EPJ Web of Conferences **337**, 01322 (2025) https://doi.org/10.1051/epjconf/202533701322
CHEP 2024

Monte Carlo efficiency via negative weight reduction

James Whitehead[1]*

[1]Jagiellonian University, ul. prof. Stanislawa Łojasiewicza 11, 30-348 Kraków, Poland

Abstract. The generation of large event samples with Monte Carlo Event Generators is expected to be a computational bottleneck for precision phenomenology at the HL-LHC and beyond. This is due in part to the computational cost incurred by negative weights in 'matched' calculations combining NLO perturbative QCD with a parton shower: for the same target statistical precision, a larger sample must be generated. We summarise two approaches taken to tackle this problem: the development of the KrkNLO matching method, which uses a redefinition of the PDF factorisation scheme to guarantee positive weights by construction, and the restructuring of the Matchbox module to reduce the fraction of negative weights for Mc@Nlo matching.

1 Introduction

Parton shower algorithms provide an invaluable complement to fixed-order perturbation theory for the prediction of LHC observables. While perturbative QCD excels for low-multiplicity observables, parton showers use the universal factorisation of matrix elements in their singular limits to compute approximate contributions to high-multiplicity observables, at higher formal orders than are yet calculable. Combining the two approaches while retaining strict perturbative accuracy is called 'matching': the correct fixed-order cross-section (to a given order of perturbative QCD) is smeared into other higher-multiplicity bins.

This problem of matching parton showers to NLO calculations has long been solved in general by the Mc@Nlo [1] and Powheg [2–4] methods. Both have achieved success as workhorses of LHC phenomenology. Mc@Nlo uses a modified form of subtraction, in which the parton shower itself acts as a local counterterm. Whilst theoretically unproblematic, in practice this risks oversubtraction. If 'unweighted' events are being generated, of constant weight, for example for detector simulation, any negative weights spoil the statistical convergence of the sample, requiring many more events to be generated to achieve the same overall uncertainty. This problem is entirely an artefact of the method, and a variety of techniques have been proposed to address it [5–11].

The Powheg method, which by construction generates only positive weights, requires a modification of the shower algorithm for the generation of the first emission and hence unphysically exponentiates the real-emission matrix-element far from the singular regions [12]. Whilst this can be mitigated by so-called 'damping factors' [4, 13], this introduces an additional theoretical uncertainty, from the dependence of the resulting prediction on the (unphysical) choice of function used for the damping.

*e-mail: james.whitehead@uj.edu.pl

Recently, the KrkNLO method [14–16] has been proposed, at present restricted to the production of colour-singlet final-states. Like Powheg, it generates positive-weighted events by construction; like Mc@Nlo, it uses an unmodified shower algorithm.

In this proceedings, in section 2 we summarise the effort to develop the KrkNLO method and investigate its properties through our implementation in Herwig 7; in section 3 we summarise ongoing efforts to modify the implementation of Mc@Nlo in Matchbox [17], the matching module within Herwig 7, to reduce the fraction of negative weights.

2 KrkNLO for negative-weight reduction[1]

The KrkNLO method for NLO parton shower matching was introduced in [14, 15]. It exploits the freedom, at NLO and beyond, to choose a factorisation scheme for the parton distribution functions (PDFs) to absorb the collinear contributions arising from the phase-space integral of the parton shower-kernels into 'Krk'-scheme PDFs [14, 18, 19]. This is sufficient to satisfy the NLO matching condition, for colour-singlet final-states [16]. In addition to its negative-weight reduction, the KrkNLO method also enjoys the advantage of simplicity and efficiency, and the theoretical advantage of applying a consistent Sudakov factor corresponding to emissions from the underlying Born to all contributions no matter their origin.

As an algorithm, the KrkNLO method may be summarised as:

---

for all Born events **do** shower
  if first emission generated, from kernel (α) **then**
    $w \leftarrow w \times \frac{\mathrm{R}(\Phi_{m+1})}{P_m^{(\alpha)}(\Phi_{m+1})}$
  end if
  $w \leftarrow w \times \left[1 + \frac{\alpha_s(\mu_\mathrm{R})}{2\pi}\left(\frac{\mathrm{V}(\Phi_m;\mu_\mathrm{R})}{\mathrm{B}(\Phi_m)} + \frac{\mathrm{I}(\Phi_m;\tilde{\mu}_\mathrm{R})}{\mathrm{B}(\Phi_m)} + \Delta_0^{\mathrm{FS}}\right)\right]$
end for

The shower is then allowed to run to completion. Here Φ_m denotes the Born phase-space and Φ_{m+1} the real-emission phase-space. B, R and V denote the relevant Born, real and virtual matrix elements respectively; $P_m^{(\alpha)}$ the emission kernel used in the shower algorithm to generate the chosen branching of generalised type (α); μ_R the renormalisation scale, and I the contribution from the shower Sudakov, integrated over the radiative phase-space, containing no residual collinear dependence. Δ_0^{FS} is a factorisation-scheme-dependent virtual-type contribution which compensates for any residual delta-function contributions in the factorisation-scheme transformations.[2] Further details are given in [16].

In order for the full real-emission phase-space Φ_{m+1} to be populated by the reweighted first shower emission, the shower must be constructed to have full phase-space coverage (i.e., no 'dead zones').[3] This is true, for example, of the Herwig dipole-shower, but not the angular-ordered shower. The shower starting-scale must also be chosen to be unrestricted, i.e. to coincide with the maximum kinematically-attainable scale for each dipole. In practice, this amounts to allowing all kinematically-valid splittings, so need not be calculated and set explicitly for each phase-space-configuration.

[1] Based on work in collaboration with Andrzej Siódmok and Pratixan Sarmah.

[2] This is the only type of contribution from the factorisation scheme transformation that is allowed to remain in the hard-process calculation, subject to the constraint of no collinear convolution terms, since it carries no x-dependence.

[3] This can be relaxed by using the Mc@Nlo method to fill in the remaining phase-space [20].

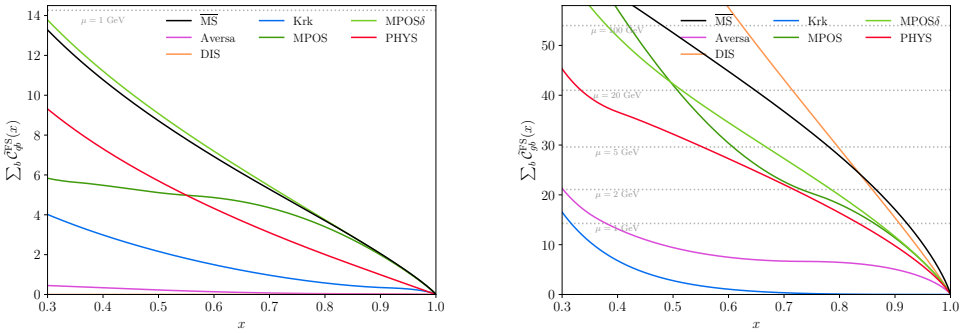

Figure 1. Cumulant plots from [19], using an argument adapted from [21]. Following the reasoning of [21], the intersections of the grey scale-markings with the coloured lines give indicative scales above which the specified PDF in the given scheme may be expected to be positive, for x-values greater than those of the point of intersection. A lower coloured line therefore implies a lower scale, which implies improved positivity characteristics; note the favourable position of the Krk scheme. Refer to [19] for further details.

2.1 PDFs in the Krk scheme[4]

The ability of the KrkNLO method to prevent the emergence of negative weights from within the parton-shower matching process would be undone if the PDF factorisation scheme transformation made the transformed PDFs negative, introducing a different source of negative weights into the calculation from the convolution with the incoming PDFs.

For that reason, in [19] we studied the properties, including the positivity properties, of the Krk scheme (amongst others). By adapting an argument advanced for the $\overline{\text{MS}}$ scheme in [21], we investigated the Krk scheme, both theoretically, and empirically through the study of transformed $\overline{\text{MS}}$ PDFs from the major fitting groups. We found that the theoretical arguments advanced for $\overline{\text{MS}}$ positivity implied a stronger bound (that is, a lower scale above which positivity should hold) for the Krk scheme than for the $\overline{\text{MS}}$ scheme. The 'cumulants' which imply this scale are shown in fig. 1.

This argument was developed primarily for massless quarks (using massless coefficient functions), and empirically the main origin of possible negativity arose in a narrow scale interval above the mass threshold for heavy-quark flavours. This is an artefact of the perturbative transformation from the $\overline{\text{MS}}$ scheme, which introduces a discontinuity at the threshold, and not necessarily an intrinsic property of the scheme.

2.2 Phenomenology of KrkNLO

The Mc@Nlo and Powheg methods for NLO parton-shower-matching both often rely in practice on unphysical choices. For Mc@Nlo, the starting-scale of the shower represents a choice which in practice can dramatically affect the predictions generated; for Powheg the functional-form chosen as a 'damping factor' to suppress the exponentiation of the real-emission matrix element away from the singular limit likewise introduces an uncertainty arising from an unphysical choice. While reasonable arguments on physical grounds can be made for the choices commonly adopted in the literature, there remains a family of reasonable choices whose impact must be estimated quantitatively by parameter variation.

[4]Based on work in collaboration with Stéphane Delorme, Aleksander Kusina and Andrzej Siódmok.

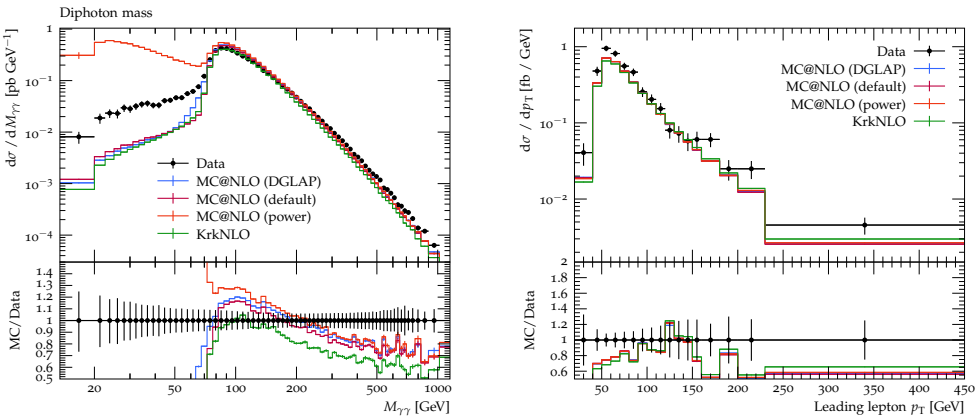

Figure 2. Matched NLO-plus-parton-shower predictions for ATLAS data for the diphoton [16] and *ZZ*-production processes respectively, for KrkNLO and for Mc@Nlo with several alternative choices for the shower starting-scale.

The KrkNLO method is unique in that, by contrast, it involves no choice of starting scale (NLO accuracy requires that the first-emission be unrestricted, a so-called 'power shower'), and requires no damping factor (since only the 'true' dipole-shower kernels are exponentiated). This provides a useful independent cross-check on common choices of parameters within the alternative methods.

In [16] we presented a detailed comparison of predictions made with KrkNLO to those produced with three variants of Mc@Nlo. These predictions were configured to be maximally comparable, using the same PDF set, the same renormalisation and factorisation scale choices (for consistency with KrkNLO, set to the invariant mass of the colourless final-state system), and the same parton shower algorithm (the Herwig 7 dipole shower) with the same infrared cut-off, and with hadronisation disabled. The difference can therefore be identified as arising solely from the matching method.

We present a small selection of plots highlighting the differences between the KrkNLO and Mc@Nlo methods for the diphoton and *ZZ*-production process in fig. 2.

2.3 KrkNLO weight distribution

In contrast with Mc@Nlo, the KrkNLO method does not tame the underlying singularities of the real-emission matrix-element by subtraction. While this allows it to avoid over-subtraction, hence preventing the introduction of negative weights, it also leads to a broader weight-distribution than that typically encountered within Mc@Nlo. In principle this reduces the unweighting efficiency, implying that the elimination of negative weights could be at the cost of additional computational cost in the event generation stage. For purposes such as detector simulation, in which the computational cost per event exceeds the generation cost by several orders of magnitude, this could nevertheless represent a worthwhile trade-off.

We have not yet performed a detailed study of this trade-off in practice, but present plots of the weight distribution for diphoton production, *WW*-production and *ZZ*-production in fig. 3.

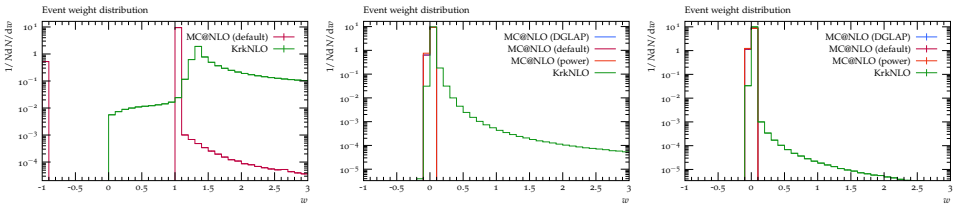

Figure 3. The raw weight distributions for (weighted) NLO-matched full-shower events generated by KrkNLO and Mc@Nlo, for the diphoton-, WW- and ZZ-production processes respectively. For diphoton-production, the Mc@Nlo prediction is unweighted.

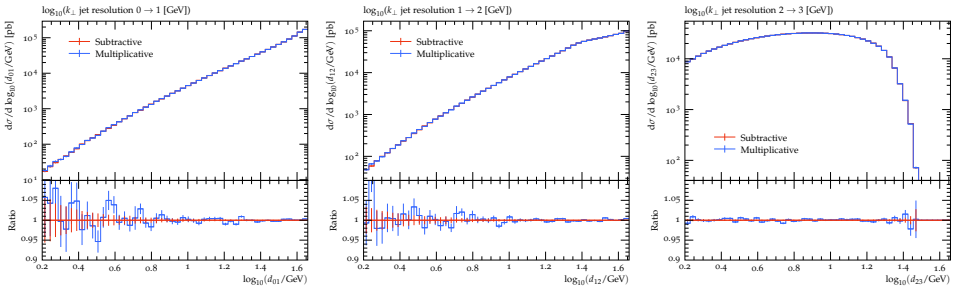

Figure 4. Preliminary validation plots for the alternative implementation of the Mc@Nlo 'bridge' cross-section within Matchbox. These plots correspond to a matched NLO-plus-parton-shower calculation of $e^+e^- \to 2j$.

3 MC@NLO with Matchbox[5]

The Mc@Nlo method is implemented within Herwig 7 through the Matchbox module [17]. This is designed to provide a flexible implementation to allow NLO matching with parton showers through both Mc@Nlo and Powheg, for both the Herwig 'default' angular-ordered shower [22] and the Herwig dipole shower.

By design, it generates the shower-subtracted real-emission and the so-called 'bridge' cross-section in the same way within Matchbox, starting from the real-emission kinematics. 'Bridge' contributions, the difference between the shower subtraction and the exact (dipole-type) subtraction contributions are commonly negative, and as a separately-generated independent process, these contributions can produce negative weights upon unweighting.

We are therefore exploring an alternative approach to the generation of such 'bridge' contributions, combining them instead with the Born process as a multiplicative weight. The positivity of the tree-level Born matrix-element can then counteract the possible negativity of the 'bridge' contribution. Preliminary validation plots for the alternative implementation are shown in fig. 4 for NLO-plus-parton-shower Mc@Nlo predictions for $e^+e^- \to 2j$.

Whilst this is primarily an equivalent alternative implementation of the same underlying method, it should additionally allow the further exploration of matching uncertainties.

[5] Based on work in collaboration with Simon Plätzer and Andrzej Siódmok.

4 Conclusions

In this contribution I have summarised ongoing work to pursue theoretical methods for NLO parton shower matching which, in addition to being of purely theoretical interest, serve to reduce the negative-weight fraction of matched parton-shower predictions. While such work has to date been confined to NLO calculations, and in the case of the KrkNLO method to the production of colour-singlet final-states, we anticipate that the underlying ideas admit extension to more complicated processes and higher orders of perturbation theory.

The theoretical issues raised, such as the role, utility, and universality of alternative factorisation schemes, are likely to prove fruitful avenues for exploration in their own right.

References

[1] S. Frixione, B.R. Webber, Matching NLO QCD computations and parton shower simulations, JHEP **06**, 029 (2002), hep-ph/0204244. 10.1088/1126-6708/2002/06/029

[2] P. Nason, A New method for combining NLO QCD with shower Monte Carlo algorithms, JHEP **11**, 040 (2004), hep-ph/0409146. 10.1088/1126-6708/2004/11/040

[3] S. Frixione, P. Nason, C. Oleari, Matching NLO QCD computations with Parton Shower simulations: the POWHEG method, JHEP **11**, 070 (2007), 0709.2092. 10.1088/1126-6708/2007/11/070

[4] S. Alioli, P. Nason, C. Oleari, E. Re, A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX, JHEP **06**, 043 (2010), 1002.2581. 10.1007/JHEP06(2010)043

[5] K. Danziger, S. Höche, F. Siegert, Reducing negative weights in Monte Carlo event generation with Sherpa (2021), 2110.15211.

[6] R. Frederix, S. Frixione, S. Prestel, P. Torrielli, On the reduction of negative weights in MC@NLO-type matching procedures, JHEP **07**, 238 (2020), 2002.12716. 10.1007/JHEP07(2020)238

[7] R. Frederix, P. Torrielli, A new way of reducing negative weights in MC@NLO, Eur. Phys. J. C **83**, 1051 (2023), 2310.04160. 10.1140/epjc/s10052-023-12243-x

[8] J.R. Andersen, C. Gütschow, A. Maier, S. Prestel, A Positive Resampler for Monte Carlo events with negative weights, Eur. Phys. J. C **80**, 1007 (2020), 2005.09375. 10.1140/epjc/s10052-020-08548-w

[9] J.R. Andersen, A. Maier, Unbiased elimination of negative weights in Monte Carlo samples, Eur. Phys. J. C **82**, 433 (2022), 2109.07851. 10.1140/epjc/s10052-022-10372-3

[10] J.R. Andersen, A. Maier, D. Maître, Efficient negative-weight elimination in large high-multiplicity Monte Carlo event samples, Eur. Phys. J. C **83**, 835 (2023), 2303.15246. 10.1140/epjc/s10052-023-11905-0

[11] P. Shyamsundar, ARCANE Reweighting: A Monte Carlo Technique to Tackle the Negative Weights Problem in Collider Event Generation (2025), 2502.08052.

[12] S. Höche, F. Krauss, M. Schönherr, F. Siegert, A critical appraisal of NLO+PS matching methods, JHEP **09**, 049 (2012), 1111.1220. 10.1007/JHEP09(2012)049

[13] S. Alioli, P. Nason, C. Oleari, E. Re, NLO Higgs boson production via gluon fusion matched with shower in POWHEG, JHEP **04**, 002 (2009), 0812.0578. 10.1088/1126-6708/2009/04/002

[14] S. Jadach, W. Płaczek, S. Sapeta, A. Siódmok, M. Skrzypek, Matching NLO QCD with parton shower in Monte Carlo scheme — the KrkNLO method, JHEP **10**, 052 (2015), 1503.06849. 10.1007/JHEP10(2015)052

[15] S. Jadach, G. Nail, W. Płaczek, S. Sapeta, A. Siódmok, M. Skrzypek, Monte Carlo simulations of Higgs-boson production at the LHC with the KrkNLO method, Eur. Phys. J. C **77**, 164 (2017), 1607.06799. 10.1140/epjc/s10052-017-4733-9

[16] P. Sarmah, A. Siódmok, J. Whitehead, KrkNLO matching for colour-singlet processes, JHEP **01**, 062 (2025), 2409.16417. 10.1007/JHEP01(2025)062

[17] S. Plätzer, S. Gieseke, Dipole Showers and Automated NLO Matching in Herwig++, Eur. Phys. J. C **72**, 2187 (2012), 1109.6256. 10.1140/epjc/s10052-012-2187-7

[18] S. Jadach, W. Płaczek, S. Sapeta, A. Siódmok, M. Skrzypek, Parton distribution functions in Monte Carlo factorisation scheme, Eur. Phys. J. C **76**, 649 (2016), 1606.00355. 10.1140/epjc/s10052-016-4508-8

[19] S. Delorme, A. Kusina, A. Siódmok, J. Whitehead, Factorisation schemes for proton PDFs (2025), 2501.18289.

[20] P. Nason, G.P. Salam, Multiplicative-accumulative matching of NLO calculations with parton showers, JHEP **01**, 067 (2022), 2111.03553. 10.1007/JHEP01(2022)067

[21] A. Candido, S. Forte, T. Giani, F. Hekhorn, On the positivity of $\overline{\text{MS}}$ parton distributions, Eur. Phys. J. C **84**, 335 (2024), 2308.00025. 10.1140/epjc/s10052-024-12681-1

[22] S. Gieseke, P. Stephens, B. Webber, New formalism for QCD parton showers, JHEP **12**, 045 (2003), hep-ph/0310083. 10.1088/1126-6708/2003/12/045