 lists the number of events after the merging scale cut.

- a) *W boson transverse momentum*. We preserve $\frac{d\sigma}{dp_{W\perp}}$, in bins with a width of 5 GeV each. In the peak region $p_{W\perp} < 125$ GeV this coincides with a histogram from the MC_WINC Rivet analysis.
- b) *Parton shower evolution variable t_1* . This unmeasurable parameter is defined as the DIRE parton shower ordering variable [23] at which a zero-parton state transitions to a one-parton state. For events with more than zero partons before showering, we use the t -value of the first node in the reconstructed parton-shower history that is employed when calculating no-emission probabilities within UNLOPS. If no transition from a zero- to a one-parton state exists (e.g. if the shower did not perform an emission, or if the history favoured electroweak clusterings), we set $t_1 = -1$. For $t_1 > 0$, we consider $\frac{d\sigma}{d\log t_1}$ in $\sqrt{N}$ bins, where N is the number of events in each sample. We add one more bin for events with $t_1 < 0$.

Our aim here is to demonstrate the viability of positive resampling independently of the concrete analysis. For this reason we consider distributions for the complete event samples and refrain from applying any analysis cuts at this stage when using the resampling prescription (b). In contrast to this, option (a) relies on the Rivet analysis for identifying the momentum of the W boson and analysis cuts have to be applied before resampling.

In the next step, we apply the following procedure to each bin in the chosen distribution. These are the same steps later applied in each bin of multi-dimensional sampling.

1. *Change weights*. We turn negative-weight into positive-weight events as described in Sect. 3. This preserves the height of each bin exactly.
2. *Partial unweighting*. We choose a target weight W_t and apply standard unweighting to all events with weights $w_i < W_t$, i.e. we keep the event with probability $p = \frac{w_i}{W_t}$ and then adjust the event weight to W_t . We choose W_t in such a way that approximately 10% of the original events are kept. This number can straightforwardly be adjusted.
3. *Bin restoration*. We rescale the weights of all events in the bin such that the original bin height is restored. If the stochastic unweighting in the previous step resulted in an empty bin, we first pick a random discarded event and add it back to the bin prior to restoring the height. The probability to recover a discarded event is chosen to be proportional to its weight.

We also compare our approach to more traditional unweighting, where the signs of the event weights are unchanged, and the modulus of the event weight is used when deciding whether an event should be kept.

As is illustrated on Fig. 1, the results obtained using both traditional unweighting and the Positive Resampler are consistent with the result using all events, well within the statistical fluctuation of that largest sample. It is also clear that the statistical uncertainty associated with the distributions obtained with the Positive Resampler is similar to that of the input *weighted* distribution. The central values and associated uncertainty is for both calculated by Rivet.

To some degree, the power of the method is in the reduction of events necessary to achieve a certain statistical accuracy.

Table 1 lists the number of events in each sample; the input test sample to the Positive Resampler contains 5.3M events, of which 195k pass the cuts of the analysis. Of these, 3.2M and 121k respectively have positive weights. The simple unweighting procedure leaves 1.5M (52k passing cuts of the analysis). The Positive Resampler according to t_1 reduces this to 659k and 25k respectively. The Positive Resampler($p_{\perp}^W$) by construction accepts events only if they pass the cuts of the analysis, which with this resampling is 33k. The true power of the positive resampling is in the reduction from 52k to 25k or 33k of events. This reduction in the number of events is brought about by the cancellation between positive and negative contributions, illustrated in Fig. 1 by the contribution of “positive only”. The fact that this contribution is roughly twice the cross section at the peak of the distribution (most easily checked for y_w) leads directly to the reduction in the number of events obtained with the Positive Resampler. The increasing ratio of the “positive only” to the “weighted” result at large $p_{\perp}^W$ and $p_{\perp}^l$ is a result of the increasing cancellations taking place within the input sample as discussed in Sect. 2. It is these cancellations that the Positive Resampler implements effectively by reducing the event count. The number of events left after the resampling can be adjusted and tuned – the largest possible number of events per unit of cross section is given by the number of events (positive and negative) in the bin with the least cancellation, divided by the cross section in this bin. The distributions expose large cancellations in some regions of phase space, indicating that many more events are required to obtain statistically meaningful spectra. Improvements of the example NLO-merging implementation used – to reduce the amount of cancellation already at the “weighted” stage – would clearly be beneficial [9].

4.1 Positive resampling in higher dimensions

The distributions studied in the previous section all relate to the momenta of the W and its decay products. Figure 2 presents results for the transverse momentum distribution of the three hardest jets produced in association with the W . Following the discussion in Sect. 3.1 on multi-dimensional resampling for processes with many final state momenta of




Fig. 1 The distributions of $p_{\perp}^W$, y_W and $p_{\perp}^l$ obtained using the various samples. The sample of “positive only” events is included to illustrate the scale of the contribution from negative weight events

Table 1 The number of events generated in the sample of $W + 0j, 1j, 2j$ @NLO merged with UNLOPS (*weighted*), and the number of events arising from these by using the algorithms discussed in the paper

Sample	Total number of events	Events included in analysis
Weighted	5.3M	195k
Positive only	3.2M	121k
Unweighted	1.5M	52k
Positive Resampler(t_1)	659k	25k
Positive Resampler($p_{\perp}^W$)	33k	33k




Fig. 2 The transverse momentum of the hardest jet ($p_{\perp}^{j_1}$), second hardest jet ($p_{\perp}^{j_2}$), and third hardest jet ($p_{\perp}^{j_3}$). The lower right plot shows the convergence of the result of the Positive Resampler for an increasing number of dimensions using the shower evolution variable of up to three emissions

leptons and jets, it is unsurprising that a good description of the transverse momentum of the third jet requires sampling in more than just one dimension. Perhaps the surprising result is how few extra dimensions are required for a good description. The lower right pane on Fig. 2 shows the description of $p_{\perp}^{j_3}$ when the Positive Resampler is applied in an increasing number of dimensions from 1 to 3, using the shower evolution variable for the first, second and third hardest emission. The result converges to the *weighted* result (the input to the resampler) as the number of dimensions used in the resam-

pling is increased, and using three dimensions it is already well within the statistical fluctuations of the sample (some statistical fluctuation is obviously expected from the stochastic process of unweighting).

The result for the transverse momentum spectrum for the leading, subleading and the third leading jet is presented on the three remaining plots of Fig. 2. The result for the Positive Resampler in $p_{\perp}^W$ uses just one-dimensional resampling, whereas that of *Positive Resampler*(t_1, t_2, t_3) uses the three-dimensional resampling in the shower evolution variable.




Fig. 3 The transverse momentum of the hardest jet in three rapidity slices. See text for discussion

The curve for *positive only* illustrates again the significance of the negative weight events in the sample, and the portion of the events that can be removed. Resampling in three scalar dimensions is sufficient for the description of all the momenta.

As discussed in Sect. 3.3, the method can also be applied for multi-differential distributions. As an example, we show the distribution of the leading transverse jet momentum differentially in rapidity in Fig. 3. In spite of large fluctuations in the negative-weight contribution, the distributions agree in all bins with the original distribution obtained from weighted events (within the statistical variation introduced by the unweighting).

5 Conclusions

We presented the *Positive Resampler*, a method for modifying the weights of events drawn from a sample with both positive and negative weights, such that all events enter with positive weights, while kinematic distributions and observables are preserved exactly or within the statistical variations in the sample. The method was demonstrated using

reweighting in three different distributions and the impact on 6 independent observables studied. Since weight cancellations are handled explicitly, the reweighting can be implemented through a reduction of the number of events, thus allowing to significantly lower the cost and time associated with post-processing the event sample with subsequent analysis and modelling. The implementation of the reweighter used in this study will be publicly available after publication of this manuscript.

Acknowledgements The ideas presented in this work arose from the inspiring environment and discussions of the 2019 Les Houches Workshop “Physics at TeV Colliders”. The authors would like to express their gratitude to the organisers for their continued efforts in creating the stimulating milieu.

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), the EU TMR network SAGEX agreement No. 764850 (Marie Skłodowska-Curie), COST action CA16201: Unraveling new physics at the LHC through the precision frontier, and by the Swedish Research Council, contract number 2016-05996.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: This article describes an algorithm to transform pseudo-data. A prototype implementation of the algorithm can be made available upon request to the authors.]

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 SCOAP³.

References

1. J. Bellm et al., Eur. Phys. J. C **76**(4), 196 (2016). <https://doi.org/10.1140/epjc/s10052-016-4018-8>
2. T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C.O. Rasmussen, P.Z. Skands, Comput. Phys. Commun. **191**, 159 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024>
3. E. Bothmann et al., SciPost Phys. **7**(3), 034 (2019). <https://doi.org/10.21468/SciPostPhys.7.3.034>
4. S. Frixione, B.R. Webber, JHEP **06**, 029 (2002). <https://doi.org/10.1088/1126-6708/2002/06/029>
5. P. Nason, G. Ridolfi, JHEP **08**, 077 (2006). <https://doi.org/10.1088/1126-6708/2006/08/077>
6. S. Höche, F. Krauss, M. Schönherr, F. Siegert, JHEP **04**, 027 (2013). [https://doi.org/10.1007/JHEP04\(2013\)027](https://doi.org/10.1007/JHEP04(2013)027)
7. L. Lönnblad, S. Prestel, JHEP **03**, 166 (2013). [https://doi.org/10.1007/JHEP03\(2013\)166](https://doi.org/10.1007/JHEP03(2013)166)

8. L. Lönnblad, S. Prestel, JHEP **02**, 094 (2013). [https://doi.org/10.1007/JHEP02\(2013\)094](https://doi.org/10.1007/JHEP02(2013)094)
9. R. Frederix, S. Frixione, S. Prestel, P. Torrielli, (2020)
10. D. Rubin, J. Am. Stat. Assoc. **82**, 543 (1987)
11. J. Olsson, S. Plätzer, M. Sjö Dahl, (2019)
12. S. Plätzer, JHEP **08**, 114 (2013). [https://doi.org/10.1007/JHEP08\(2013\)114](https://doi.org/10.1007/JHEP08(2013)114)
13. K. Hamilton, P. Nason, C. Oleari, G. Zanderighi, JHEP **05**, 082 (2013). [https://doi.org/10.1007/JHEP05\(2013\)082](https://doi.org/10.1007/JHEP05(2013)082)
14. S. Alioli, C.W. Bauer, C.J. Berggren, A. Hornig, F.J. Tackmann, C.K. Vermilion, J.R. Walsh, S. Zuberi, JHEP **09**, 120 (2013). [https://doi.org/10.1007/JHEP09\(2013\)120](https://doi.org/10.1007/JHEP09(2013)120)
15. R. Frederix, S. Frixione, JHEP **12**, 061 (2012). [https://doi.org/10.1007/JHEP12\(2012\)061](https://doi.org/10.1007/JHEP12(2012)061)
16. J. Bellm, S. Gieseke, S. Plätzer, Eur. Phys. J. C **78**(3), 244 (2018). <https://doi.org/10.1140/epjc/s10052-018-5723-2>
17. S. Alioli, A. Broggio, S. Kallweit, M.A. Lim, L. Rottoli, Phys. Rev. D **100**(9), 096016 (2019). <https://doi.org/10.1103/PhysRevD.100.096016>
18. R. Frederix, K. Hamilton, JHEP **05**, 042 (2016). [https://doi.org/10.1007/JHEP05\(2016\)042](https://doi.org/10.1007/JHEP05(2016)042)
19. J.R. Christiansen, S. Prestel, Eur. Phys. J. C **76**(1), 39 (2016). <https://doi.org/10.1140/epjc/s10052-015-3871-1>
20. N. Fischer, S. Prestel, Eur. Phys. J. C **77**(9), 601 (2017). <https://doi.org/10.1140/epjc/s10052-017-5160-7>
21. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.S. Shao, T. Stelzer, P. Torrielli, M. Zaro, JHEP **07**, 079 (2014). [https://doi.org/10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079)
22. C. Bierlich et al., SciPost Phys. **8**, 026 (2020). <https://doi.org/10.21468/SciPostPhys.8.2.026>
23. S. Höche, S. Prestel, Eur. Phys. J. C **75**(9), 461 (2015). <https://doi.org/10.1140/epjc/s10052-015-3684-2>