

Fast interpolation grids for the Drell–Yan process

Juan Cruz–Martinez¹, Alexander Huss¹, and Christopher Schwan²

¹ CERN, Theoretical Physics Department, CH-1211 Geneva 23, Switzerland

² Universität Würzburg, Institut für Theoretische Physik und Astrophysik, 97074 Würzburg, Germany

Date: 24th January 2025

Abstract Modern analyses of experimental data from hadron colliders rely on theory predictions at high orders in perturbation theory and a variety of input settings. Interpolation grids facilitate an almost instant re-evaluation of theory predictions for different input parton distributions functions (PDFs) or scale settings and are thus indispensable in the study of the parton content of the proton. While interpolation grids at next-to-next-to-leading order (NNLO) exist for some key processes relevant for PDF determinations, a notable exception is the Drell–Yan process that constitutes the production of electroweak gauge bosons at hadron colliders and provides important constraints on the quark content of the proton. To address this gap, we report on a new interface between the parton-level Monte Carlo generator NNLOJET and the interpolation grid library PINEAPPL and demonstrate its use for the Drell–Yan process. Accompanying this note, we release Drell–Yan grids covering a wide range of measurements that commonly enter global determinations of PDFs. We use the grids to study accidental cancellation between partonic channels at NNLO and inspect the validity of a K -factor approximation that was widely employed previously.

1 Introduction

Over the past decades, the unprecedented precision achieved by the experiments at the Large Hadron Collider (LHC), as well as the projections of its high-luminosity upgrade (HL-LHC), has driven remarkable theoretical advances, pushing the frontier of perturbative Quantum Chromodynamics (pQCD) to next-to-next-to-leading order (NNLO) accuracy and beyond. Given the high computational resources consumed by such state-of-the-art calculations, efficient methods to re-evaluate predictions for different inputs are of critical importance. Of particular interest are fast interpolation grids, which preserve the dependence of the partonic cross sections on the renormalisation and factorisation scales, the longitudinal momentum fractions carried by the partons within the colliding hadrons, and the strong coupling α_s . The generation and usage of interpolation grids are facilitated by libraries such as APPLgrid [1], fastNLO [2,3], and PINEAPPL [4,5]. These grids enable fast evaluations of the cross section for different choices of parton distribution functions (PDFs) that would otherwise require a costly re-calculation, and are thus indispensable for the determination of PDFs.

Interpolation grids are available at NNLO for various processes, e.g., for jet production in deep-inelastic scattering (DIS) [6,7], jet production in hadron collisions [8], top-quark pair production [9,10], and inclusive DIS production [11,12]. A notable absence in this list is the Drell–Yan (DY) process, which constitutes the production of electroweak gauge bosons at hadron colliders. While grids were made available for inclusive fixed-target DY [13], no

differential predictions are available despite it being a process of high importance for PDF determinations, accounting for roughly 20% of the total dataset in recent global fits [14–16].

Moreover, the Drell–Yan process [17] is a high-precision probe at hadron colliders and it has been known fully differentially at NNLO for a long time [18–21], and most recently also computed at N³LO both at the inclusive [22–25] and differential level [26–30]. Despite the availability of predictions, no differential grids are available at NNLO and so current PDF determinations commonly rely on an approximation based on NLO grids supplemented with NNLO K -factors obtained for a specific choice of PDF. In this note we interface the parton-level event generator NNLOJET [31] that is based on the antenna subtraction formalism [32–36] with the PINEAPPL [4,5] grid-interpolation library. This enables the production of interpolation grids for all processes included in the NNLOJET code in a variety of scenarios. We demonstrate the use of this new interface using the DY process as an example. Together with this paper we make available a wide range of grids that correspond to the entire set of DY data that enters the NNPDF4.0 [16] determination of PDFs. The grids are available to download at:¹

<https://ploughshare.web.cern.ch>

This note is structured as follows: in Section 2 we describe the grids that we make available and provide a

¹ Grids are provided in the PINEAPPL format; conversion *from* and *to* (`pineappl [import|export] --help`) this format facilitate its use as an universal converter.

summary of the metadata that is included. We demonstrate the accuracy of the interpolation and provide some examples of studies that can be readily performed with the PINEAPPL command-line interface (CLI). We end in Section 3 with new insights into the known feature of accidental cancellations at NNLO and assess the impact of the widely used K -factor approximation in the context of global PDF analyses.

2 Interpolation grids

The general idea of interpolation grids follows from the property of QCD factorisation, where hadronic cross sections can be written in terms of a convolution of the partonic cross section with the PDFs. Approximating the PDFs through a decomposition into a set of eigenfunctions, a grid can be produced that corresponds to the convolution of said eigenfunctions with the partonic cross section. With this at hand, the evaluation of a hadronic cross section reduces to performing a sum over the grid entries weighted by the PDFs at the given nodes, thus, substantially speeding up the evaluation. Further details on the grid-interpolation techniques as implemented in the PINEAPPL library can be found in Ref. [4].

Accompanying this manuscript, we release a large set of interpolation grids for the Drell–Yan process corresponding to the measurements summarised in Table 1 that were produced using the implementation of this process within the NNLOJET framework [59]. All predictions are produced for the central scale choice $\mu_R = \mu_F = E_{T,V}$, where the transverse energy $E_{T,V}$ is defined in terms of the invariant mass of the intermediate electroweak gauge boson and its transverse momentum, $E_{T,V} = \sqrt{m_V^2 + p_{T,V}^2}$. This choice largely coincides with the invariant mass m_V in inclusive quantities, however, it accounts for the impact of hard QCD emissions in phase-space regions that are sensitive to it.

Lastly, all grids evaluate to absolute predictions in fb units with appropriate bin-width normalisations following the HEPData entries provided with the corresponding measurements. No composition of observables is performed in order to maintain the full flexibility and granularity that these grids offer. In particular, this means that in order to obtain normalised distributions, such as $(1/\sigma)d\sigma/d\mathcal{O}$, this has to be performed by the user by appropriately summing the cross section of the individual bins.

2.1 Metadata

The interpolation grids made available alongside this publication contain the following metadata that are accessible through the PINEAPPL CLI:²

`nnlojet_runcard` — A sample NNLOJET runcard used to produce the grid.

² `pineappl read <grid> --get <key>`

`nnlojet_version` — Version of NNLOJET used.

`pineappl_gitversion` — Version of PINEAPPL used.

`result` — Reference numbers as reported by NNLOJET broken down into different perturbative orders together with associated Monte Carlo integration errors.

`results_pdf` — The PDF set used to generate the results. This information is necessary in order to perform closure tests shown in Section 2.2.

`hepdata` — HEPData entry of the dataset for which the grids were produced, c.f. Table 1.

2.2 Closure tests

One of the main advantages of interpolation grids is the relatively small storage requirement associated with them. This is in contrast to approaches based on storing separate collision events, such as Ntuples [60, 61] or HighTEA [62], which often have a significantly larger storage footprint. The smaller size is achieved by both fixing the setup (event selection and binning of histograms) and thus giving up on some flexibility, as well as the approximation through interpolation. The latter demands that any systematic uncertainty introduced by the interpolation to be well below other uncertainties in the calculation. In the case of the Drell–Yan process, where NNLO corrections are known to be very small and plagued by large numerical cancellations, assessing the quality of the grids is particularly important.

In Fig. 1 we present closure tests that compare predictions obtained from evaluating interpolation grids against exact reference numbers from the calculation used to generate them. Any observed difference is thus solely due to interpolation errors. To this end, we decompose the NNLO predictions into separate contributions

$$d\sigma^{\text{NNLO}} = d\sigma^{\text{LO}} + d\sigma^{\delta\text{NLO}} + d\sigma^{\delta\text{NNLO}},$$

and perform the closure test on each of the perturbative *coefficients*. Figure 1a highlights that the interpolation errors are typically well below the per-mille level in the bulk of the phase space. However, deviations can become as large as few per-mille in extreme cases such as the forward region of the LHCb experiment that probes a much larger range of the momentum fractions as seen in Fig. 1b. It should be emphasized that $d\sigma^{\delta\text{NNLO}}$ typically amounts to a few percent of the full NNLO prediction and an interpolation error of 1‰ on the coefficient translates to a 0.01‰ level of exactitude on the final results. In contrast, the residual Monte Carlo errors on the predictions are at the few ‰ level, which in turn are already negligible with respect to the experimental uncertainties of the measurement. For any phenomenological application, the impact of interpolation errors in the provided grids is thus completely negligible.

2.3 Uncertainties

A common use-case for grid files is the study of various sources of uncertainties that require re-evaluating predictions for different input parameters. The PINEAPPL CLI

| Dataset | Ref. | N_{dat} | Cuts |
|--|------|------------------|---|
| CDF Z differential | [37] | 29 | $0.0 \leq y_{\ell\ell} \leq 2.9, 66 \leq m_{\ell\ell} \leq 116$ |
| D0 Z differential | [38] | 28 | $0.0 \leq y_{\ell\ell} \leq 2.8, 66 \leq m_{\ell\ell} \leq 116$ |
| D0 W electron asymmetry | [39] | 13 | $0.0 \leq y_e \leq 2.9$ |
| D0 W muon asymmetry | [40] | 10 | $0.0 \leq y_\mu \leq 1.9$ |
| ATLAS low-mass DY 7 TeV | [41] | 6 | $ \eta_\ell \leq 2.1, 14 \leq m_{\ell\ell} \leq 56$ |
| ATLAS high-mass DY 7 TeV | [42] | 13 | $ \eta_\ell \leq 2.1, 116 \leq m_{\ell\ell} \leq 1500$ |
| ATLAS W, Z 7 TeV ($\mathcal{L} = 35 \text{ pb}^{-1}$) | [43] | 30 | $ \eta_\ell, y_Z \leq 3.2$ |
| ATLAS W, Z 7 TeV ($\mathcal{L} = 4.6 \text{ fb}^{-1}$) | [44] | 61 | $ \eta_\ell, y_Z \leq 2.5, 3.6$ |
| ATLAS W 8 TeV | [45] | 22 | $ \eta_\ell < 2.4$ |
| ATLAS low-mass DY 2D 8 TeV | [46] | 84 | $ y_{\ell\ell} < 2.4, 46 \leq m_{\ell\ell} \leq 200$ |
| ATLAS high-mass DY 2D 8 TeV | [47] | 48 | $ y_{\ell\ell} < 2.4, 116 \leq m_{\ell\ell} \leq 1500$ |
| CMS W electron asymmetry 7 TeV | [48] | 11 | $ \eta_\mu \leq 2.4$ |
| CMS W muon asymmetry 7 TeV | [49] | 11 | $ \eta_\mu \leq 2.4$ |
| CMS DY 2D 7 TeV | [50] | 132 | $ \eta_{\ell\ell} \leq 2.2, 20.0 \leq m_{\ell\ell} \leq 200$ |
| CMS W rapidity 8 TeV | [51] | 22 | $ \eta_\ell \leq 2.3$ |
| LHCb $Z \rightarrow ee$ 7 TeV | [52] | 9 | $2.0 \leq \eta_\ell \leq 4.5$ |
| LHCb $W, Z \rightarrow \mu$ 7 TeV | [53] | 33 | $2.0 \leq \eta_\ell \leq 4.5$ |
| LHCb $Z \rightarrow ee$ 8 TeV | [54] | 17 | $2.00 < \eta_e < 4.25$ |
| LHCb $W, Z \rightarrow \mu$ 8 TeV | [55] | 34 | $2.00 < \eta_\mu < 4.25$ |
| LHCb $W \rightarrow e$ 8 TeV | [56] | 8 | $2.00 < \eta_e < 4.25$ |
| ATLAS $\sigma_{W,Z}^{\text{tot}}$ 13 TeV | [57] | 3 | — |
| LHCb $Z \rightarrow ee$ 13 TeV | [58] | 17 | $2.00 < y_Z < 4.25$ |
| LHCb $Z \rightarrow \mu\mu$ 13 TeV | [58] | 18 | $2.00 < y_Z < 4.50$ |

Table 1: Collider Drell–Yan datasets considered in the NNPDF4.0 global PDF determination and selected for this release of interpolation grids. Table of kinematic cuts from Ref. [16].

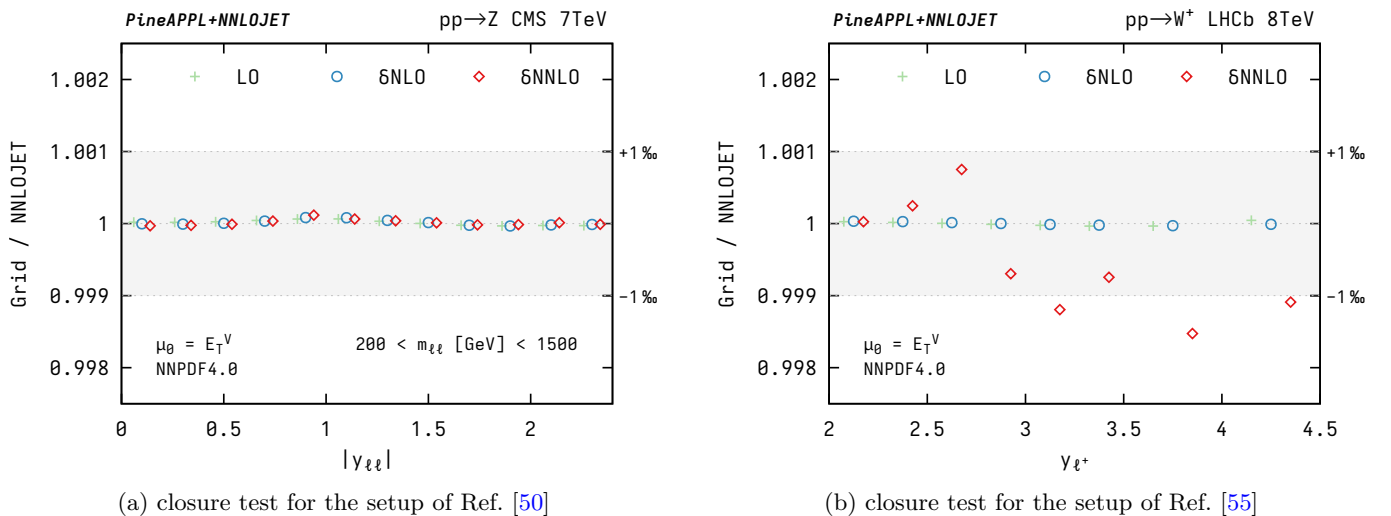


Figure 1: Grid closure between the interpolation and the exact NNLOJET reference numbers for the datasets of Refs. [50] (a) and [55] (b) at LO and the NLO and NNLO coefficients.

provides efficient tools for analysing grids in various scenarios. Below, we demonstrate some representative analyses using the built-in utilities.

Scale variations are the most widely used approach to estimate theory uncertainties from missing higher orders in the perturbative expansion. The PINEAPPL grids allow to vary the renormalisation and factorisation scales, μ_R and μ_F respectively, through any linear transformation of

the original scale used for the grid generation. This offers the possibility to explore the (μ_R, μ_F) space with high granularity.

The conventional variation by factors of $(\frac{1}{2}, 2)$ is directly provided within the CLI through the command³

³ Instead of the 9-point variation, also the 3- and 7-point variations are available, which can be specified through the `--scale-abs` option.




Figure 2: 9-point scale variations for the Z data from Ref. [43], evaluated using the PDF set (NNPDF40_nnlo_as_01180) at all orders.

```
pineappl uncert <grid> <pdf>
--scale-abs=9 --orders <orders>
```

where `<orders>` specifies the perturbative orders to be considered, e.g. `--orders a2,as1a2` for NLO in the DY process. In Fig. 2 we show as an example the 9-point scale variation uncertainties for the Z production calculation for the measurement of Ref. [43].

PDF uncertainties are derived from a collection of predictions obtained via a convolution with PDFs of an associated error set. Interpolation grids provide detailed information across several parameters, not only the renormalisation and factorisation scales, to facilitate the efficient a posteriori re-evaluation of the predictions for arbitrary PDF sets and thus the study of PDF uncertainties as well as their use in PDF fits. To this end, interpolation grids additionally retain the information on the longitudinal momentum fractions of the colliding partons, x_1 and x_2 , separately for each independent combination of partonic channels of the corresponding process.

In Fig. 3, we contrast the PDF uncertainties of different sets at NNLO from the MHST [15], CTEQ [14] and NNPDF [16] fitting groups. The figure is obtained with the CLI command

```
pineappl plot -s 3 <grid> <pdf1> <pdf2>...
```

which is used together with the option `-s 3` to specify the 3-point scale variation in this case. The output of the command is a Python script to generate a figure as shown in

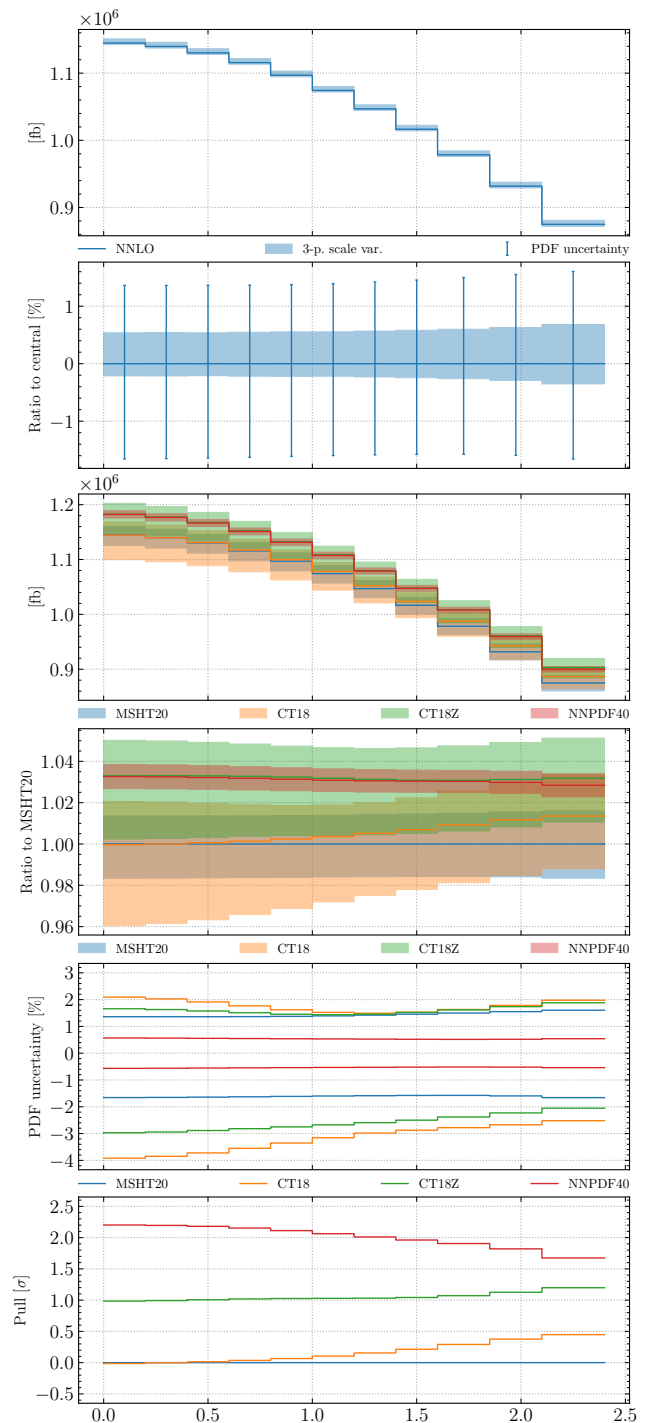


Figure 3: Scale and PDF uncertainties for the W^- data of Ref. [51]. Using interpolation grids we can obtain several different analyses with a single command of `PineAPPL` in a matter of seconds.

Fig. 3 and is composed of several panels: The first two panels show separately the scale- and PDF-uncertainties of the predictions using `<pdf1>` (in this case MSHT20) for the absolute prediction and the relative uncertainties. This is followed by three panels that show the PDF uncertainties for all PDF sets that were specified, which are provided as absolute predictions and ratio plots with respect to the central `<pdf1>`. The final two panels directly contrast the PDF uncertainties of the different sets together with the pull in units of σ for each PDF, using `<pdf1>` as reference.

By utilising interpolation grids, analyses can be made with multiple PDF set at zero added computational cost. The importance of these extra studies is clear e.g., in the 4th and 6th panels of Fig. 3. While the pull between sets can be as large as 2σ , and thus the choice of PDF might have phenomenological implications, this difference is washed out once several sets are included in the analysis.

3 Drell–Yan phenomenology

3.1 Accidental cancellations

One peculiar feature of the DY process that has been the subject of various investigations are the large accidental cancellation between the $q\bar{q}$ and $(qg + \bar{q}g)$ partonic channels. This cancellation is especially severe at NNLO, in part causing the scale uncertainties to be underestimated and for the N³LO predictions to often lie outside of the NNLO uncertainty estimates [22]. This striking trait of non-overlapping scale-uncertainty bands is particularly pronounced at higher hadron-collider energies and was found to be largely independent of whether pp or p \bar{p} collisions are considered [25]. These observations hint towards cancellations that are likely driven by correlations between the gluon and sea-quark distributions.

Interpolation grids offer the unique opportunity to study such hypotheses in more detail by providing the possibility to evaluate the predictions at different orders and broken apart into individual partonic channels (using different bases), combined with the flexibility to evolve the prediction to any value for the scales. Practically, this is achieved by evolving the grids with the DGLAP evolution library EKO [63, 64] to different values in μ_F , a feature provided by the PINEAPPL CLI:⁴

```
pineappl evolve <grid>
  <evolution_kernel_operator.tar>
  <output_evolved_grid>
  <pdf_to_check> --orders a2as2
```

where the evolution kernel operator is generated with the `pineko` program [13] and the option `--orders a2as2` ensures that only the NNLO QCD contribution is being considered.

⁴ Note that it might be necessary to set a high value for the accuracy if the `<pdf_to_check>` used for the check has not been evolved with the exact same settings as the kernel operator.

Figure 4a shows the NNLO contribution $\sigma^{\delta\text{NNLO}}$ separated into partonic channels as a function of the factorisation scale μ_F . The results correspond to the total cross section obtained with the setup of the measurement in Ref. [57]. The figure clearly highlights the large cancellation between the $q\bar{q}$ and $(qg + \bar{q}g)$ channels, with a compensation of almost two orders of magnitude around $\mu_F \sim M_Z$. We further observe a strong dependence on the factorisation scale that enhances the cancellation for larger μ_F values, which hints at an underlying correlation in the DGLAP evolution that drives this feature.

In order to gain further insights into this correlation and its impact, we transform to the so-called evolution basis⁵ by rotating the evolution operator with EKO, which serves the purpose of decoupling the evolution of the different independent components. In particular, the gluon evolves together with the combination of quark distributions $\Sigma = \sum_{i=1}^6 (q_i + \bar{q}_i)$ as the so-called singlet contribution. The non-singlet part includes contributions that in EKO follow the common notation of V_i and T_i , while “mixed” refers to the contribution where both a singlet and non-singlet part enter the convolution. The dependence on μ_F for a decomposition in this basis is shown in Fig. 4b. Indeed, in this basis two points are immediately apparent: first, the cancellation is much less pronounced even at high values of μ_F by approximately an order of magnitude and; second, the dependence on the factorisation scale is rendered almost flat. This demonstrates that much of the cancellations occurs *within* the singlet sector and that the large cancellations observed in the flavour basis are largely an artefact of the choice of basis and the correlations introduced by the DGLAP evolution.

3.2 Stability of K -factors

The accidental cancellation between partonic channels highlighted in the previous section could, in principle, be either amplified or diminished by the choice of PDF, as they control the relative size of each partonic channel. In turn one might expect the NNLO contribution to the cross section to be significantly impacted by the choice of PDF.

Figure 5a provides a detailed comparison of NNLO predictions for a variety of PDF sets, where the bands correspond to the respective PDF uncertainties. In the bottom panel, predictions are also shown for approximate N³LO sets in order to assess the impact from the (approximate) higher-order DGLAP evolution implemented in these sets. The latter is particularly important as the previous section exposed the correlation induced by the evolution to be a main driver of the cancellations observed at NNLO. The spread between the different PDFs is at the few percent level and largely covered by the respective uncertainties. The largest deviations are seen between the NNPDF and ABM sets of about 4%, which can be attributed to the different datasets included in the fits and

⁵ The exact definition of the evolution basis is provided in the [EKO documentation](#).

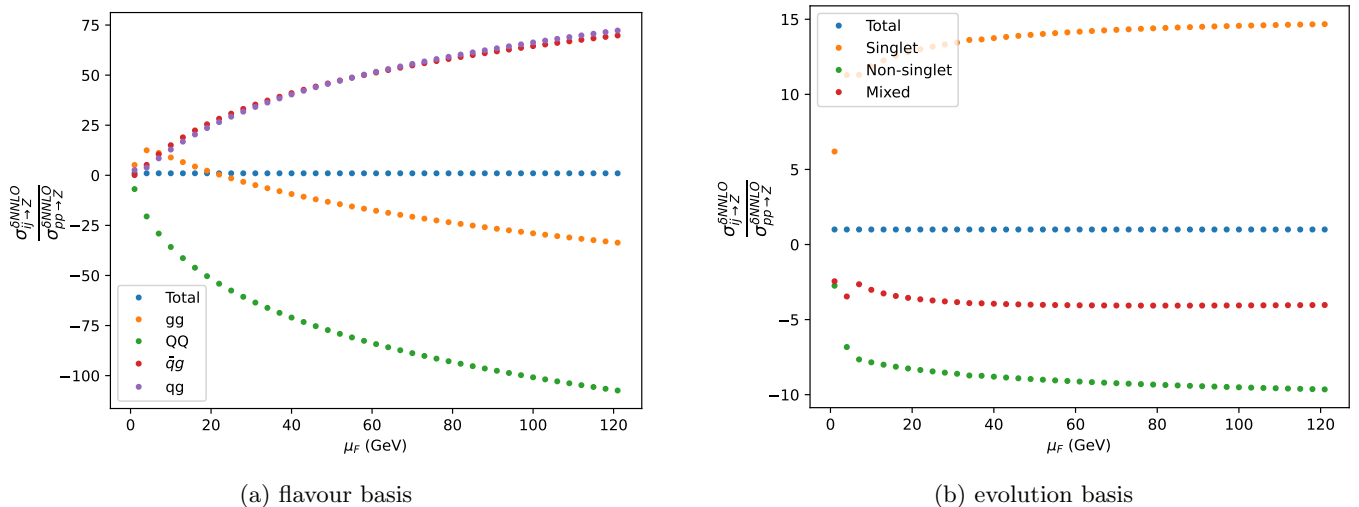


Figure 4: Channel decomposition of the NNLO contribution $\sigma^{\delta\text{NNLO}}$ to the total cross section for Z production as a function of the scale, normalised to the full NNLO contribution. Decomposition is shown both for the flavour basis (a) and the evolution basis (b). The “QQ” channel includes all quark and anti-quark initiated contributions.

the different methodologies employed rather than the details of this calculation. The impact of the approximate $N^3\text{LO}$ evolution is found to be small and well within the PDF uncertainties.

Without the availability of NNLO grids, approximations were constructed based on NLO grids supplemented by NNLO K -factors that were obtained for a specific choice of PDF,

$$\sigma^{\text{NNLO}} \approx \sigma^{\text{NLO}} \times K_{\text{PDF}}^{\text{NNLO}}, \quad K_{\text{PDF}}^{\text{NNLO}} \equiv \frac{\sigma_{\text{PDF}}^{\text{NNLO}}}{\sigma_{\text{PDF}}^{\text{NLO}}}.$$

For this approximation to be valid, the relevant property to inspect is not the total cross section but the K -factors and how stable they are with respect to the choice of the underlying PDF set used. In Fig. 5b we show the K -factors for the same PDF sets as in Fig. 5a including in the top panel a comparison of the (correlated) PDF uncertainties against the size of the NNLO scale variation. With the K -factors typically at the percent level, we observe that variation of the K -factors under the change of the PDF sets is very stable and only at the level of a few per-mille. The final impact on the full NNLO cross section σ^{NNLO} from the K -factor approximation is thus estimated to be at the order of 0.01‰ and thus negligible.

3.3 K -factor approximation and PDF fits

Given the current claimed level of accuracy of PDF global analyses ($N^3\text{LO}$, albeit approximated) [65, 66] and the uncertainty achieved in regions well constrained by data (close to 1%) [16] it is critical to study and remove any possible sources of bias in order to ensure the accuracy and robustness of the PDFs and its associated uncertainties. While the previous section supports the robustness of the K -factor approximation for the DY process at the

level of the full cross section, it is important to also verify the impact of this approximation on global PDF fits. This is because PDF fits are sensitive to corrections for the *individual* channels that are not correctly captured by a global K -factor nor are partonic channels that only open up at NNLO (such as gg and qq').

Moreover, PDF analyses are among the main use-cases of interpolation grids: Determining PDFs relies heavily on comparing data and predictions for many observables and varying inputs. This process involves evaluating hundreds of differential cross-sections at the PDF fitting scale for each step in the optimisation procedure. A global analysis of PDFs requires computing predictions for approximately 5000 datapoints to NNLO [16]. Given the complexity and the computational cost, handling this vast amount of data without the speed-up that the grids provide is practically infeasible.

Since grids include detailed information about scales and orders, they can be convolved with evolution operators, to generate optimized Fast-Kernel tables. These tables optimize away information about the orders and scales such that they can be directly convolved with PDFs at the fitting scale [11]. This makes them particularly well suited for PDF determination. We have used the tools outlined in Ref. [13] to prepare Fast-Kernel tables for all grids provided together with this paper.

In order to examine the impact of the K -factor approximation on state-of-the-art PDF analyses, we have utilised the open-source NNPDF fitting framework [67] to perform a series of fits under different assumptions: DY theory predictions are varied between the exact NNLO grids and the K -factor approximation, and the data entering the fits either include the full NNPDF4.0 dataset or a reduced dataset based solely on collider DY measurements (“DY only fit”). The latter fit based on the restricted data is intended to act as the worst-case scenario where *all* datasets

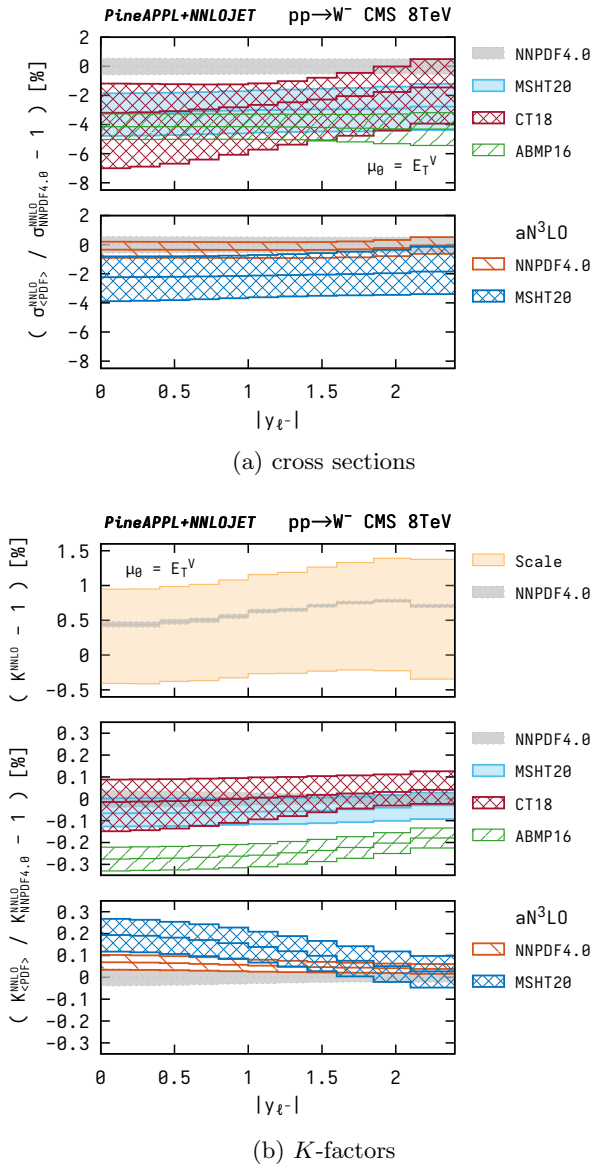


Figure 5: Comparison of different PDF sets for NNLO cross sections (a) and K -factors (b) in units of % with respect to the NNPDF4.0 set.

are impacted by the approximation in the theory predictions. To closely mimic the procedure of an actual PDF fit, we have computed the K -factor with the closest PDF, in this case NNPDF4.0 [16].

Note that for the results presented in this section, aside from the grids, all other conditions and settings are kept identical between the two fits, including initialisation and seeding of random numbers.

In Fig. 6 we compare the effect of the K -factor approximation in the context of a global fit by explicitly showing the gluon distribution at the fitting scale. We consider two situations, a fit based exclusively on DY data from LHC and Tevatron (top) and a fit using a global dataset (bottom). In the first case the impact of the K -factor approximation is visible even in the data region, with a small shift

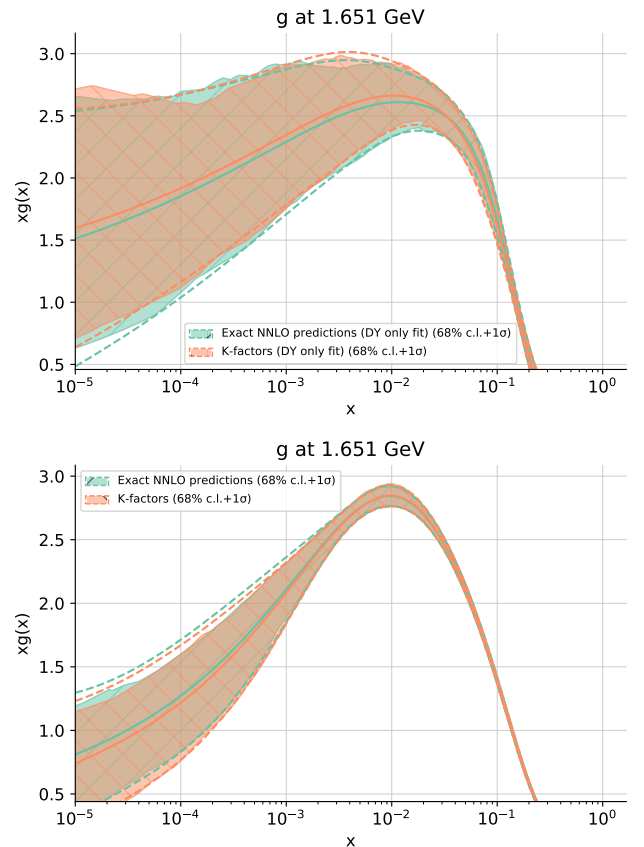


Figure 6: Comparison of the gluon PDF between two PDF fits, one with exact NNLO grids and the other based on the K -factor approximation. The fit is repeated using only DY data (top) and the global dataset (bottom).

in the data region which however is below a percent and safely well within the uncertainties of the determination. In the global fit instead we consider the NNPDF4.0 settings, i.e., all DIS data is exact up to NNLO and double-hadronic data instead uses the K -factor approximation. The better constraints introduced by the global dataset make it so that when using exact grids for the DY data, the effect is completely negligible and only visible in the extrapolation region, where instabilities introduced by the small numerical differences between the K -factor and exact approximation won't be compensated by the fit. While we only consider the gluon in Fig. 6, the same is true for other partons.

In Fig. 7 we display a comparison of the qq luminosity in the same scenarios as Fig. 6. We observe how the impact of the K -factor approximation affects the entire range of the phase space in a DY only fit, while the inclusion of the rest of the datasets commonly considered in a global fit eliminates the impact (note the different range on the y -axis).

From the comparison of the different fits presented in this section, we can conclude that the K -factor approximation is a valid and safe approximation for the DY process in the context of PDF fits. Even in the specially




Figure 7: Comparison of the qq luminosity between two PDF fits, one with exact NNLO grids and the other based on the K -factor approximation. The fit is repeated using only DY data (top) and the global dataset (bottom).

tuned scenario where the impact of the approximation was maximised by restricting the fit data only to hadron-collider Drell–Yan measurements, the impact on the resulting PDFs was found to be minimal and well within uncertainties.

4 Conclusions and outlook

In this note we have reported on a new interface between the PINEAPPL interpolation grid library and the NNLO parton-level Monte Carlo generator NNLOJET. This interface was used to produce interpolation grids for a wide range of DY measurements performed at Tevatron and the LHC that commonly enter global PDF analyses.

These grids were used to gain further insights into the accidental cancellations that occur in the Drell–Yan process at NNLO, where strong correlations induced by the DGLAP evolution within the singlet sector were identified as a main driver of the cancellations. We further performed a detailed study of the K -factor approximation that was employed in the PDF fits so far. While few-% variations are found from the variation of PDF sets in the NNLO

cross section, the K -factor is found to be very stable with only changes at the per-mille level. A set of PDF fits based on varying theory predictions as well as fit data further support the conclusion that the impact of the K -factor approximation is minimal and well below the quoted PDF uncertainties.

The independence of PDF fits on the K -factor approximation for the DY data is a non-trivial consequence of the observations detailed in sections 3.1 and 3.2. While there are big channel-by-channel cancellations, these are dominated by DGLAP evolution. While PDF determinations might differ widely in data and methodology, at a fixed order the evolution is fixed, and so the pattern of cancellations is preserved regardless of the PDF in use.

The grids provided in this work not only allowed to establish the validity of the K -factor approximation for DY at NNLO but also paves the way towards incorporating N^3 LO predictions into PDF fits in the future. Full N^3 LO grids are still beyond reach due to computation costs, however, the NNLO grids provide one ingredient to construct approximate N^3 LO predictions based on N^3 LO K -factors.

Acknowledgements

C.S. is supported by the German Research Foundation (DFG) under reference number DE 623/6-2.

References

1. T. Carli, D. Clements, A. Cooper-Sarkar, C. Gwenlan, G. P. Salam, F. Siegert, P. Starovoitov and M. Sutton, *Eur. Phys. J. C* 66 (2010) 503–524, [arXiv:0911.2985](#).
2. T. Kluge, K. Rabbertz and M. Wobisch, “FastNLO: Fast pQCD calculations for PDF fits,” in *14th International Workshop on Deep Inelastic Scattering*, pp. 483–486. 9 2006. [arXiv:hep-ph/0609285](#).
3. fastNLO Collaboration, D. Britzger, K. Rabbertz, F. Stober and M. Wobisch, “New features in version 2 of the fastNLO project,” in *20th International Workshop on Deep-Inelastic Scattering and Related Subjects*, pp. 217–221. 2012. [arXiv:1208.3641](#).
4. S. Carrazza, E. R. Nocera, C. Schwan and M. Zaro, *JHEP* 12 (2020) 108, [arXiv:2008.12789](#).
5. C. Schwan, *SciPost Phys. Proc.* 8 (2022) 079, [arXiv:2108.05816](#).
6. D. Britzger *et al.*, *Eur. Phys. J. C* 79 (2019) 845, [arXiv:1906.05303](#). [Erratum: *Eur.Phys.J.C* 81, 957 (2021)].
7. H1, ZEUS Collaboration, I. Abt *et al.*, *Eur. Phys. J. C* 82 (2022) 243, [arXiv:2112.01120](#).
8. D. Britzger *et al.*, *Eur. Phys. J. C* 82 (2022) 930, [arXiv:2207.13735](#).
9. M. Czakon, D. Heymes and A. Mitov, [arXiv:1704.08551](#).
10. M. V. Garzelli, J. Mazzitelli, S. O. Moch and O. Zenaiev, *JHEP* 05 (2024) 321, [arXiv:2311.05509](#).
11. V. Bertone, S. Carrazza and N. P. Hartland, *Comput. Phys. Commun.* 212 (2017) 205–209, [arXiv:1605.02070](#).
12. A. Candido, F. Hekhorn, G. Magni, T. R. Rabemananjara and R. Stegeman, *Eur. Phys. J. C* 84 (2024) 698, [arXiv:2401.15187](#).

13. A. Barontini, A. Candido, J. M. Cruz-Martinez, F. Hekhorn and C. Schwan, *Comput. Phys. Commun.* **297** (2024) 109061, [arXiv:2302.12124](#).
14. T.-J. Hou *et al.*, *Phys. Rev. D* **103** (2021) 014013, [arXiv:1912.10053](#).
15. S. Bailey, T. Cridge, L. A. Harland-Lang, A. D. Martin and R. S. Thorne, *Eur. Phys. J. C* **81** (2021) 341, [arXiv:2012.04684](#).
16. NNPDF Collaboration, R. D. Ball *et al.*, *Eur. Phys. J. C* **82** (2022) 428, [arXiv:2109.02653](#).
17. S. D. Drell and T.-M. Yan, *Phys. Rev. Lett.* **25** (1970) 316–320. [Erratum: *Phys.Rev.Lett.* **25**, 902 (1970)].
18. C. Anastasiou, L. J. Dixon, K. Melnikov and F. Petriello, *Phys. Rev. D* **69** (2004) 094008, [arXiv:hep-ph/0312266](#).
19. K. Melnikov and F. Petriello, *Phys. Rev. D* **74** (2006) 114017, [arXiv:hep-ph/0609070](#).
20. S. Catani, L. Cieri, G. Ferrera, D. de Florian and M. Grazzini, *Phys. Rev. Lett.* **103** (2009) 082001, [arXiv:0903.2120](#).
21. S. Catani, G. Ferrera and M. Grazzini, *JHEP* **05** (2010) 006, [arXiv:1002.3115](#).
22. C. Duhr, F. Dulat and B. Mistlberger, *Phys. Rev. Lett.* **125** (2020) 172001, [arXiv:2001.07717](#).
23. C. Duhr, F. Dulat and B. Mistlberger, *JHEP* **11** (2020) 143, [arXiv:2007.13313](#).
24. C. Duhr and B. Mistlberger, *JHEP* **03** (2022) 116, [arXiv:2111.10379](#).
25. J. Baglio, C. Duhr, B. Mistlberger and R. Szafron, *JHEP* **12** (2022) 066, [arXiv:2209.06138](#).
26. X. Chen, T. Gehrmann, N. Glover, A. Huss, T.-Z. Yang and H. X. Zhu, *Phys. Rev. Lett.* **128** (2022) 052001, [arXiv:2107.09085](#).
27. X. Chen, T. Gehrmann, E. W. N. Glover, A. Huss, P. F. Monni, E. Re, L. Rottoli and P. Torrielli, *Phys. Rev. Lett.* **128** (2022) 252001, [arXiv:2203.01565](#).
28. X. Chen, T. Gehrmann, N. Glover, A. Huss, T.-Z. Yang and H. X. Zhu, *Phys. Lett. B* **840** (2023) 137876, [arXiv:2205.11426](#).
29. T. Neumann and J. Campbell, *Phys. Rev. D* **107** (2023) L011506, [arXiv:2207.07056](#).
30. J. Campbell and T. Neumann, *JHEP* **11** (2023) 127, [arXiv:2308.15382](#).
31. NNLOJET Collaboration, A. Huss *et al.*, “NNLOJET: a parton-level event generator for jet cross sections at NNLO QCD accuracy.” 2025.
32. A. Daleo, T. Gehrmann and D. Maitre, *JHEP* **04** (2007) 016, [arXiv:hep-ph/0612257](#).
33. R. Boughezal, A. Gehrmann-De Ridder and M. Ritzmann, *JHEP* **02** (2011) 098, [arXiv:1011.6631](#).
34. T. Gehrmann and P. F. Monni, *JHEP* **12** (2011) 049, [arXiv:1107.4037](#).
35. A. Gehrmann-De Ridder, T. Gehrmann and M. Ritzmann, *JHEP* **10** (2012) 047, [arXiv:1207.5779](#).
36. A. Gehrmann-De Ridder, T. Gehrmann and E. W. N. Glover, *JHEP* **09** (2005) 056, [arXiv:hep-ph/0505111](#).
37. CDF Collaboration, T. A. Aaltonen *et al.*, *Phys. Lett. B* **692** (2010) 232–239, [arXiv:0908.3914](#).
38. D0 Collaboration, V. M. Abazov *et al.*, *Phys. Rev. D* **76** (2007) 012003, [arXiv:hep-ex/0702025](#).
39. D0 Collaboration, V. M. Abazov *et al.*, *Phys. Rev. D* **88** (2013) 091102, [arXiv:1309.2591](#).
40. D0 Collaboration, V. M. Abazov *et al.*, *Phys. Rev. D* **91** (2015) 032007, [arXiv:1412.2862](#). [Erratum: *Phys.Rev.D* **91**, 079901 (2015)].
41. ATLAS Collaboration, G. Aad *et al.*, *JHEP* **06** (2014) 112, [arXiv:1404.1212](#).
42. ATLAS Collaboration, G. Aad *et al.*, *Phys. Lett. B* **725** (2013) 223–242, [arXiv:1305.4192](#).
43. ATLAS Collaboration, G. Aad *et al.*, *Phys. Rev. D* **85** (2012) 072004, [arXiv:1109.5141](#).
44. ATLAS Collaboration, M. Aaboud *et al.*, *Eur. Phys. J. C* **77** (2017) 367, [arXiv:1612.03016](#).
45. ATLAS Collaboration, G. Aad *et al.*, *Eur. Phys. J. C* **79** (2019) 760, [arXiv:1904.05631](#).
46. ATLAS Collaboration, M. Aaboud *et al.*, *JHEP* **12** (2017) 059, [arXiv:1710.05167](#).
47. ATLAS Collaboration, G. Aad *et al.*, *JHEP* **08** (2016) 009, [arXiv:1606.01736](#).
48. CMS Collaboration, S. Chatrchyan *et al.*, *Phys. Rev. Lett.* **109** (2012) 111806, [arXiv:1206.2598](#).
49. CMS Collaboration, S. Chatrchyan *et al.*, *Phys. Rev. D* **90** (2014) 032004, [arXiv:1312.6283](#).
50. CMS Collaboration, S. Chatrchyan *et al.*, *JHEP* **12** (2013) 030, [arXiv:1310.7291](#).
51. CMS Collaboration, V. Khachatryan *et al.*, *Eur. Phys. J. C* **76** (2016) 469, [arXiv:1603.01803](#).
52. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **02** (2013) 106, [arXiv:1212.4620](#).
53. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **08** (2015) 039, [arXiv:1505.07024](#).
54. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **05** (2015) 109, [arXiv:1503.00963](#).
55. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **01** (2016) 155, [arXiv:1511.08039](#).
56. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **10** (2016) 030, [arXiv:1608.01484](#).
57. ATLAS Collaboration, G. Aad *et al.*, *Phys. Lett. B* **759** (2016) 601–621, [arXiv:1603.09222](#).
58. LHCb Collaboration, R. Aaij *et al.*, *JHEP* **09** (2016) 136, [arXiv:1607.06495](#).
59. A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, C. T. Preuss and D. M. Walker, *JHEP* **05** (2023) 002, [arXiv:2301.11827](#).
60. Z. Bern, L. J. Dixon, F. Febres Cordero, S. Höche, H. Ita, D. A. Kosower and D. Maitre, *Comput. Phys. Commun.* **185** (2014) 1443–1460, [arXiv:1310.7439](#).
61. D. Maître, *J. Phys. Conf. Ser.* **1525** (2020) 012014.
62. M. Czakon, Z. Kassabov, A. Mitov, R. Poncelet and A. Popescu, *J. Phys. G* **51** (2024) 115002, [arXiv:2304.05993](#).
63. A. Candido, F. Hekhorn and G. Magni, “N3pdf/eko: Paper,” Mar. 2022. <https://doi.org/10.5281/zenodo.6340153>.
64. A. Candido, F. Hekhorn and G. Magni, *Eur. Phys. J. C* **82** (2022) 976, [arXiv:2202.02338](#).
65. J. McGowan, T. Cridge, L. A. Harland-Lang and R. S. Thorne, *Eur. Phys. J. C* **83** (2023) 185, [arXiv:2207.04739](#). [Erratum: *Eur.Phys.J.C* **83**, 302 (2023)].
66. NNPDF Collaboration, R. D. Ball *et al.*, *Eur. Phys. J. C* **84** (2024) 659, [arXiv:2402.18635](#).
67. NNPDF Collaboration, R. D. Ball *et al.*, *Eur. Phys. J. C* **81** (2021) 958, [arXiv:2109.02671](#).