

Precision phenomenology with MCFM

John Campbell^a and Tobias Neumann^{a,b}

^a*Fermilab,*

PO Box 500, Batavia, Illinois 60510, U.S.A.

^b*Department of Physics, Illinois Institute of Technology,
Chicago, Illinois 60616, U.S.A.*

E-mail: johnmc@fnal.gov, tneumann@fnal.gov

ABSTRACT: Without proper control of numerical and methodological errors in theoretical predictions at the per mille level it is not possible to study the effect of input parameters in current hadron-collider measurements at the required precision. We present the new version of the parton-level code MCFM-9.0 that achieves this requirement through its highly-parallelized nature, significant performance improvements and new features. An automatic differential jetiness slicing cutoff extrapolation is introduced to assess the cutoff dependence of all results, thus ensuring their reliability and potentially improving fixed-cutoff results by an order of magnitude. The efficient differential study of PDF uncertainties and PDF set differences at NNLO, for multiple PDF sets simultaneously, is achieved by exploiting correlations. We use these improvements to study uncertainties and PDF sensitivity at NNLO, using 371 PDF set members. The work described here permits NNLO studies that were previously prohibitively expensive, and lays the groundwork necessary for a future implementation of NNLO calculations with a jet at Born level in MCFM.

KEYWORDS: QCD Phenomenology

ARXIV EPRINT: [1909.09117](https://arxiv.org/abs/1909.09117)

Contents

1	Introduction	1
2	New and improved features in MCFM	3
2.1	User interface	4
2.2	Correlated calculations with multiple scales, PDF sets and τ_{cut} values	5
3	Integration performance and uncertainty estimation	8
3.1	Low discrepancy sequence integration	9
3.2	Parallelization and integration uncertainty estimation	11
4	Updated NNLO benchmarks and τ_{cut}-dependence	17
4.1	Inclusive τ_{cut} benchmarks	18
4.2	Differential τ_{cut} extrapolation and assessment of τ_{cut} uncertainty	21
4.3	Examples of automatic differential τ_{cut} fitting	24
5	Precision PDF uncertainties and PDF set differences	31
5.1	Rapidities in single- and diboson production	34
5.2	Higgs boson transverse momentum	37
6	Precision studies of W and Z production	46
7	Conclusions	51
A	Detailed description of new features	54
A.1	Compatibility with the Intel compiler and benchmarks	57
A.2	Remarks on memory bound performance issues	58
B	Supporting plots for the jackknife-after-bootstrap procedure	59

1 Introduction

The production of bosons, either singly or in pairs, provides the bread and butter for LHC analyses that perform precision tests of the Standard Model (SM). Consequently, they also serve as probes of physics beyond the Standard Model (BSM), as well as arenas in which to perform resilient extractions of fundamental parameters of the SM and non-perturbative inputs such as parton density functions (PDFs). For this to be the case it is essential that the fixed-order calculations, upon which these analyses rely, are pushed to as high order in the perturbative expansion as possible. After many years of effort, QCD corrections have reached the level of NNLO and electroweak effects are commonly available

at NLO. Broadly speaking, these push the perturbative truncation uncertainty of these fixed order predictions to the percent level. At the same time, uncertainties related to the determination of PDFs and α_s have been reduced to the same level. To make further progress it is therefore imperative to have robust tools that can systematically compute predictions at this level of precision, whilst maintaining sensitivity to possible differences between theoretical inputs such as the choice of renormalization scale or PDF set.

Although by now many calculations have been performed at NNLO in QCD, far fewer have resulted in public codes.¹ Most of these are restricted to a single process, although some codes do offer additional specialized features such as the inclusion of electroweak corrections or resummation. A notable exception is MATRIX [17], which in its current release features fixed order NNLO implementations of single boson processes and a subset of diboson processes. Related, but more complicated, final states that also include the presence of an additional hard jet have also been computed at NNLO but are only available as private codes, see for example ref. [18] and references therein.

MCFM is a publicly available code, with version 1.0 released in 2001 focusing on NLO corrections to vector boson pair production processes [19]. Since then the code has been continuously maintained with updated and new processes at NLO and beyond. In 2015 multi-threading capability using OpenMP was added [20], enabling multi-core desktop systems to compute the most complicated NLO processes. In 2016 an initial set of color singlet NNLO processes was included [21] together with MPI capability, allowing full use of cluster systems. MCFM is now capable of computing W^\pm , Z , H as well as $\gamma\gamma$, $W^\pm H$, ZH and $Z\gamma$ production processes at NNLO, as well as hundreds of processes at NLO. Some processes include NLO electroweak corrections [22], while others account for contributions from BSM sources and anomalous couplings, as well as NLO corrections for the SMEFT [23]. All leptonic decay channels of Z and W^\pm are included as well as the Higgs boson decays into $\gamma\gamma$, W^+W^- , ZZ , $Z\gamma$, $\tau^+\tau^-$ and $\bar{b}b$.

With its flexibility, ease of use and performance, MCFM has been an indispensable ingredient in hundreds of experimental studies for the comparison of theory predictions and nature. However, the utility of the code is not confined to just the hands of experimentalists. One of the goals of MCFM was always to provide a collection of analytic results to serve as a platform for further work, where others can extend, modify, or reuse parts without significant help from the authors. In the past, elements of the MCFM code have been used extensively for this purpose. For example, codes such as DYNNLO [1], DYRes [24, 25], HNNLO [2–4], HRES [4, 26] as well as most recently MCFM-RE [27] are all based on the MCFM framework. Other public codes benefit from various parts of MCFM and its efficient implementation of amplitudes.²

This paper describes an update of the parton level code MCFM that includes significant improvements in its usability, reliability, maintainability and performance. For example, the integration has been rewritten to adaptively reach a specified precision goal.

¹Fully differential codes for various boson and diboson processes are for example DYNNLO [1], HNNLO [2–4], SusHi [5, 6], ggHiggs [7], FehiPro [8, 9], proVBFH(H) [10, 11], 2 γ NNLO [12], FEWZ [13–15] and GENEVA [16].

²See e.g. refs. [28–31] for its use in other codes, and <https://inspirehep.net/search?ln=en&p=find+fulltext+MCFM> for many more examples.

To establish trust in our results at the per mille level, we perform a detailed study of the performance of the Vegas [32] integration algorithm in our code. In particular, we focus on the convergence of the integral and the reliability of its error estimates.

With improvements in the jettiness slicing cutoff (τ_{cut}) dependence we improve and update previous benchmark results. We furthermore introduce the automatic sampling of additional τ_{cut} values fully differentially, making use of the correlations to decrease numerical uncertainties by orders of magnitude. The τ_{cut} dependence and its automatic fitting to the known asymptotic behavior can be used to reliably assess systematic τ_{cut} errors.

For the first time, multiple PDF sets can be used at the same time for the evaluation of PDF uncertainties in the same correlated way. We demonstrate the calculation of PDF set differences and uncertainties at NLO and NNLO for differential distributions with sub per mille level numerical accuracies. We then compare PDF uncertainties obtained using lower order matrix elements for a broad range of processes to understand the level of precision that may be expected when estimating PDF uncertainties at NNLO through this procedure. Making use of correlations, we can furthermore study the *differences* between any number of PDF sets in our improved setup at the sub per mille level.

In section 2 we introduce the new and improved features of MCFM-9.0, and support them with technical data, details and benchmarks. In section 3 we study the performance of the Vegas integration routine, comparing the use of a newly introduced low-discrepancy sequence with that of a pseudo-random number sequence. We focus on issues regarding the estimation of numerical integration uncertainties, comparing against approaches in the literature and suggesting improvements. In section 4 we report on the performance gains resulting from using the boosted definition of the jettiness slicing variable and the inclusion of power corrections differentially at NNLO. We also present our fully differential automatic τ_{cut} extrapolation based on the theoretical asymptotic τ_{cut} dependence and discuss its use and limitations. In section 5 we study PDF uncertainties at NNLO using six PDF sets simultaneously for all NNLO processes in MCFM and compare it against the use of lower order matrix elements. The correlated multi-PDF-set integration allows for per mille comparisons between different PDF sets, which are also covered for the Higgs transverse momentum distribution at large values. In section 6 we comment on issues in precision studies in W and Z -boson physics, where current experimental data now requires per mille level theory predictions. We compare benchmark results in the literature with our predictions. Our conclusions are summarized in section 7.

2 New and improved features in MCFM

A number of features have been introduced to simplify the operation of the MCFM code. In this section we summarize the most important new and modified features. In passing, we note that this version represents an overhaul of many key components of the code and, for the sake of clarity, we have removed a number of features introduced in previous versions that had been largely unused.³ This version aims to be compliant with Fortran 2008 and

³These include the ntuple interface, as well as the option to write output to LHEF files. These can easily be added back at a later time, according to user demand. In the meantime older versions of MCFM can still be downloaded to make use of these features.

fully supports GCC versions newer than 7 and Intel compilers newer than 19. Benchmark comparisons of compiler versions and optimization flags are described in appendix A. A technical description of the new features presented in this section and their configuration within MCFM is also given in appendix A.

MCFM is distributed in a fully self-contained form, where all amplitudes are bundled as optimized explicit expressions, making use of the included QCDDLoop [33, 34] distribution for the evaluation of one-loop scalar integrals. Expressions for multi-loop integrals in terms of harmonic polylogarithms are evaluated with TDHPL [35]. While a number of PDF sets are included through a native interface, the LHAPDF library [36] can also be linked. This enables using PDF uncertainties and a larger number of PDF sets. Through LHAPDF one also gains the flexibility of easily interchanging grid interpolation routines. While we do not enforce citations for using MCFM or results obtained with MCFM (MCFM is now released under the GPL-3.0 license), we encourage citing at least the papers printed at the beginning of a run. In practice, the results in MCFM depend on a deep tree of results (for instance, for the calculation of the matrix elements) and we hope that users use the appropriate citation depth to acknowledge research appropriately.

Already with MCFM-8.0 inclusive cross sections can be computed precisely at NNLO on a modern multi-core desktop computer in a few hours. Achieving sub percent level precision also in tails of distributions requires more computing resources. Furthermore, adding features like automatic scale variation or automatic computation of PDF uncertainties increases required computational resources. Apart from buying more computers, improvements in various parts of the theory predictions and the code, both which are covered in this study, can be made to make these computations feasible on smaller sized clusters.

2.1 User interface

For this version we have first introduced a new, more flexible input file format. The options from the input file can now be over-ridden via command line arguments as well, which can be useful for batch parameter run scripts. Second, we have re-implemented the Vegas algorithm [37] and the surrounding integration routines, including a new alternative to the pseudo-random numbers used in previous versions of code. By default, we now use the Sobol low discrepancy sequence [38–42] that is described in detail in section 3.1.

With the new integration routines all parts of a NLO or NNLO calculation are now chosen adaptively based on the largest absolute numerical uncertainty. A precision goal can be set in the input file as well as a χ^2 /iteration goal and a precision goal for the warmup run. If the goals for the warmup are not reached, the warmup repeats with twice the number of calls. Our new version also allows the integration to be resumed from any point from a previous run, using snapshots that save the whole integration state automatically. This allows the precision goal to be reduced in a future run without starting the whole integration from scratch. Further configuration options have been introduced to control the stages of the integration that can provide benefits over the default settings in certain situations, such as when calculating PDF uncertainties, as described in detail in appendix A.

MCFM allows to easily specify the most common kinematical cuts in the input file and automatically fills a pre-defined set of histograms. The cuts from the input file can be modified and augmented in a prepared subroutine for user cuts. Additionally, a re-weighting function has been introduced that multiplies all events in the integration. This feature can be used as a manual importance sampling technique to give tails of distributions a larger weight so they are integrated with the same relative precision as numerically larger contributions. For example in transverse momentum distributions one can reweight with an exponential function of the transverse momentum to approximately flatten the distribution and obtain an equal relative precision in all bins. Last, histograms of arbitrary kinematical variables can easily be added for full flexibility.

We have implemented a new suite of histogram routines that allows for any number of histograms with any number of bins, each of which is dynamically allocated. Furthermore, everything is also handled in a fully multi-threaded approach within the integration. For each OMP thread temporary histograms are allocated and filled that are then reduced to a single one after each integration iteration. These histograms are also written out at every intermediate stage of the integration, and are updated appropriately when an integration run is resumed with a new precision goal. This allows the user to inspect the results already during the integration and gives the possibility to interactively stop the integration as soon as the results are satisfactory.

2.2 Correlated calculations with multiple scales, PDF sets and τ_{cut} values

When integrating multiple functions at once one can make use of the correlations between the integrands, and straight away obtain significantly lower integration uncertainties on the differences and ratios between the different integrands. Generally one expects the difference between two perfectly correlated integrands to be computed with the same relative uncertainty as that of either integral. In practice, we find that the absolute numerical uncertainty computed for the difference between an integral and the central value turns out to be roughly an order of magnitude lower than the numerical uncertainty on the central value itself. This idea had already been employed in previous versions of MCFM for the calculation of renormalization and factorization scale uncertainties. In this version we extend the treatment to the calculation of NLO and NNLO predictions with different values of the 0-jettiness cutoff τ_{cut} , and to calculations with different PDF sets and members.

PDF uncertainties and PDF set differences. Apart from FEWZ [13], no (public) code is known to the authors that computes PDF uncertainties automatically while taking into account correlations between PDF set members. Indeed, studies commonly avoid the expensive calculation of fixed order NNLO results convoluted with different PDF sets for central values and PDF uncertainties, for example [43, 44]. Instead, relative PDF uncertainties or differences between sets are calculated using fixed order NLO matrix elements convoluted with NNLO PDFs. Other alternatives include frameworks like fastNLO [45] and applGRID [28], which were developed to accelerate PDF fits and to reduce the burden of computing PDF uncertainties.

In MCFM-8.1 PDF uncertainties could successfully be computed inclusively at NLO. In MCFM-9.0 we enable studying calculations of PDF uncertainties differentially at NNLO on smaller sized clusters with just a few hundred cores in total.⁴ This is achieved through a fully parallelized OMP+MPI interface to LHAPDF that is based on the new object oriented treatment of PDFs in LHAPDF 6. This interface thereby avoids the limitation to have at most one simultaneous call to the library from the OMP threads. We also added the capability to handle any number of PDF sets, with or without PDF uncertainties, limited only by the available system memory. Studying precise differences between PDF sets, at the sub per mille level, can then be performed with only little computational overhead compared to using just one central value.

Note that to enable PDF uncertainties, MCFM has to be compiled with LHAPDF [36] support. Any sets available for LHAPDF can be used, and the uncertainties are estimated according to the provided LHAPDF uncertainty routines for replica and Hessian sets. In addition, as a further improvement to previous versions, we also allow for sets with additional members that use different values of $\alpha_s(m_Z)$, so that combined PDF + α_s uncertainties can be computed at the same time.

To demonstrate this procedure at work, in figure 1 we show normalized e^+ rapidity distributions in W^+ production. This calculation has been performed at NNLO, using six NNLO PDF sets simultaneously and with PDF uncertainties for all of them, so calls to 371 PDF members for each phase space point. The normalization is with respect to the central value of the PDF4LHC [46] set PDF4LHC15_nnlo_30 to show the differences between this set and the other sets. Cuts and parameters are standard as in MCFM-8.0 and 9.0, introduced later in section 5. Only the positron rapidity y has been constrained to $2.0 < y < 4.5$. It is clear that this feature allows precision studies of differences between predictions of various PDF sets and their uncertainty bands.

We hope that our fast multi-set implementation in MCFM, that can inexpensively study 0.1% differences between PDF sets and their associated uncertainties, will help facilitate further theoretical work in these directions. For instance, differences such as those shown in figure 1 and section 5 can be explored and criticism of the PDF4LHC paradigm [47] can be supported or refuted directly at the level of NNLO kinematical predictions. Furthermore, recent work that incorporates theory uncertainties more rigorously in PDF fits [48] can again be studied directly and efficiently with MCFM-9.0.

Multi- τ_{cut} integration and automatic asymptotic fitting. For 0-jettiness calculations, either at NLO or NNLO, we now allow an array of τ_{cut} values to be specified in addition to the nominal value of τ_{cut} . The Vegas integration grid is still adjusted according to the nominal value, and the remaining values are sampled on the fly with little computational overhead. This means that values smaller than the nominal value of τ_{cut} are only computed with relatively large uncertainties, although this may still yield useful information. In contrast, any values larger than the nominal one are computed with approximately the same

⁴The per-mille level differential computations with uncertainties from six PDF sets simultaneously in this paper were performed on a cluster using at most 16 nodes, where the nodes have AMD and Intel CPUs from 2010, for details see appendix A.




Figure 1. NNLO e^+ rapidity distributions for W^+ production in the forward region, computed with uncertainties from a variety of PDF sets, normalized to the central PDF4LHC prediction.

precision as the nominal value, and are therefore highly reliable. If no values of additional values of τ_{cut} are specified, the code automatically chooses further values, see section 4.

In all cases an automatic fit to the known asymptotic behavior is performed for the total cross section as well as for all histograms differentially. The histograms for the nominal and individual τ_{cut} values are written separately from the histogram with just the fitted corrections. With this procedure one can quickly and easily check the τ_{cut} dependence of the result and estimate the effect of a non-zero value of τ_{cut} . See section 4 for details on this procedure.

All τ_{cut} dependence plots in this paper are computed with the multi- τ_{cut} feature and its automatic fitting. For example figure 2 displays the hardest photon p_T spectrum in diphoton production at NNLO using the fully automatic differential fit for two runs with nominal τ_{cut} values of 10^{-3} GeV and 10^{-4} GeV. The latter choice is our default and is expected to result in systematic cutoff effects of less than one percent. In both cases the fit significantly improves the fixed τ_{cut} predictions, and for $\tau_{\text{cut}} = 10^{-4}$ GeV it allows one to estimate that indeed the residual dependence is at most 1%, while for the choice of $\tau_{\text{cut}} = 10^{-3}$ GeV the effects are up to 4.5%. The fitted results agree with each other at around the one half percent level (up to numerical noise), so allow one to choose a nominal τ_{cut} one order of magnitude larger. For practical applications the purpose of the additionally sampled τ_{cut} values and the fit is that no separate runs with different nominal τ_{cut} values have to be run. In order to make use of the fit its quality — separately reported by MCFM — must be inspected according to section 4.




Figure 2. The transverse momentum distribution of the hardest photon in diphoton production, computed at NNLO. Results are shown for $\tau_{\text{cut}} = 10^{-3}$ GeV (blue) and $\tau_{\text{cut}} = 10^{-4}$ GeV (red), as well as the results obtained using automatic fitting (darker blue and red). All results are normalized to the fit from $\tau_{\text{cut}} = 10^{-4}$ GeV.

3 Integration performance and uncertainty estimation

The numerical integration procedure in this version of MCFM has been overhauled and extended, with the aim of providing a platform with which to perform detailed high-precision studies of the most complicated NNLO processes. In this section we describe a number of the new features and benchmark their performance.

First we study our implementation of the Sobol low discrepancy sequence as an alternative to the MT19937 pseudo random number generator. Although Cuba [49], a popular integration library, has used the Sobol sequence as default for a long time already, to our knowledge no systematic study of its impact in Monte Carlo generators has been performed so far. We therefore describe a series of benchmarks quantifying the behavior of the Sobol low discrepancy sequence in MCFM and compare the performance of both sequences.

Second, we investigate the reliability of the result of the integration, including the uncertainty estimate. Modern complicated NLO and NNLO calculations are computationally intensive and typically require of the order of many CPU months. Since one does not want to wait months and years to run a NNLO calculation on a single core, parallelization is a straight-forward, easy and highly-efficient approach for Monte Carlo integrations. An outstanding feature of MCFM is its parallelization in both OMP and MPI on multi-core machines and for cluster setups. With this version the integration is also fully resumable. Complications arise when the Vegas integration is not parallelized, and one tries to combine many independent integrations with low statistics where uncertainties are typically underestimated. In order to prevent such problems, various approaches have been taken in the literature that go beyond a naïve combination. We study approaches performed in the

literature, suggest improvements, and furthermore compare these with the fully parallelized Vegas integration in MCFM that outright avoids such complications.

3.1 Low discrepancy sequence integration

If an integrand has a so-called bounded variation, implying certain smoothness properties, see for example the pioneering publications for the Koksma-Hlawka inequality [50, 51], the Monte Carlo integration error has a bound that is proportional to this bounded variation. The bounded variation is just a property of the integrand and the discrepancy of the sequence that defines the points being sampled. Since it is a fixed property, one can try to construct sequences that have a lower discrepancy than true or pseudo random numbers to improve the Monte Carlo integration. In this sense the asymptotic discrepancy of zero for equidistributed numbers is optimal. The problem for the rectangle rule with points $x_i = i/N$ for example, where N is the total number of calls, is that it has a fixed length and any increase of the number of calls to get an improved estimate leads to a recomputation that is not statistically independent. However, sequences have been constructed that keep the benefits of random numbers (statistically independent estimates) with the additional benefit of a lower discrepancy and additional uniformity properties in higher dimensions. Their discrepancies are generally bounded by $\sim (\log N)^d/N$, where N is the number of integrand evaluations and d the dimensionality. This is to be compared with true random numbers, where the asymptotic uncertainty in a Monte Carlo integration decreases only as $1/\sqrt{N}$.

In practice, it is unclear whether a bounded variation for our integrands exists or how large it is. Therefore, it is also unclear if one can benefit from using a low discrepancy sequence instead of pseudo random numbers. To test this, we use our implementation of the Sobol sequence that is based on the code `sobseq` [52], extended to 64 bits with a maximum sequence length of 2^{63} instead of $2^{31} \simeq 2.15 \cdot 10^9$. The extension to 64 bits is important because the 32-bit limit can quickly be saturated in NNLO calculations, and even in very precise NLO calculations. We use the initialization numbers from refs. [53–55].

To test the performance of the sequence, compared to the usual pseudo random numbers, we run benchmarks for the NNLO Higgs production double real emission calculation. We focus on the real emission since these contributions are usually the most computationally intensive ingredients in higher order calculations. The cuts are standard, and described later in section 5, but do not matter for the discussion here. The jettiness slicing cutoff τ_{cut} is set to 0.002 GeV, corresponding to a systematic precision goal of 0.2% for the full cross section with MCFM-8.0 and $\ll 0.1\%$ with MCFM-9.0.⁵ The dipole subtraction α parameter restricting the dipole phase space is set to 1.0 [56, 57], which corresponds to the original unrestricted formulation. To achieve a 0.2% precision goal also numerically the double real emission has to be computed with an uncertainty better than 4 (fb) or a relative precision of 10^{-4} .

Figure 3 shows the results of a comparison between the Sobol and one seeded MT19937 sequence, as a function of the accumulated number of Vegas calls. For this study we have multiplied the number of calls per iteration by 1.4 after every five iterations up to $500 \cdot 10^6$

⁵For the default dynamic choice of τ_{cut} as in section 4 this corresponds to $\tau_{\text{cut}}/m_{\text{Born}} = 1.6 \times 10^{-5}$.




Figure 3. The dependence of the integral for the NNLO Higgs production double real emission contribution (for $\alpha = 1$ and $\tau_{\text{cut}} = 0.002 \text{ GeV}$) on the accumulated number of calls. The red data corresponds to using the Sobol sequence and blue to a MT19937 sequence. Each point represents a new estimate from a new iteration.

calls per iteration, and from that point on we multiply by 1.4 after every single iteration.⁶ By using this factor, after two iterations (or two batches of iterations) the number of calls per iteration has been (approximately) doubled. This exponential increase allows successive Vegas iterations to make sufficient progress in reducing the uncertainty estimate. From the figure, the approach to the horizontal line — that represents the final average obtained using both sequences (31552 with an uncertainty of ± 2.5) — is equally good with both sequences. However, we note that for a smaller number of calls (not displayed in the figure) the unwritten rule to multiply Monte Carlo integration uncertainties by a factor of two is useful to see that the two are equivalent. We note that for the largest number of accumulated calls shown here, the number of calls per iteration reached $20 \cdot 10^9$ and one might benefit from increasing the number of Vegas grid subdivisions beyond their default value of 100.

As an alternative measure of performance, in figure 4 we directly show the dependence of the integration uncertainty on the number of accumulated calls. The dashed line represents the reported uncertainty estimate from the Vegas routine, while the points and solid lines indicate the “true” error, the distance between the reported integral and the true (final Vegas) value of 31552. Again, both sequences perform equally well from this point of view.

In conclusion, we find that the Sobol sequence and the MT19937 pseudo random sequence perform equally well in practice, at least for the benchmarks performed here for one sequence evaluation each. Additional samples could be constructed by using alternative seeds for the MT19937 sequence and by skipping points, or using different initialization numbers, for the Sobol sequence. Theoretically one should benefit from a more even sam-

⁶For the first five warmup iterations the Vegas grid adjustment parameter was set to 1.5, afterwards to 0.8, discarding the warmup estimate but keeping the grid. The number of grid subdivisions was kept at the default of 100.

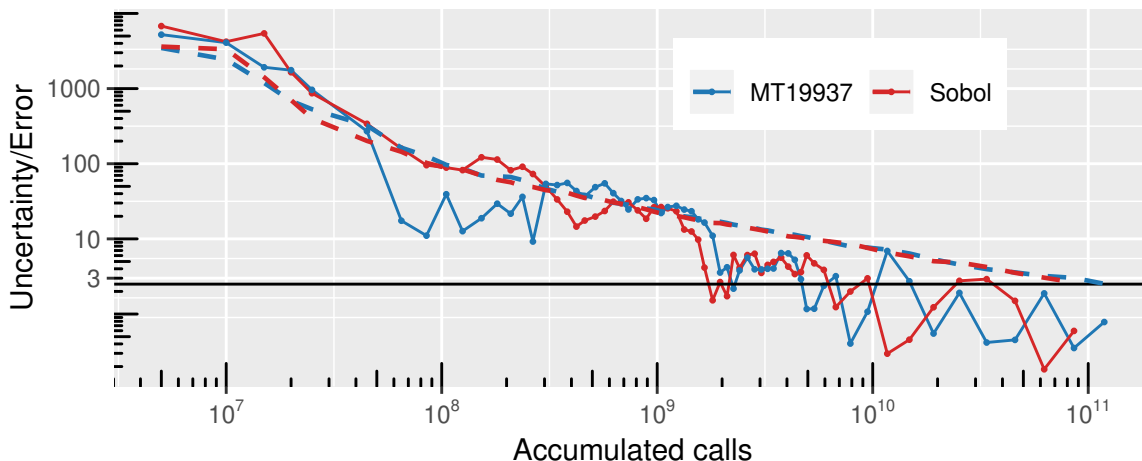


Figure 4. The dependence of the estimated uncertainty of the NNLO Higgs production double real emission contribution (for $\alpha = 1$ and $\tau_{\text{cut}} = 0.002$ GeV) on the accumulated number of calls. The red data corresponds to using the Sobol sequence and blue to the MT19937 sequence. Each point represents a new estimate from a new iteration. The dashed line represents the uncertainty from the Vegas routine, while the points and solid line represent the “true” error, assuming that the average from the last reported numbers of sequences is the most precise with an uncertainty of ± 2.5 .

pling from the Sobol sequence but we have not been able to identify conclusive evidence for this. Nevertheless, we expect that this could be the case for some calculations, particularly ones targetting high-precision kinematical distributions. With this in mind, and given the fact that it performs at least as well as the MT19937 pseudo random sequence, we use the Sobol sequence as default in MCFM-9.0.

3.2 Parallelization and integration uncertainty estimation

In many calculations beyond leading order, and in particular at NNLO, it is important to control numerical uncertainties. Small differences between contributions that should exactly cancel can easily lead to incorrect results and misinterpreted conclusions. To avoid such issues it is essential to have good control of the uncertainties on the numerical integration. In MCFM we have chosen to use OMP and MPI to produce a highly parallelized code, allowing us to reach high levels of precision in a fast and convenient manner. In this subsection we compare with possible alternative approaches, and describe and benchmark the procedure for controlling integration uncertainties.

To assess the benefit of our parallelized code we can contrast it with a possible alternative. A degree of parallelization can be achieved by simply performing many runs with lower Monte Carlo statistics, limited by reasonably-achievable run-times. A naïvely weighted combination is not recommended since the uncertainties of individual runs cannot be trusted for low statistics. We demonstrate this in the following paragraph, then discuss alternative statistical methods for the case when many low statistics runs have to be combined. We again choose the NNLO Higgs production double real emission component with $\tau_{\text{cut}} = 0.002$ GeV and $\alpha = 1.0$ as a representative example for a complicated integration.

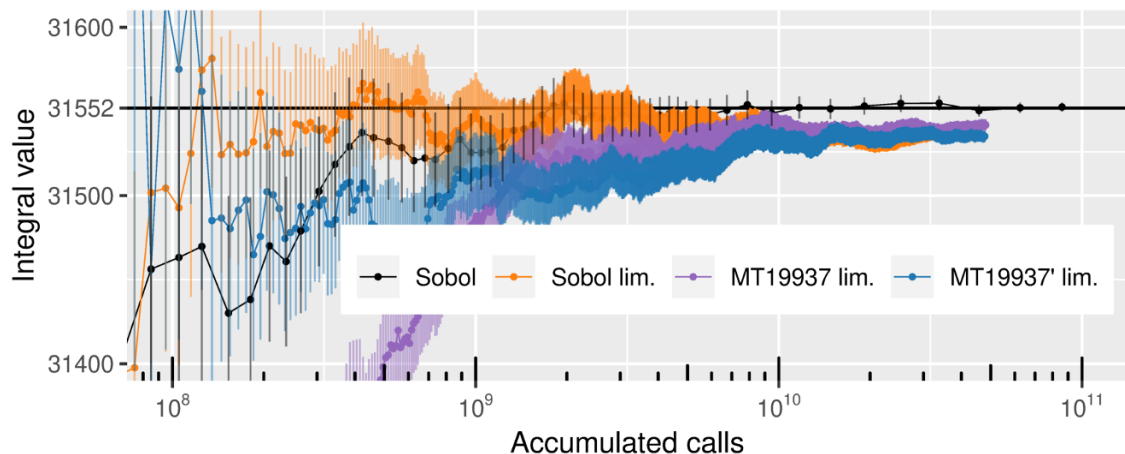


Figure 5. The dependence of the integral for the NNLO Higgs production double real emission contribution (for $\alpha = 1$ and $\tau_{\text{cut}} = 0.002$ GeV) on the accumulated number of calls. The black points correspond to our default setup, while the others (labelled by “lim”) are obtained by limiting the calls per iteration to 10^6 .

Naïvely weighted combination of low statistics runs. Our low statistics runs correspond to running the integration for about 4500 iterations, limiting the calls per iteration to 10 million calls each for the Sobol and MT19937 sequences. This is demonstrated in figure 5, where the curves “lim.” have been obtained with the call limit applied. Two runs for the MT19937 sequence with different seeds are displayed and one run with the Sobol sequence. For comparison the Sobol run of the previous subsection with a steady increase of the number of calls per iteration is also shown. In this scenario the limited Sobol sequence (orange) vastly outperforms the MT19937 sequences (blue and purple) in the region of a low number of total calls, still giving reliable error estimates. However, even though individual iteration uncertainties are $\lesssim 1\%$, in the final stages *both* sequences underestimate the true error.

When we directly consider the dependence of the integration error on the number of calls, in similar fashion as figure 4, we obtain the results shown in figure 6. It becomes clear that when limiting the number of calls the estimated error is severely underestimated for the MT19937 sequences, in comparison with the true error. The limited Sobol sequence performs somewhat better in this case. Nevertheless, all limited sequences lead to significantly underestimated errors and systematically biased results in the final integration stages. This procedure essentially corresponds to a naïvely weighted combination of many low statistics runs and its outcome is more than clear: the integral estimates lock in at the wrong value with underestimated errors.

Although the absolute effects of numerical precision appear small here, to match the expected systematic cutoff effect for the Higgs NNLO cross section (0.2% uncertainty in MCFM-8.0) they must be controlled at the same level. This study demonstrates that, indeed, limiting the number of calls per iteration is insufficient to achieve the desired precision. For this reason we believe that it is important for the integration procedure to




Figure 6. The dependence of the estimated uncertainty of the NNLO Higgs production double real emission contribution (for $\alpha = 1$ and $\tau_{\text{cut}} = 0.002 \text{ GeV}$) on the accumulated number of calls. Each point represents a new estimate from a new iteration. The dashed line represents the uncertainty from the Vegas routine, while the points and solid line represent the “true” uncertainty, assuming an estimate of 31552 ± 2.5 from the previous section.

not only be highly parallelized, but also resumable in order that the number of integration points can be continually increased as required.

Weighted combination of “pseudoruns”. An example of a statistical combination that avoids the weighted combination of runs with underestimated uncertainties is provided in ref. [58], in the context of a NNLO calculation of Z +jet production using the code NNLOJET. The prescription for validating and testing the combination of results is presented in detail there; here we replicate this methodology in MCFM and compare it with our usual, fully parallelized, approach.

The study described in ref. [58] proceeds as follows. After a warmup run to produce a Vegas grid, this grid is used as the basis for thousands of subsequent low-statistics runs. From these, a sample of k individual runs can be combined in an unweighted average to form a “pseudorun” with an associated variance. Pseudoruns can then be combined in a statistical manner, using a weighted average, to obtain the final result. The dependence of this result on the size of the pseudoruns, i.e. on k , is then used to assess the validity of the computed quantity. For small k ($k \gtrsim 1$) it can be very easy to obtain an incorrect result for the integral since each pseudorun may lack sufficient statistics — but the final reported error is small, due to the large number of pseudoruns. Conversely, for large k the estimate of the integral is more reliable, but the uncertainty estimate is much larger. The method is based on identifying a region of k for which results agree within the computed uncertainties. Although this appears reasonable, in principle it requires a careful case-by-case study of each observable to justify the choice of k .

We replicate this method for the calculation of the double real emission contribution to NNLO Higgs production, where in our study of the previous subsection we generated about 4500 iterations with 10 million calls each for the Sobol sequence and one seed of the

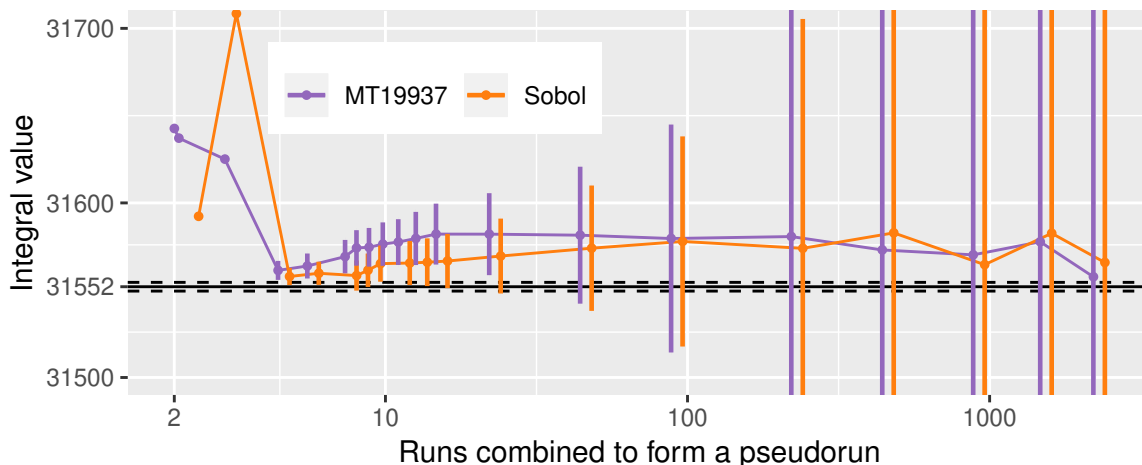


Figure 7. Study similar to ref. [58]: we recombine our data of about 4500 iterations with 10 million calls each into unweighted pseudoruns, computing the uncertainty using the sample variance. These pseudoruns are then combined into the integral estimate through a weighted sum. The true result, computed from a separate, longer run, is indicated by the black line with dashed uncertainties.

MT19937 sequence. From these we can then combine runs into pseudoruns as described above. Our results are shown in figure 7, as a function of the number of runs combined (k). This displays the features noted above, at both small and large values of k , as well as a relatively stable central plateau. The true result for the integral, obtained in the previous subsection, is also displayed. We see that, while the value of the integral is in reasonable agreement for a wide range of k , it is not clear exactly what value to use. Moreover, it is not clear which value of the uncertainty to trust. From figure 7 it is clear that even if a prescription for picking the integral value and its uncertainty can be given, any reasonable uncertainty estimate is significantly larger than the uncertainty estimate of ± 2.5 from the integration in MCFM with the same number of calls.

Bootstrap. To improve upon the method for combining runs into pseudoruns, and to extract a reliable uncertainty, one can use the well-established Bootstrap and Jackknife techniques.⁷ One starts with N estimates of the integral (data points) forming a set, where N should be sufficiently large. Generally, N should be at least at the order of 10, and in our situation below $N \sim 4500$. Within the bootstrap method one then constructs replica sets and obtains statistical properties from the replicas sets. For example, instead of obtaining one unweighted integral average from the original set, with k replica sets one obtains k unweighted integral averages.

To construct each replica set one randomly picks N times from the N original data points so that each point of the original set could occur multiple times in the resulting replica set. This corresponds to taking a combination with repetition in the language of combinatorics, or random sampling with replacement. In total, the procedure is repeated k times, so one obtains k combinations with repetition of the original set, and each one

⁷We use the bootstrap functions implemented in the GNU R boot package [59–61].

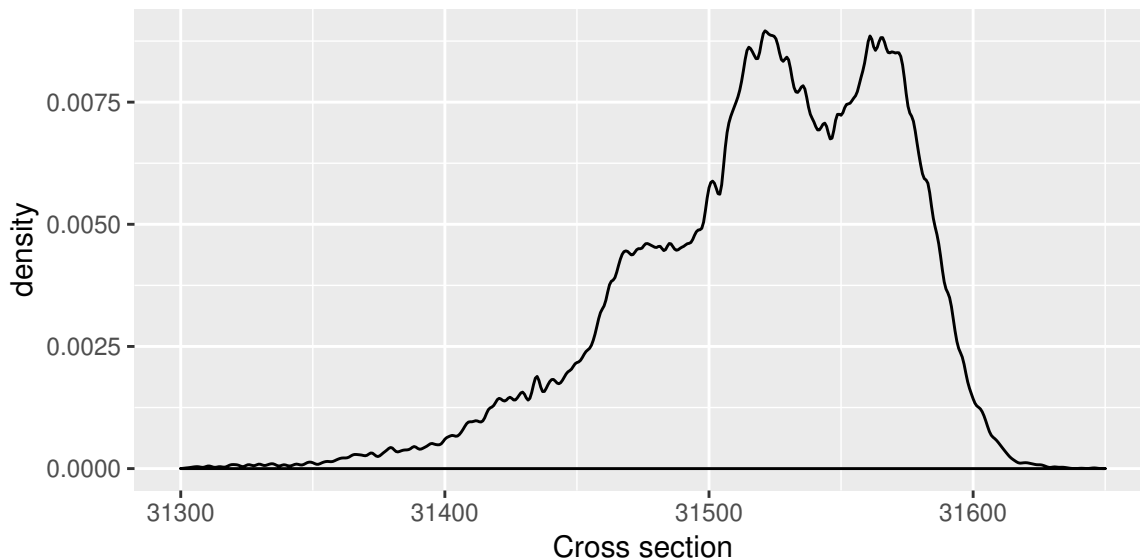


Figure 8. Distribution of the cross section averages from 50,000 bootstrap replicas.

of the k newly resampled sets has N data points. The number of bootstrap replica sets k should be at least of the order of 1000–10000 and we have chosen $k = 50000$.

From the bootstrap replica sets one can then obtain unweighted averages of the data points as estimates for the integral. The generated distribution from the k averages is Gaussian if N and k are sufficiently large. Bootstrap confidence intervals can be constructed by looking at the percentiles of the distribution or using more sophisticated methods [59, 60]. In practice, when N is not asymptotically large, one or a few outliers can cause a non-Gaussian distribution. In that case the bootstrap results are not accurate and can be sensitive to changes in the data points.

Such single outliers can be identified with the jackknife-after-bootstrap method [61]. Within jackknife resampling one leaves out one specific data point of the N data points and computes the integral average using the leftover points. In total $N - 1$ resamplings are obtained, generating a distribution from which percentiles can be obtained. For the jackknife after bootstrap method the k bootstrapped replica sets are taken, and for each data point x out of N , the sets with x removed are taken and analyzed with the jackknife technique. One can then directly analyze the influence each data point x has on the bootstrapped results. While a detailed description of the jackknife-after-bootstrap method is beyond the scope of this paper, it is a standard procedure. We refer to ref. [61] and the software package in ref. [59] for details of the method, and provide some details for our illustrative example below in appendix B.

The distribution of the cross section obtained from a bootstrap with 50,000 replica sets from the full ~ 4500 MT19937 integral estimates is shown in figure 8. We immediately see that this resampled distribution is not Gaussian so that the interpretation of an uncertainty on the average based on percentiles, 31521 ± 49 , is not clear. However, this multi-peaked structure is caused by just a few outliers that can be easily identified using the jackknife-after-bootstrap technique described above, and illustrated in more detail in




Figure 9. Distribution of the cross section averages from bootstrap replica, after removing the worst outlier from the full set in figure 8.

figures 32 and 33 in appendix B. Using such a procedure we can identify a single outlier of 192155 ± 223717 , whose uncertainty is not properly taken into account in this method, and remove it. After doing so we observe the distribution of cross section averages shown in figure 9, that displays a very good Gaussian shape. The resulting estimate of the cross section and uncertainty is 31568 ± 17 , where the error is now expected to be Gaussian and reliable. Removing the worst three outliers gives a further modest improvement, with an estimate of 31559 ± 13 . With the ~ 4500 integral estimates obtained with the Sobol sequence, we obtain a Gaussian distribution outright and an estimate of 31570 ± 28 . Removing the worst four outliers in the same way we obtain an estimate of 31572 ± 15 . The uncertainty could be reduced further to 10 by removing 10 outliers, but in the end the sample size should be increased.

We conclude that the bootstrap technique is a more robust way to obtain an average and uncertainty when limited to a large sample size with a limited number of calls per sample. For example, it allows for a systematic identification and removal of outliers and it can be used to obtain Gaussian confidence intervals of various degrees. Furthermore, it can be automatized straightforwardly and gives enough indicators to estimate the quality of the results. In fact, we have successfully implemented and tested such an approach in MCFM as an alternative to the default Vegas integration uncertainties.

However, this procedure requires a large number of iterations or independent runs. Moreover, a completely parallelized — and resumable — approach allows these issues to be avoided, or at least postponed, since a large number of integration points can be more easily achieved. By continually increasing the number of points used in each Vegas sweep, the code can ensure that the grid adjusts sufficiently to the phase space by monitoring the $\chi^2/\text{iteration}$. This avoids the code “locking in” to an incorrect value for the integral, with a *significantly* underestimated uncertainty. Compared to the other approaches, the

few iterations but larger statistics integration results in trustworthy uncertainties that are generally significantly smaller.

4 Updated NNLO benchmarks and τ_{cut} -dependence

The technique employed in MCFM for the calculation of NNLO processes involves a technical jettiness-slicing parameter τ_{cut} [62, 63], that needs to be chosen small enough to eliminate the dependence on it. In practice one can not choose τ_{cut} arbitrarily small because it induces large cancellations between different components that need to be integrated separately. The dependence on τ_{cut} also does not need to be uniform over the whole phase space. So in kinematical distributions the τ_{cut} dependence and its induced systematic error can be different in each bin. This also holds for cross sections with different cuts, of course. We ship MCFM with presets for τ_{cut} corresponding to conservatively-estimated systematic cutoff effects of less than one percent for total inclusive cross sections.

In this section we first demonstrate order of magnitude improvements through the use of a boosted definition of 0-jettiness and the inclusion of power corrections. We then introduce automatized methods that allow an assessment of whether the cutoff effects are indeed as small as suggested, especially in kinematical distributions in the presence of thresholds.

Improved τ_{cut} dependence. To improve the performance of the jettiness-slicing method, by default this version of the code uses a version of 0-jettiness that incorporates the boost of the Born system [64, 65]. This was introduced already with the NNLO $Z\gamma$ implementation [66] in MCFM-8.1, but here we benchmark the improvement in all of the NNLO calculations included in the code. As a further benefit, we have also implemented the leading power corrections presented in refs. [65, 67] for all appropriate NNLO processes. The combination of these two provides a substantial reduction in computational time required for a given nominal accuracy. As a result, in this section we update the benchmark results for color-singlet processes presented in ref. [21] and extend that exercise to the $\gamma\gamma$ [68] and $Z\gamma$ [66] processes. This allows us to improve the preset τ_{cut} values for our precision goals in the code accordingly.

We make use of a further feature in this version of the code that allows τ_{cut} to be defined on an event-by-event basis. This has the benefit of elucidating the nature of the power corrections more cleanly, as well as avoiding potential numerical issues related to very small values of τ_{cut} in events with very high-energy partons. This dynamic choice of τ_{cut} is defined by,

$$\tau_{\text{cut}} = \epsilon \times m, \tag{4.1}$$

where m is the invariant mass of the Born system, e.g. $m(\ell^-\ell^+)$, $m(\gamma\gamma)$ or $m(Z\gamma)$.

Automatic fitting of the τ_{cut} dependence. With this version of MCFM an array of additional τ_{cut} values can be specified that is sampled in addition to the nominal choice. Using these additional points a fit based on the leading power behavior of the cross sections' τ_{cut} dependence is automatically performed. When no additional τ_{cut} values are chosen an automatic choice is made and this will be discussed in section 4.2.

Since we use correlated τ_{cut} values, it is advantageous to just fit the *correction* with respect to the nominal value with a small uncertainty. Our weighted fit is then performed using the Minpack package [69–71]. We weight by the (small) uncertainties from the correlated τ_{cut} sampling and thus take into account the different uncertainties between larger and smaller values of τ_{cut} .

In our results we take into account the numerical integration uncertainty from the nominal value just as if one τ_{cut} value was chosen. The additional τ_{cut} values are sampled as differences to the nominal value. Since they are highly correlated to the nominal value, their uncertainties are in practice an order of magnitude smaller. From the uncertainty-weighted fit we obtain an uncertainty estimate of the fitted asymptotic $\tau_{\text{cut}} \rightarrow 0$ value. Further below we introduce a prescription to evaluate the goodness of fit and reliability of these results. In the final results one can then combine the result for the nominal value of τ_{cut} with the fitted corrections. The uncertainty of that value is then the combination of the numerical integration uncertainty of the nominal value and the fit uncertainty, which we both separately output in MCFM.

At NLO the leading power τ_{cut} dependence of the cross section starts with $\tau_{\text{cut}} \log(\tau_{\text{cut}})$. We use the following form for fitting, including an additional linear term:

$$\sigma(\tau_{\text{cut}})^{\text{NLO}} = \sigma_0 + c_1 \cdot \tau_{\text{cut}} \cdot \log(\tau_{\text{cut}}/m) + c_2 \cdot \tau_{\text{cut}}, \quad (4.2)$$

where m is the invariant mass of the Born system as in eq. (4.1). At NNLO the dependence starts with $\tau_{\text{cut}} \log^3(\tau_{\text{cut}})$ and we include the subleading term $\tau_{\text{cut}} \log^2(\tau_{\text{cut}})$ as well as an optional linear term:

$$\sigma(\tau_{\text{cut}})^{\text{NNLO}} = \sigma_0 + c_1 \cdot \tau_{\text{cut}} \cdot \log^3(\tau_{\text{cut}}/m) + c_2 \cdot \tau_{\text{cut}} \cdot \log^2(\tau_{\text{cut}}/m) + c_4 \cdot \tau_{\text{cut}}. \quad (4.3)$$

From the fit we extract a reduced χ^2 per degree of freedom value that should be small compared to one for a good fit. The linear subleading term with coefficient c_4 is included if the fit without the linear term has a $\chi^2/\text{d.o.f.}$ of more than one and the inclusion of the linear term improves the χ^2 . We note that while the inclusion of power corrections in principle removes the leading coefficient c_1 at each order, we still include it in order to remain valid for processes where such corrections are not known, as well as to be robust against the effect of cuts.

Our fitting procedure can easily be extended to successively include a tower of sub-leading terms. Such a fit could be relevant for theory applications, for instance to examine residual subleading corrections at extremely high precision [65]. In practice, to obtain the most reliable results at the highest level of precision it is better to choose a smaller value of τ_{cut} , closer to the asymptotic behavior, rather than using very high numerical precision far from the asymptotic region.

4.1 Inclusive τ_{cut} benchmarks

We first show results for the processes for which the effect of leading power corrections is known — namely $gg \rightarrow H, W, Z, ZH$ and WH production — in figure 10. For this inclusive study all the settings and parameters are chosen according to ref. [21]. The ratio

of the MCFM calculations to the known NNLO results, for the correction itself and for the total rate, are shown as a function of the 0-jettiness variable τ_{cut} . In each case the dependence on τ_{cut} is shown using the unboosted definition of 0-jettiness as in MCFM-8.0 (corresponding exactly to the results in ref. [21]), with the boosted definition implemented in this version, and after the further addition of the leading power corrections.

In all cases the improvement from v8.0 to this version is dramatic. Equivalent residual cutoff effects are obtained for τ_{cut} values that are around two orders of magnitude larger. The improvement from using the boosted definition is most dramatic for W and Z production, for which a significant portion of the inclusive cross-section results from events at large rapidities where the unboosted definition performs poorly.

We now turn to the remaining processes, $\gamma\gamma$ and $Z\gamma$ production, for which the leading power corrections are unknown due to their t -channel nature [66]. However, we can still examine the improvement due to the use of the boosted definition of 0-jettiness. For this study we use the setup of the validation section of ref. [68] for the $\gamma\gamma$ process and the validation cuts for the $Z \rightarrow \nu\bar{\nu}$ decay specified in ref. [66] for the $Z\gamma$ process. The results of this study are shown in figure 11. Here the improvement from using the boosted definition is relatively mild due to the fact that the Born system tends to be quite centrally produced.

We now summarize our findings by specifying the default values of the dynamic τ_{cut} values that are used in the code. Previous version of MCFM included presets for 1% and 0.2% cutoff effects for the total inclusive cross sections. In this version we devise a different scheme that improves the quality and allows for an automatic inclusive and differential τ_{cut} fitting with automatically chosen τ_{cut} values to assess the τ_{cut} uncertainty. Our process-dependent preset τ_{cut} values are overall chosen to first satisfy better than 1% cutoff effects in the total inclusive cross section. While a one percent error on the total cross section can be considered small for many processes with larger theory uncertainties, this is not a sufficient criterion that allows for an automatic reliable estimation of systematic τ_{cut} errors. In MCFM the NLO component is independent of the value of τ_{cut} and only the NNLO corrections are computed with jettiness subtractions that include a τ_{cut} dependence. Since the size of the NNLO corrections varies widely between different processes, this means that the NNLO coefficient itself must be probed in the asymptotic regime described by eq. (4.3), independent of the size of the contribution to the total cross section.

To elaborate further, we put an emphasis on probing the asymptotic τ_{cut} regime for the NNLO coefficient itself for the following reason. In the case when the NNLO coefficient is small or even zero one would not have to calculate it very precisely and a failed fit resulting from working in the non-asymptotic regime would be harmless. However in this regime no conclusions about the behavior of the τ_{cut} -dependence can be drawn, so one has to reach and check the asymptotic behavior at least once. It is conceivable that one is far from the asymptotic regime, so that the shape of the power corrections is unknown, but that the NNLO coefficient with the given choice of (large) τ_{cut} is accidentally small. At the same time the true result from the asymptotic regime could mean a sizable NNLO contribution. To ensure that no accidentally small NNLO corrections occur, one thus has to determine or probe the asymptotic regime at least once. Of course such a case would be detected by a bad fit and could be ignored by the user when it is known that the NNLO

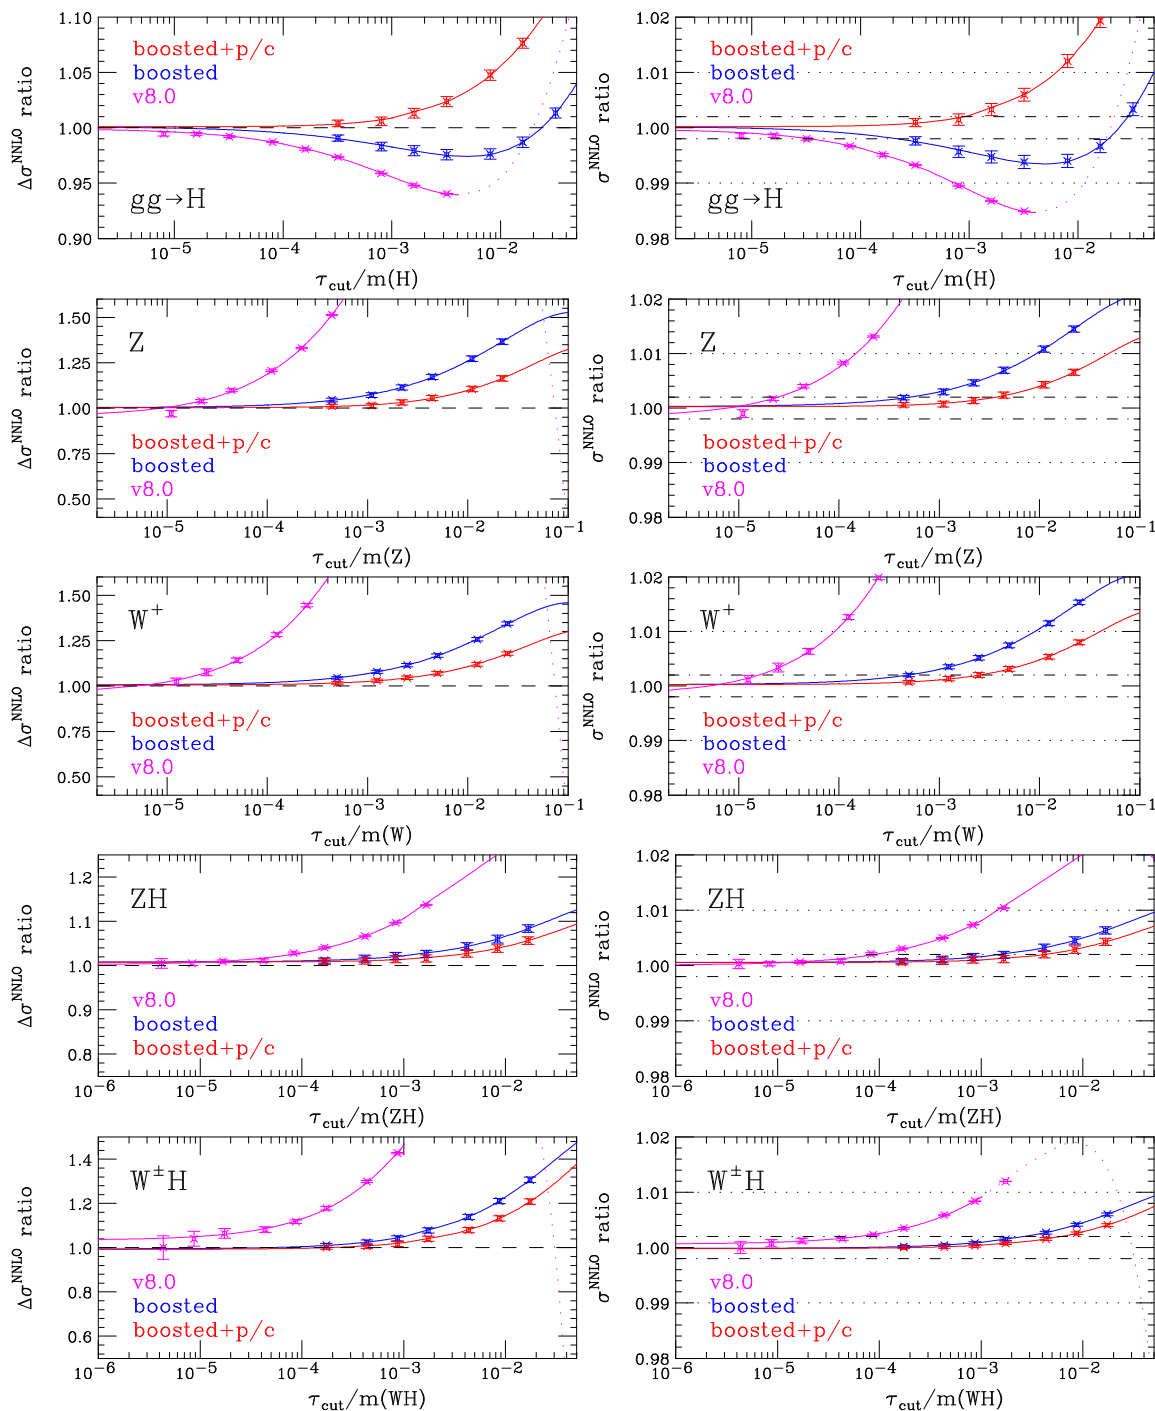


Figure 10. The ratio of the NNLO correction (left) and total NNLO (right), calculated using jettiness slicing, to the known result. The ratios are plotted as a function of the jettiness resolution parameter, normalized to the mass of the Born system. The results using the previous (unboosted) definition of 0-jettiness, as implemented in MCFM v8.0, are shown in magenta. Results using the boosted definition of 0-jettiness are shown in blue, while results after the further addition of power corrections are shown in red. For reference, horizontal lines are shown in the total NNLO plots that represent accuracies of 1% and 2 per mille.

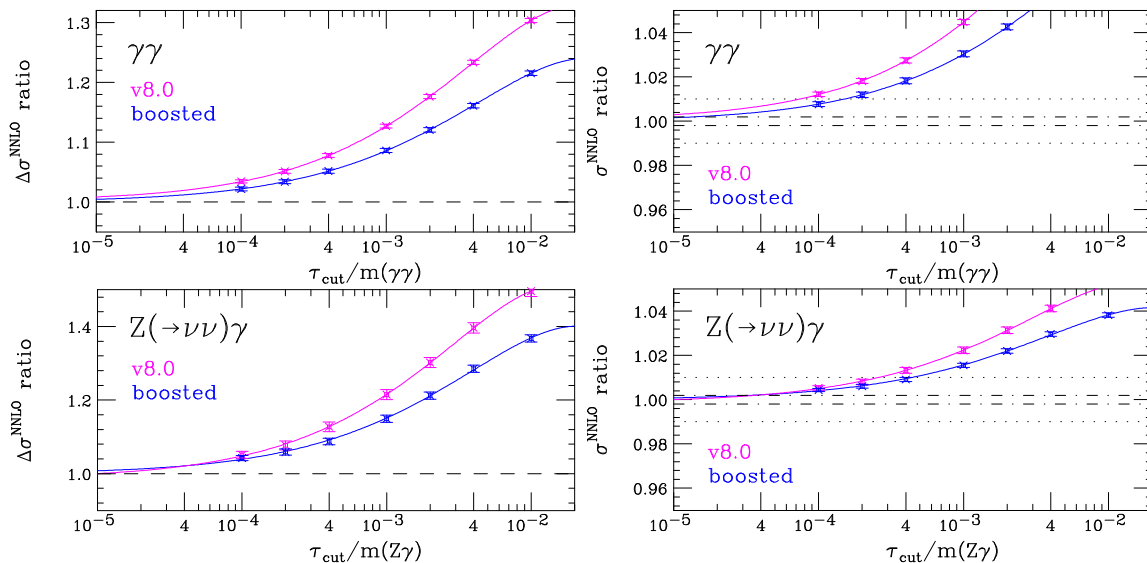


Figure 11. The ratio of the NNLO correction (left) and total NNLO (right), calculated using jettiness slicing, to the asymptotic result. The ratios are plotted as a function of the 0-jettiness resolution parameter scaled to the invariant mass of the Born system, cf. eq. (4.1). The results using the previous (unboosted) definition of 0-jettiness, as implemented in MCFM v8.0, are shown in magenta while results using the boosted definition of 0-jettiness are shown in blue. For reference, horizontal lines are shown in the total NNLO plots representing accuracies of 1% and 2 per mille.

coefficient is small, as discussed below. To ensure precise results in general and make use of the automatic fitting we therefore use a small enough τ_{cut} value that the NNLO coefficient itself is probed in the asymptotic regime.

For example in Drell-Yan production the NNLO coefficient is less than one percent of the total cross section, while for Higgs production it is about 25%. Our goal is to automatically include a sufficient range of larger τ_{cut} values in the fit to improve results and allow for a systematic assessment of τ_{cut} dependence and uncertainties. By default, when enabling a fit, in addition to the nominal value τ_{cut} we use an array of additional values $\kappa \times \tau_{\text{cut}}$ with $\kappa \in \{2, 4, 8, 20, 40\}$. Therefore, we further require that τ_{cut} values as large as 40 times the nominal choice, used for the fitting of the NNLO coefficient, are still in the asymptotic regime and thus described by eq. (4.3).

Our settings are given in table 1. With these criteria we obtain τ_{cut} presets that always give precision better than 1%, yet are orders of magnitude larger than the presets in MCFM-8.0 due to the boosted definition and power corrections. To achieve an even smaller systematic error one can choose values of τ_{cut} according to figures 10 and 11 above.

4.2 Differential τ_{cut} extrapolation and assessment of τ_{cut} uncertainty

While the results in the previous section were highly inclusive, and we believe that in most circumstances the presets of less than one percent τ_{cut} -dependence also apply to the most typically-used cuts and distributions, modern applications and benchmark comparisons require true per mille precision differentially. Then a very careful assessment of the numerical

Process	$\tau_{\text{cut}}/m_{\text{Born}}$	σ^{NNLO} precision
H	4×10^{-3}	0.7%
Z	6×10^{-3}	0.4%
W	6×10^{-3}	0.4%
ZH	3×10^{-3}	0.2%
WH	3×10^{-3}	0.2%
$\gamma\gamma$	10^{-4}	0.8%
$Z\gamma$	3×10^{-4}	0.8%

Table 1. Default τ_{cut} values for each process, together with the expected level of precision when computing the total cross section with that value. As indicated, all τ_{cut} values are specified relative to the invariant mass of the Born system as in eq. (4.1).

uncertainties and τ_{cut} dependence has to be performed. We addressed the former case in section 3. Here we address the latter.

In the previous section the fitted inclusive results were obtained using a wide range of τ_{cut} values to study and display the asymptotic behavior. In practice, a good fit can improve results by an order of magnitude compared to using just the smallest value of τ_{cut} . Here we benchmark such improvements fully differentially and discuss the reliability of the fit. With such a fit one can then also estimate how large the true residual τ_{cut} dependence is and whether it matches the expected uncertainty estimate of the presets in table 1. A successful fit depends on the process, the smallest value of τ_{cut} chosen, the kinematical cuts applied, and the precision of the input data. However, even a failed fit still delivers valuable information that can be used to further assess the quality of the prediction. Here we discuss both cases, describe the output of the MCFM τ_{cut} fitting procedure and provide guidance on how to obtain such improvements and assess the systematic τ_{cut} error.

To our knowledge this is the first (public) code that allows the user to automatically and systematically evaluate the reliability of the slicing calculation and assess true per mille level precision also for kinematical distributions. To perform such analyses in previous versions of MCFM required many more orders of magnitude of computational resources. Nevertheless, some manual inspection is required to be certain about the systematic τ_{cut} uncertainties. This is because any procedure with hard-coded reliability thresholds will eventually fail for some combination of processes and kinematic cuts. So instead of hard-coding fixed warnings we instead outline a systematic procedure that can always be applied to assess the reliability of the results.

Both for the inclusive cross section and fully differentially for the output histograms of the fitted corrections we include gauges for a good fit. The first gauge is the maximum relative numerical uncertainty on the additionally sampled τ_{cut} corrections, describing the quality, or rather uncertainty of the fit's input data. Second, we output the reduced χ^2 of the standard-deviation-weighted fitting, as a direct measure of the fit quality. These two numbers are in addition to the fitted correction itself and its corresponding uncertainty. A

more detailed technical description of the MCFM output files is provided in appendix A. All of these data points should be used to establish a chain of trust as described in the following three steps.

1. The associated uncertainties of the additionally sampled τ_{cut} corrections, as direct input to the fit, have to be reliable for the subsequent steps and the fit to be meaningful. If the uncertainties of the additional τ_{cut} corrections are severely underestimated, one might be led to the wrong conclusion through a fit that appears to work reliably but which is flawed from the outset.

To estimate this, we include the maximum relative numerical uncertainty on the sampled τ_{cut} corrections as information with the fit. While a $\chi^2/\text{iteration}$ could be a useful quantity to consider for that purpose, we do not store the necessary information to compute it fully differentially. Instead, we generally simply find that relative uncertainties of the integrals are underestimated when they are still large. So while a 10% uncertainty on a result indicates only questionable precision, once percent-level precision is reached we deem results generally reliable. In practice one can also check whether the smoothness of differential distribution agrees with the reported uncertainties.

2. Assuming from the first step that the additional τ_{cut} corrections and their uncertainties are reliably estimated, the input data to the fitting procedure can be trusted. For the uncertainty-weighted fit we report the χ^2 per degree of freedom (reduced χ^2), which should generally be close to one, or smaller, for the fit to be reliable. A large χ^2 value means that the fitting formula does not work well. This can either be caused by data that is so precise that additional subleading terms in τ_{cut} need to be included in the fitting form, or that additional terms should be included because the τ_{cut} -dependence is so large that it has not yet reached the asymptotic regime of eqs. (4.2) and (4.3). In practice, for phenomenology the first case is not relevant when “only” per-mille level precision is required, but additional terms can easily be added if required.⁸

This leaves the second case for usual applications, where the large χ^2 value of the fit then means that the τ_{cut} dependence is large and the default τ_{cut} setting for the total inclusive cross section can likely not be trusted in the considered kinematic region. The values of τ_{cut} being probed are not asymptotic and are not described by eq. (4.3). This does not necessarily mean that the full NNLO result has a large τ_{cut} dependence, as only the NNLO coefficient is computed using the jettiness slicing with τ_{cut} dependence. In the case of W or Z production the NNLO coefficients are generally tiny, so larger values of τ_{cut} can be chosen. Nevertheless, one should always ensure that one remains in the asymptotic regime to obtain a reliable estimate of τ_{cut} corrections.

⁸Also, instead of improving the fit with more precise data, one might rather choose a set of smaller τ_{cut} to begin with. Nevertheless, one can of course easily adjust the fitting function in the code of MCFM-9.0 for theory applications.

3. When finally the reduced χ^2 is small or close to 1, after possibly decreasing the nominal and smallest value of τ_{cut} , the fitted correction and its uncertainty can be trusted. When the uncertainty on the fitted correction is small compared to the fitted correction, the fit can often improve results with a 1% cutoff error to the per-mille level. We recommend applying the fitted correction when this chain of trust has been established. When the uncertainty on the fitted correction is large, and after the previous steps it can nevertheless be trusted, it still gives an estimate of the maximum τ_{cut} dependence that can be expected.

In the next subsection we give examples where results have been obtained with sufficient numerical precision and goodness of fit to allow for order of magnitude improvements and assessment of per mille level precision fully differentially. Performing this example first at NLO allows us to compare against cutoff-independent results. We then proceed to two examples where the automatic fitting fails. In the first example the input data is unreliable (large maximum numerical uncertainty on the additionally sampled τ_{cut} values). In the second example we consider a kinematical distribution at NNLO that contains both a threshold region with large τ_{cut} dependence and a region that has no logarithmic τ_{cut} dependence and is fully finite for $\tau_{\text{cut}} \rightarrow 0$. These examples are constructed to show that with the above procedure and all the information given together with the fitted corrections, the cutoff dependence can be studied without further runs even when the fit fails.

4.3 Examples of automatic differential τ_{cut} fitting

First we demonstrate our implementation on NLO results of diphoton production, where we can compare with the cutoff-independent dipole subtraction [72] implementation. At NNLO we can compare fitted results using larger values of τ_{cut} with small values of τ_{cut} that can be considered accurate at the per-mille level. Cuts and parameters are as described later in section 5, but do not matter for the discussion here.

In figure 12 we show the NLO *coefficient* for the transverse momentum distribution of the hardest photon normalized to the result of the exact computation. In the left panel we show results for $\tau_{\text{cut}}/m_{\gamma\gamma}$ values ranging between 0.01 and 0.08, and the fitted result using four points $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.02, 0.04, 0.08, 0.2$. The nominal value 0.01 is not included in the fit, since we just fit the τ_{cut} corrections to the nominal value in our procedure. The right panel shows results with cutoffs smaller by a factor of ten. For the nominal $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.01$ panel the fitting improves the residual systematic uncertainties by over 15%, and its error when compared to the exact result is at the level of at most 1.5%. The improvements when choosing a lower τ_{cut} as in the right panel are smaller, but the fitted result again is within 1% of the cutoff-independent result, while $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.001$ deviates as much as 4%.

From this we conclude that the fit can improve even large values of τ_{cut} with systematic errors of 10% and more. Furthermore, and quite importantly, the fit makes the τ_{cut} dependence uniform over the whole range. Without such a fitting procedure, to guarantee a systematic τ_{cut} error differentially, one has to carefully choose τ_{cut} according to the worst region, which is at small p_T . For the example above this would mean running at $\tau_{\text{cut}}/m_{\gamma\gamma} = 10^{-4}$ instead of 0.01 to achieve the same uniform level of precision for the NLO coefficient.

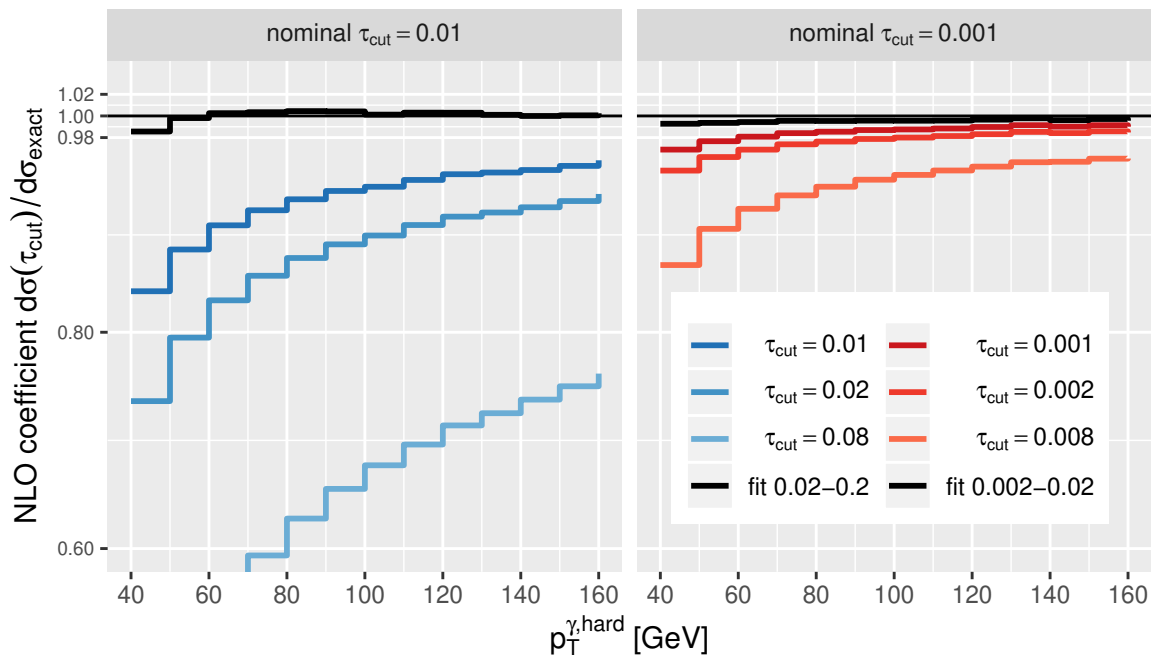


Figure 12. The NLO coefficient of the transverse momentum distribution of the hardest photon in diphoton production. The left panel shows the τ_{cut} dependence for $\tau_{\text{cut}}/m_{\gamma\gamma}$ ranging from 0.01 to 0.08, where the fitted result has been obtained by using four points $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.02, 0.04, 0.08, 0.2$, not using the nominal value of 0.01. The right panel shows the corresponding results for cutoff parameters a factor of ten smaller.

Similarly, we show in figure 13 the NLO coefficient for the diphoton rapidity distribution. While the τ_{cut} dependence is clearly more uniform in the beginning, the conclusions drawn are the same.

For the total NLO result, but also the coefficient alone, we see that one can use the fitted result as the most precise result (following the chain of trust to ensure its reliability) and compare it with the result using the smallest and nominal τ_{cut} value. If the difference is larger than the desired or expected systematic error one should consider the use of a smaller nominal τ_{cut} value.

NNLO results. Corresponding to the NLO coefficient plot in figure 12, we show in figure 14 the results for the NNLO *coefficient*. Since this time we do not have a cutoff-independent result available we directly compare results obtained with two sets of τ_{cut} . The first set ranges from the nominal value of $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.001$ through $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.04$, where the five values $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.002, 0.004, 0.008, 0.02, 0.04$ have been used for the asymptotic fit according to eq. (4.3). A second set of τ_{cut} values smaller by a factor of ten is also shown.

We first see that, as for the NLO coefficient, there is a non-uniform τ_{cut} dependence. Second, both fits agree within 1%, such that we can trust that the fit results are accurate within 1%. In both cases they significantly improve the prediction and make the τ_{cut} dependence uniform. The improvement from the fit is dramatic especially in the first bin, where for $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.001$ the systematic uncertainty without a fit is over 10%.




Figure 13. NLO coefficient for the diphoton rapidity distribution. The left panel shows the dependence for $\tau_{\text{cut}}/m_{\gamma\gamma}$ values ranging from 0.01 to 0.08, where the fitted result has been obtained by using four points $\tau_{\text{cut}}/m_{\gamma\gamma} = 0.02, 0.04, 0.08, 0.2$, not using the nominal value of 0.01. The right panel shows the corresponding results for cutoff parameters a factor of ten smaller.




Figure 14. The NNLO coefficient of the transverse momentum distribution of the hardest photon in diphoton production. All curves are normalized to the fitted result with the smaller set of five $\tau_{\text{cut}}/m_{\gamma\gamma}$ values.

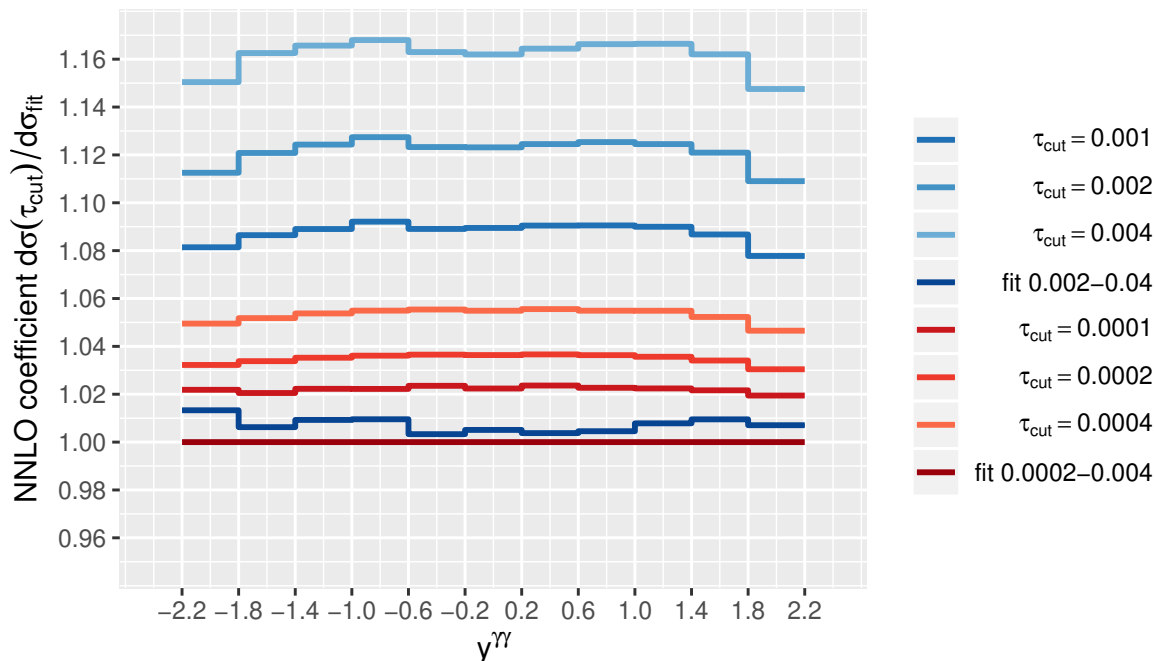


Figure 15. The NNLO coefficient of the diphoton rapidity distribution. All curves are normalized to the fitted result with the smaller set of five $\tau_{\text{cut}}/m_{\gamma\gamma}$ values.

Similar conclusions are drawn from the diphoton rapidity distribution NNLO coefficient in figure 15.

Since in diphoton production the NNLO corrections are large, we show in figure 16 the full NNLO result with corresponding τ_{cut} dependence and fitted results. Our default setting, with a residual cutoff dependence of 0.8%, is $\tau_{\text{cut}}/m_{\gamma\gamma} = 10^{-4}$. The figure shows that, up to numerical artifacts, the same level of precision can be reached using the results of a fit with $\tau_{\text{cut}}/m_{\gamma\gamma} \simeq 10^{-3}$. Nevertheless, we prefer to use the τ_{cut} values in table 1 as our default settings. The additional asymptotic fitting can then be used to check that the cutoff dependence is indeed at the percent level. In the case of a good fit it can also result in a substantial improvement over that, with the added benefit of a uniform systematic cutoff error.

We now consider another representative example, the calculation of the τ^- rapidity spectrum in the process $pp \rightarrow H \rightarrow \tau^- \tau^+$. We use the same setup as for our benchmark calculation of the inclusive $gg \rightarrow H$ cross section. In figure 17 we compare the prediction for this distribution, both in the NNLO coefficient and in the NNLO total, for two choices of τ_{cut} — our default value ($\tau_{\text{cut}}/m_H = 4 \times 10^{-3}$) and one corresponding to an expected 0.2% accuracy in the inclusive cross section ($\tau_{\text{cut}}/m_H = 4 \times 10^{-4}$). We first see that the residual τ_{cut} -dependence of these two choices of τ_{cut} agrees within the expectation from the presets in table 1, although this is of course not guaranteed since the table is based only on the inclusive cross section.

The figure also indicates the result of the fit that is automatically provided by MCFM for the case of the default value of τ_{cut} . The shaded band corresponds to the fit uncertainty




Figure 16. The full NNLO prediction for the transverse momentum distribution of the hardest photon in diphoton production. All curves are normalized to the fitted result with the smaller set of five $\tau_{\text{cut}}/m_{\gamma\gamma}$ values.

and does not include other numerical uncertainties. In the bulk of the rapidity range the fit is in excellent agreement with the 0.2% result, indicating a clear improvement from the non-fitted calculation. In the tails of the distribution, for $|y| > 3$, the fit is no longer reliable since the input data to the fit is calculated with insufficient precision only at the order of 5–10%, corresponding to a failure of step 1 of our prescription above. We have denoted this with a dashed line where the numerical uncertainty on the input data becomes more than 5%. One clearly observes that the reported fit uncertainty is too small to explain the discrepancy between the two fixed τ_{cut} values. However, when taking into account the numerical uncertainties of all contributions one still observes agreement. In this instance, instead of increasing the numerical precision in the tails for a reliable fit, either through an additional run with cuts or a reweighting function, it may be less expensive to simply run the code with a smaller value of τ_{cut} to draw better conclusions about the residual τ_{cut} dependence.

Finally, in figure 18 we show the transverse momentum distribution of the electron produced in $pp \rightarrow Z \rightarrow e^-e^+$, again using our benchmark setup. As above, we compare results using the default value $\tau_{\text{cut}}/m_Z = 6 \times 10^{-3}$ with results targeting 0.2% precision, $\tau_{\text{cut}}/m_Z = 10^{-3}$. At low transverse momenta, $p_T(e^-) \lesssim 35 \text{ GeV}$ the situation is much the same as observed above in figure 17 — the two results are in agreement within expectations. Although the quality of the fit is fine, that is we successfully checked that the input data for the fit has reliable uncertainties and that the $\chi^2/\text{d.o.f.}$ from the fit is small, the fitted corrections still have large numerical uncertainties. To improve upon this and improve the predictions with the fit would require more numerical precision. But we have still gained

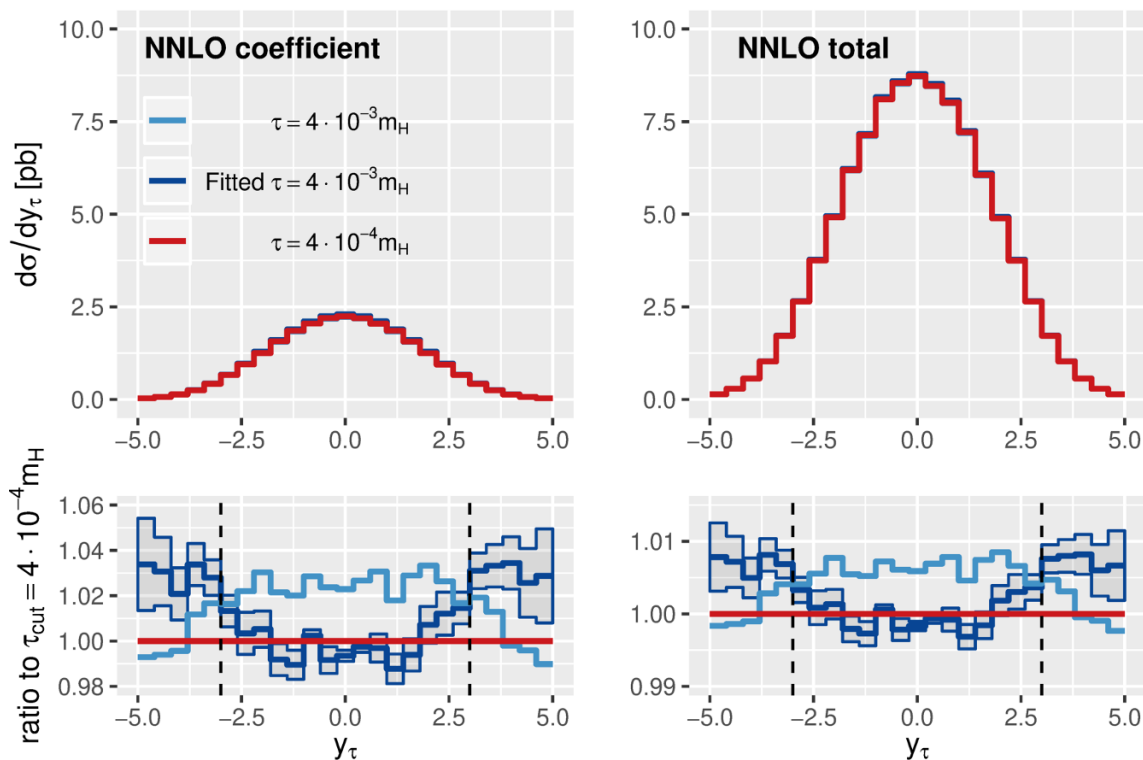


Figure 17. τ^- rapidity distributions in $pp \rightarrow H \rightarrow \tau^- \tau^+$, corresponding to the NNLO coefficient (left) and the NNLO total (right). The histograms for the fit result are included with the fit uncertainties as a shaded band, not including the numerical uncertainties (represented by the noise in the data). The dashed line indicates the point at which the fit is no longer trustworthy (see text).

valuable information: since the fit can be trusted, the fitted corrections including their uncertainty (the shaded band) give an estimate of the residual τ_{cut} dependence.

In the region around $m_Z/2$ the fixed order predictions are not reliable since this distribution is exactly zero for $p_T(e^-) > m_Z/2$ in the leading order calculation. This is because the benchmark setup of ref. [21] uses the zero-width approximation. For the same reason, the only τ_{cut} dependence beyond $m_Z/2$ comes from subleading power corrections to jettness factorization, so that there is no logarithmic dependence on τ_{cut} in this region. Since the residual τ_{cut} dependence is not captured by our fit form, the code reports poor fits in the region $p_T(e^-) \gtrsim 38$ GeV. As a result we have only displayed the fit up to this point (the dashed line in figure 18). Following our own advice we have repeated this study with an even smaller cutoff, $\tau_{\text{cut}}/m_Z = 10^{-4}$, and verified that the result shown in figure 18 is accurate at the per mille level for $p_T(e^-) \gtrsim 50$ GeV.

Summary. While in all the studies in this section we have compared results using two choices of nominal τ_{cut} , in most cases conclusions for the reliability and magnitude of residual finite τ_{cut} effects could be drawn from the fit and the automatically sampled τ_{cut} dependence alone. No further runs with additional nominal τ_{cut} values are necessary

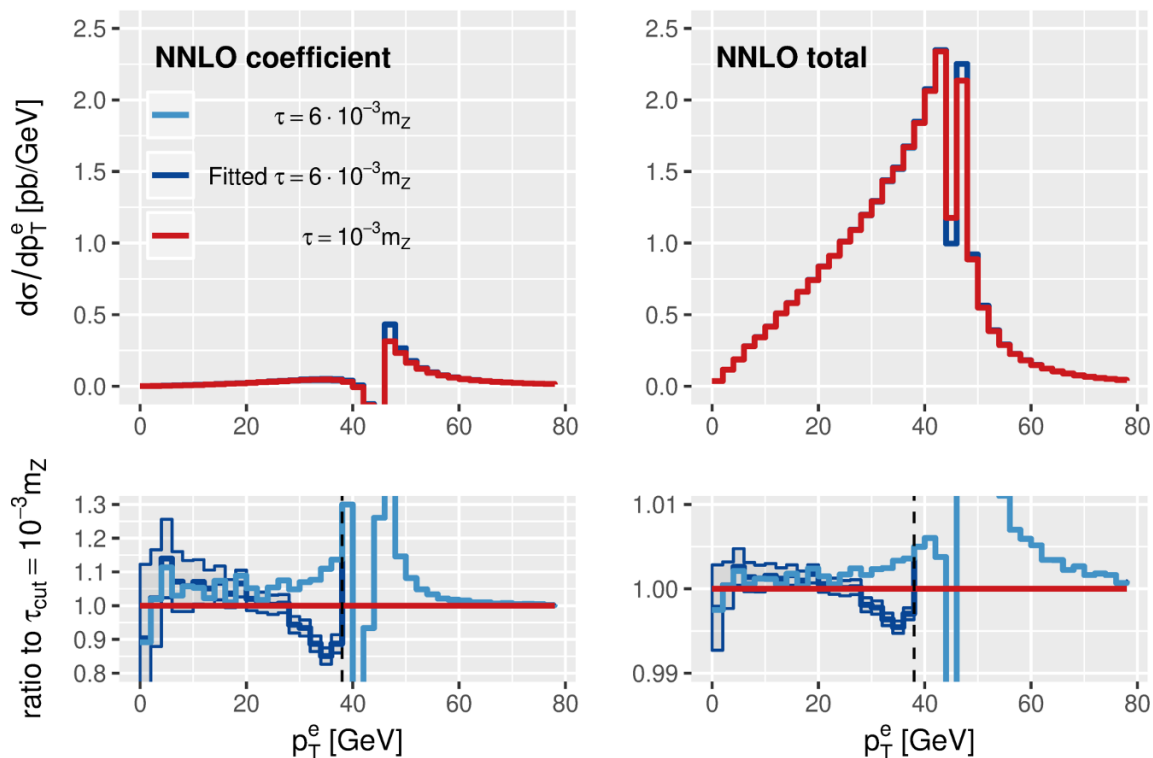


Figure 18. e^- transverse momentum distributions in $pp \rightarrow Z \rightarrow e^-e^+$, corresponding to the NNLO coefficient (left) and the NNLO total (right). The histograms for the fit result are included with the fit uncertainties as a shaded band, not including the numerical uncertainties (represented by the noise in the data). The dashed line indicates the point at which the fit is no longer trustworthy (see text).

for most practical applications. This is in contrast to previous versions of MCFM and other codes that provide no means to assess reliability of results fully differentially without additional extra runs.

MCFM outputs histograms for the additionally sampled τ_{cut} values separately in addition to the nominal τ_{cut} histogram. Furthermore, the fitted corrections are also output separately. In the best case the τ_{cut} dependence is asymptotic and the numerical precision is good, so that the fitted corrections improve the nominal τ_{cut} result and at the same time estimate the error made by using the nominal τ_{cut} value. If the numerical precision is insufficient the fit might still be good, but the fitted corrections have large uncertainties, adding noise. This is still valuable information, since it estimates the maximum τ_{cut} dependence.

The fitting procedure can fail for the following reasons. Either the input data cannot be trusted (generally uncertainties are then still large, around 5–10%), or the input data uncertainties are small, but the fit itself is poor with a large $\chi^2/\text{d.o.f.}$. In the first case one can increase the numerical precision or manually inspect the τ_{cut} dependence instead of relying on the fit. In the latter case, the calculation is not in the asymptotic regime where the logarithmic form of the fit function applies. This can signify an unusually large τ_{cut} dependence, which requires a manual inspection of the τ_{cut} dependence and/or a further

run for this kinematical region with a lower nominal τ_{cut} value. It can also signify that the region of phase space is only populated starting at NLO. For this region there is no logarithmic τ_{cut} dependence and so the fit fails. In this case a manual inspection should demonstrate a small τ_{cut} dependence.

5 Precision PDF uncertainties and PDF set differences

Within the realm of uncertainties for cross section predictions, PDF uncertainties are often large compared to residual theoretical uncertainties. For instance, uncertainties from missing higher order corrections are generally no longer leading and, even for inclusive Higgs production, the PDF+ α_s uncertainty estimates range from three [73] to thirteen percent [47]. They are in the same range as the combined theory uncertainties of about seven percent [73]. Since the LHC is able to collect data on far more than just total cross sections, the real goal is to compute PDF uncertainties differentially for kinematic distributions at NNLO precision.

However, PDF uncertainties are among the most computationally expensive to study. This is due to the need to sample many different PDF sets, or PDF set members. With PDF uncertainties at the few percent level, one wishes to have per mille level numerical uncertainties that are sufficiently small in comparison. The most complicated NNLO processes (in MCFM) require hours up to days on a 10–20 node cluster to reach the desired 1–2 per mille level of numerical precision also differentially. The most naïve implementation of a PDF uncertainty calculation then requires an individual evaluation for every PDF set member, which for our demonstrations below with PDF uncertainties for six chosen PDF sets, multiplies required resources by 371, the number of individual PDF set members. The small cluster is now no longer sufficient, and either one has to run the calculation for months on the small system or days on thousands of nodes.

It is clear that the naïve evaluation of PDF uncertainties represents a substantial additional burden. To avoid this situation, in MCFM we take into account the correlations between matrix elements evaluated with different PDFs by integrating over all PDFs at once, storing only difference information. This not only holds for one PDF set, but for any number of PDF sets, only limited by the available system memory, see appendix A.2. In practice one then gains several orders of magnitude in performance. With these improvements, the time to compute NNLO PDF uncertainties and differences between PDF sets for hundreds of PDF members simultaneously at the per mille level is comparable to the time it takes to calculate the nominal central prediction to the per mille level. This is because one needs a significantly smaller sample size to achieve the required accuracy in the difference information.

The capability of our efficient implementation of multi-set PDF uncertainties is demonstrated in figure 19. Shown are the NNLO W^+ rapidity distributions for a total of six different PDF sets, each with their own PDF uncertainties, normalized to the PDF4LHC central value.⁹ Note that we have included the full W^+ boson decay in the integration, but we choose to show the W^+ rapidity distribution since it is more directly related to

⁹See also figure 1 for the positron rapidity distributions obtained in the same run.

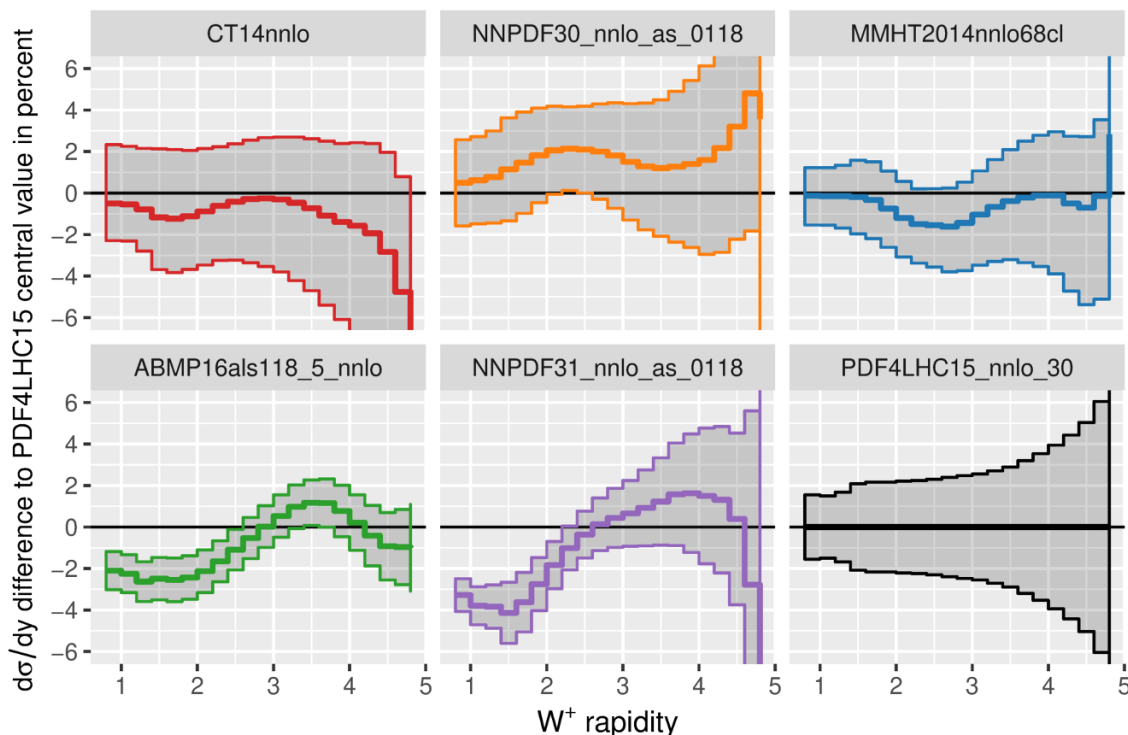


Figure 19. NNLO W^+ rapidity distributions with PDF uncertainty bands after applying a lepton rapidity cut, $2.0 < y_l < 4.5$. All curves are normalized to the central PDF4LHC prediction.

the momentum fraction variables that parametrize the PDFs. These NNLO results were obtained by running less than a day on our cluster using 16 nodes (2010 generation) with 12 cores each. A total of 371 PDFs were evaluated for each phase space point. Not only are the PDF uncertainties predicted within numerical sub per mille level precision, but one can read off sub per mille level differences between different PDF sets fully differentially. This is an invaluable benefit for detailed studies on the impact of different data and methods on predictions at NNLO. For instance we note that while use of the PDF4LHC envelope set gives a first estimate of the uncertainties, figure 19 also clearly shows its shortcomings: the averaging obscures the rich structure of the individual PDF determinations shown in the upper row of the figure. Precision studies such as this one can help to illuminate the differences between the central values, and the corresponding uncertainty bands, of all the sets.

NNLO PDF uncertainties using lower order matrix elements. Depending on the number of PDF sets, and considering bounds on CPU or memory availability, the gain from correlations might not be enough to perform a calculation on a local desktop computer. In this case it is natural to ask whether it is sufficient to compute relative PDF uncertainties at a lower perturbative order. This would cut the computing time by another few orders of magnitude, especially if it were possible to use LO matrix elements together with NNLO PDFs to estimate the PDF uncertainties for the full NNLO result at the per mille level. The difference between using NNLO matrix elements and LO matrix elements with six

PDF sets at once, for example, can be a day of run-time on 10–20 nodes on a cluster or a few minutes to hours on a desktop computer.

The question then is: when is it possible to use lower order matrix elements for the extraction of PDF uncertainties and to examine precise differences between PDF sets? While generally a flat k -factor in higher perturbative orders might indicate that it is acceptable to use the lower order matrix elements, this is neither necessary nor sufficient. A flat k -factor is certainly beneficial though, since it means that a large portion of the observable dependence essentially factorizes with respect to the lower order matrix element. In the same way the dependence on the parton momentum fractions would have to be very similar at each order of the calculation in order to reliably use the lower order matrix element to extract PDF uncertainties at a higher order.

In this section we assess the validity of such a procedure by studying the rapidity distribution of the color singlet system for all NNLO processes included in this version of MCFM. Since the rapidity of the color singlet system translates directly to the parton momentum fractions, we choose this kinematical distribution as the most representative example. Other distributions obtain some form of smearing, but of course can be computed in the same way.

Decay channels and cuts. For our demonstrations the set of decay channels and cuts is as follows. In all cases the W^+ boson decays semi-leptonically, the Z boson leptonically and the Higgs boson into $\tau^+\tau^-$. We do not explicitly present results for W^- , which are qualitatively the same as the ones shown for W^+ production.

Except for the single boson processes W^+ , Z and H , all cuts are standard cuts as in MCFM-8.0 and MCFM-9.0. For the single boson processes W^+ and Z we additionally restrict the rapidity of the leptons to lie between 2.0 and 4.25, motivated by the LHCb experiment. For H production we relax the τ -rapidities y to $-5 \leq y \leq 5$ to probe extreme regions of the phase space. We perform all calculations with the default set of EW parameters as in MCFM-9.0 and for the LHC operating at $\sqrt{s} = 14$ TeV. We allow all vector bosons to be off-shell (`zerowidth` is `.false.`) and include their decays (`removebr` is `.false.`). For parameters that are set in the input file we use,

$$m_H = 125 \text{ GeV}, \quad m_t = 173.3 \text{ GeV}, \quad m_b = 4.66 \text{ GeV}, \quad (5.1)$$

and we use $\mu_F = \mu_R = Q^2$ (i.e. we set `dynamicsscale` equal to either `m(34)` or `m(345)` or `m(3456)`, as appropriate). Our generic set of cuts is,

$$\begin{aligned} p_T(\text{lepton}) &> 20 \text{ GeV}, & |\eta(\text{lepton})| &< 2.4, \\ p_T(\text{photon 1}) &> 40 \text{ GeV}, & p_T(\text{photon 2}) &> 25 \text{ GeV}, \\ |\eta(\text{photon})| &< 2.5, & \Delta R(\text{photon 1, photon 2}) &> 0.4, \\ \Delta R(\text{lepton, photon}) &> 0.7, \\ E_T^{\text{miss}} &> 30 \text{ GeV}. \end{aligned} \quad (5.2)$$

For Z as well as ZH and $Z\gamma$ production we also impose a minimum Z^* virtuality (`m34min`) of 40 GeV.

5.1 Rapidities in single- and diboson production

To study and exemplify the use and calculation of differential NNLO PDF uncertainties we first show the rapidity distributions of W and Z bosons in the forward region, where one might expect to see the most substantial differences between PDFs and the largest uncertainties. This region is of particular interest for the W [74, 75] and Z [74, 76, 77] production measurements in LHCb.

W^+ rapidities. We first compute the W^+ rapidity distributions with PDF uncertainties for the six PDF sets ABMP16als118_5_nnlo [78], CT14NNLO [79], MMHT2014nnlo68c1 [80], NNPDF30_nnlo_as_0118 [81], NNPDF31_nnlo_as_0118 [82] and PDF4LHC15_nnlo_30 [46]. We compute the distributions using NLO and LO matrix elements and with full NNLO matrix elements using two values of τ_{cut} .

Our results are shown in figure 20 where, as in figure 19, we have normalized all distributions to the PDF4LHC central value. The first column shows the normalized NNLO cross section for $\tau_{\text{cut}} = 2$ GeV, corresponding to slicing cutoff effects of about 1% in the total cross section. The second column shows the difference between the NNLO $\tau_{\text{cut}} = 2$ GeV result and the NNLO result with $\tau_{\text{cut}} = 0.1$ GeV, which corresponds to slicing cutoff effects that are smaller than 0.2%. NLO and LO matrix elements are used for the last two columns, respectively. All these differences are magnified by a factor of 10. Difference values above one mean that the nominal NNLO uncertainties used for normalization are larger, whereas values smaller than one mean that the uncertainties are smaller. The latter case can be seen, for example, in the comparison with NLO matrix elements at the lowest displayed rapidities. While for most considered PDF sets the differences just result in a line, the CT14 and MMHT PDF sets show a band. This is because these sets allow for asymmetric PDF uncertainties, and the upper and lower bounds behave slightly differently with respect to the NNLO result.

We first observe that the differences between the two NNLO calculations are at the 1–2 per mille level, entirely compatible with numerical noise. In principle one might only expect half-percent or so agreement due to the finite cutoff values corresponding to better than 1% and 0.2% effects, but this does not appear to be the case. The difference between using the NLO and NNLO calculations reaches about the half a percent level at the largest rapidities, and is only a little larger again when using just the LO matrix element. Of course the W^+ rapidity distribution is not a directly-observable quantity, but we have checked that the same conclusions apply for the lepton rapidity distribution.

We are now in a position to judge whether, for this process, it is acceptable to use LO or NLO matrix elements to estimate NNLO PDF uncertainties. The answer of courses depends on the level of precision that is required. LO matrix elements do not change the PDF uncertainty predictions in W^+ production for lower rapidities significantly, but change uncertainties by up to one percent for most PDF sets towards the end of the rapidity range shown here. While using NLO matrix elements reduces that difference a little, one ultimately has to rely on the full NNLO calculation for uncertainty estimates that are more precise than 0.5%. When using lower order matrix elements one can still compare features of different PDF sets at the per mille level, since these are relatively unaffected by the




Figure 20. Normalized W^+ rapidity distributions in the forward region. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}} = 2$ GeV. The other three columns show the difference of the PDF uncertainties when using NNLO ($\tau_{\text{cut}} = 0.1$ GeV), NLO and LO matrix elements, respectively. The differences are magnified by a factor of 10. Rows represent different PDF sets.

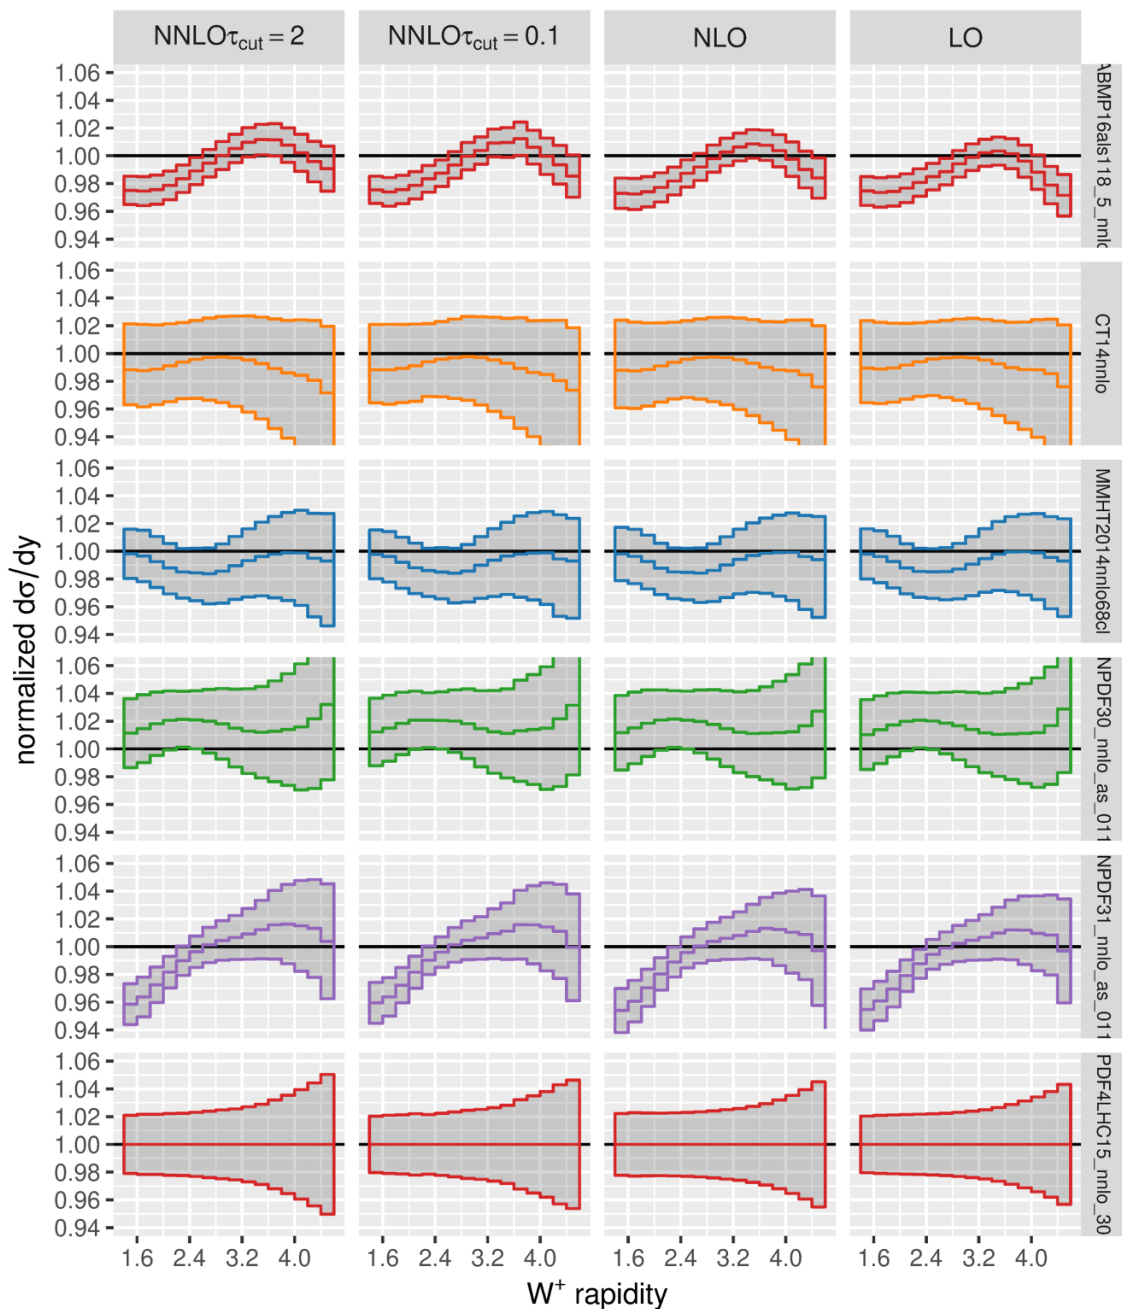


Figure 21. Normalized W^+ rapidity distributions in the forward region. The columns represent using NNLO ($\tau_{\text{cut}} = 2 \text{ GeV}$), NNLO ($\tau_{\text{cut}} = 0.1 \text{ GeV}$), NLO and LO matrix elements, respectively, each one normalized by the PDF4LHC PDF set.

shifts induced at higher orders. This is illustrated in figure 21, which shows the raw results entering figure 20, but without taking differences with respect to the NNLO $\tau_{\text{cut}} = 2 \text{ GeV}$ result. It is clear that features corresponding to individual PDF sets are retained, both qualitatively and quantitatively, when using lower order matrix elements.

Other processes. In order to make a more general statement, we repeat this exercise for all the other processes that may be computed at NNLO in MCFM. We focus on the equivalent of figure 20 in order to readily assess the impact of using matrix elements of different orders to estimate PDF uncertainties at NNLO. Our results are shown in figures 22 to 27. As a reminder, these are all normalized by the NNLO result with a τ_{cut} setting representing 1% cutoff effects on the total cross section and include the 0.2% τ_{cut} setting for comparison.

Although there are qualitative differences between the impact of individual PDF sets on W^+ and Z production, the behavior of the uncertainty bands at each order is rather similar, as shown in figure 22. We are therefore led to similar conclusions for Z production as already discussed above. In order to allow for predictions of the Higgs rapidity also in the very forward region, in figure 23 we have relaxed the lepton rapidity to $|y| < 5$. We then observe that here the differences between using different order matrix elements are also small and only grow systematically to a few per mille towards large rapidities. In all cases the differences between the two τ_{cut} values at NNLO are consistent with one per mille numerical noise for moderate rapidities.

The results for the $Z\gamma$ (figure 24), WH (figure 26) and ZH figure 25 diboson processes are more or less similar, with per mille level differences at most. The situation for $\gamma\gamma$ production (figure 27) is a notable exception, where using LO matrix elements leads to noticeable shape and uncertainty differences of up to one percent compared to the NNLO result. Of course this is not unexpected since, in this case, there is a significant contribution from gluon-initiated processes that only enters at NLO and whose effects are clearly not quite captured with enough precision at LO. In this case one must use at least NLO matrix elements in order to properly assess PDF uncertainties, especially since they themselves are only at the level of $\sim 2\%$.

Summary. In general, we find that NLO matrix elements can be used in many cases for precision studies, sometimes even just LO matrix elements — even to determine per mille level differences between different PDF sets. The difference induced by using a lower order matrix element is far smaller than the differences between using different PDF sets. Therefore, at least for the processes and observables studied here, it is completely sufficient for per mille level applications to use NLO matrix elements for the computation of relative PDF uncertainties. Nevertheless, unless theoretically motivated or plainly demonstrated, we still advocate to use the highest order matrix element possible given the computational resources available. Such a calculation can then be compared with results obtained using lower order matrix elements, which can then be used for further runs or more PDF sets, as required. For such a test it might be enough to perform a run using just a set of central values from different PDF sets, a capability for which MCFM-9.0 allows.

5.2 Higgs boson transverse momentum

As a further example calculation we now consider the Higgs boson transverse momentum distribution. This distribution can be used to constrain the effects of new physics contri-




Figure 22. Normalized Z rapidity distributions in the forward region. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}} = 2$ GeV. The other three columns show the difference of the PDF uncertainties when using NNLO ($\tau_{\text{cut}} = 0.1$ GeV), NLO and LO matrix elements, respectively. The differences are magnified by a factor of 10. Rows represent different PDF sets.




Figure 23. Normalized H rapidity distribution with the lepton rapidities y relaxed to $|y| < 5$. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}} = 0.5 \text{ GeV}$ (1%). The other three columns show the difference of the PDF uncertainties magnified by a factor of 10. Rows represent different PDF sets.

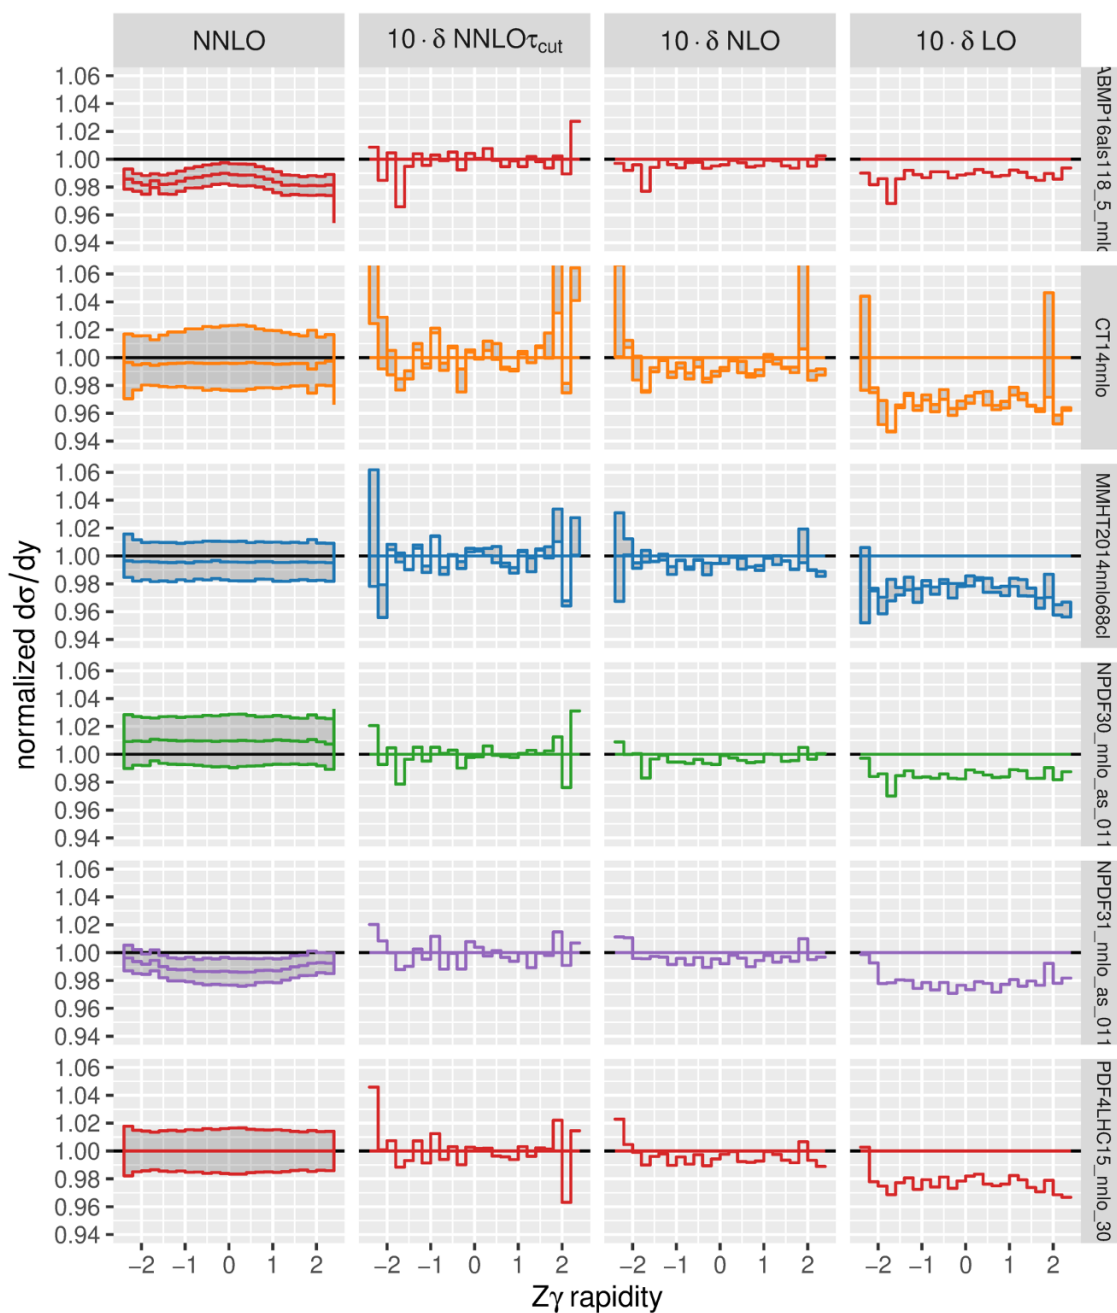


Figure 24. Normalized $Z\gamma$ rapidity distribution. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}}/m_{Z\gamma} = 3 \cdot 10^{-4}$ (1%). The other three columns show the difference of the PDF uncertainties magnified by a factor of 10. Rows represent different PDF sets.




Figure 25. Normalized ZH rapidity distribution. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}} = 4 \text{ GeV}$ (1%). The other three columns show the difference of the PDF uncertainties magnified by a factor of 10. Rows represent different PDF sets.




Figure 26. Normalized WH rapidity distribution. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}} = 4 \text{ GeV}$ (1%). The other three columns show the difference of the PDF uncertainties magnified by a factor of 10. Rows represent different PDF sets.

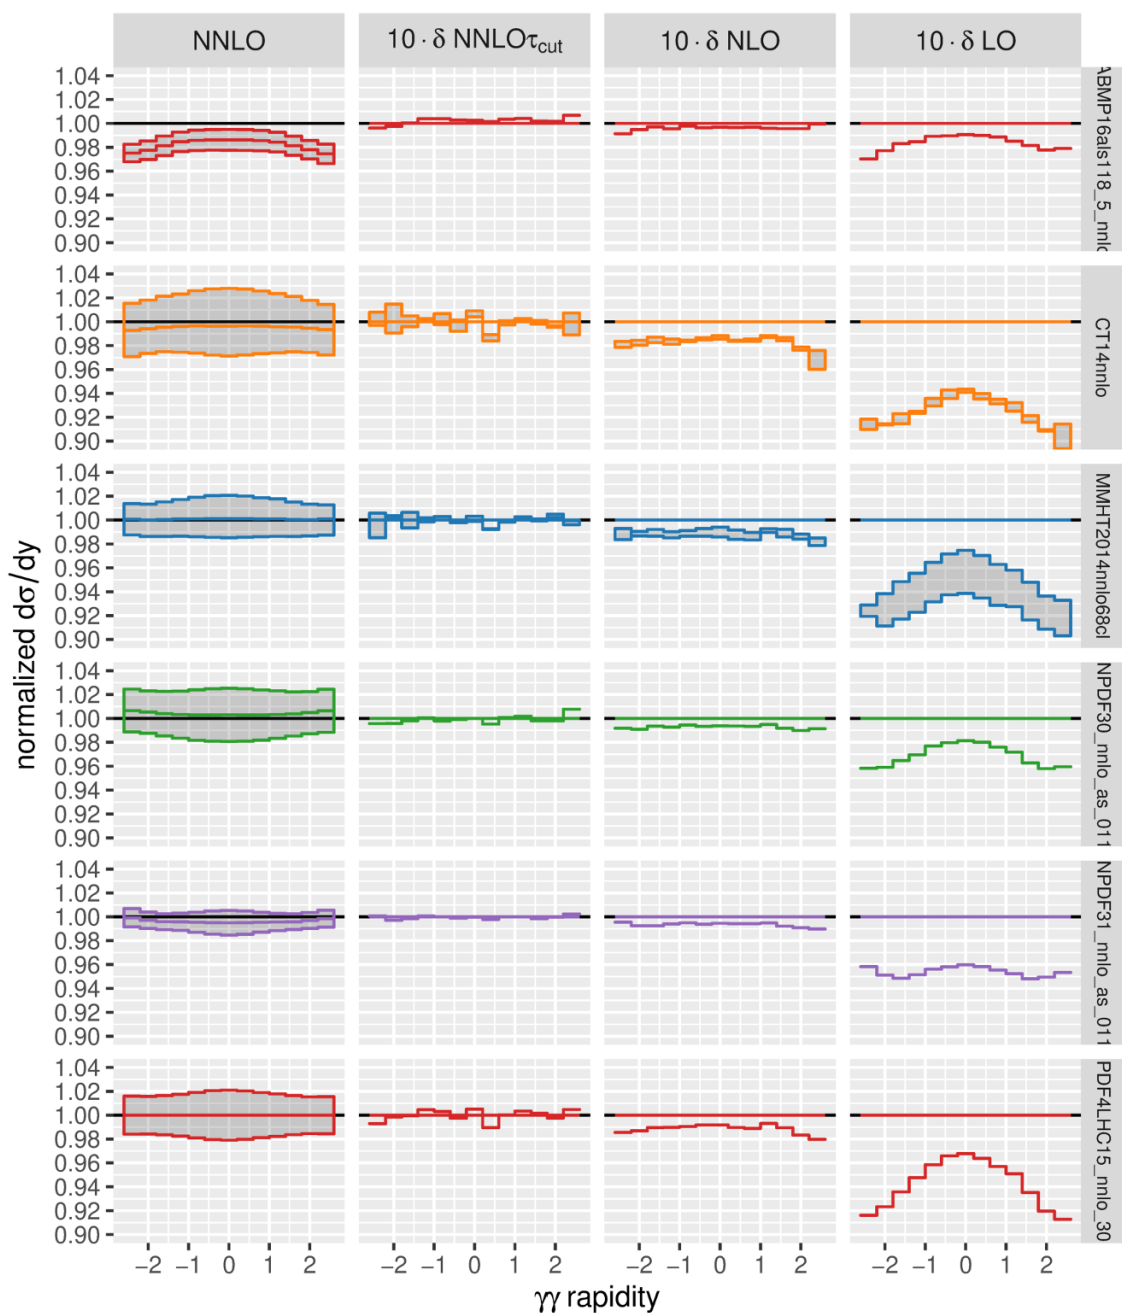


Figure 27. Normalized $\gamma\gamma$ rapidity distribution. The first column shows the NNLO result normalized to the central value of the PDF4LHC set with a value of $\tau_{\text{cut}}/m_{\gamma\gamma} = 10^{-4}$ (1%). The other three columns show the difference of the PDF uncertainties magnified by a factor of 10. Rows represent different PDF sets.

butions in the Higgs sector, and has been used previously in the highly boosted regime to measure the $H \rightarrow \bar{b}b$ channel at the LHC [83].

Perturbative corrections to Higgs+jet production have been computed in an EFT with an integrated-out top-quark up to NNLO [84–90], and also in the full theory with a finite top-quark mass up to NLO [91]. Scale uncertainties at NNLO are about 10% inclusively and differentially. With theory uncertainties at the 10% level, it is hard to imagine that studying PDF uncertainties at the few percent level is a priority. But even for inclusive Higgs production, where theory uncertainties are 5–7% and PDF uncertainties have been estimated to be about 3% [73], recent studies question this precision and suggest PDF uncertainties at the level of thirteen percent [47].

Since predictions for Higgs+jet production at NNLO are computationally highly challenging, one ideally wishes to obtain information for different PDFs using lower order matrix elements. Therefore, in this section we study the Higgs transverse momentum distribution at NLO using NNLO and NLO PDFs. The perturbative NNLO/NLO k -factors are remarkably flat, being almost constant for the most common kinematical distributions. Finite top-quark-mass corrections are also flat with respect to rescaling with the top-quark mass dependent Born cross section, even in the threshold region. See ref. [91] for explicit results at NLO in the Higgs transverse momentum distribution. The combination of these effects means that we anticipate the calculation of PDF uncertainties for this process to be excellent in an EFT description and even when just using NLO matrix elements. The calculation in MCFM provides NLO predictions that go beyond the EFT and take into account mass effects, with residual mass effect uncertainties of 1–2% [92, 93].

To stress test the calculation of PDF uncertainties across kinematical thresholds, we have applied a jet cut of 150 GeV, which means that the region of Higgs boson transverse momentum below this value is not described at LO. Our results for the PDF uncertainties in the Higgs boson transverse momentum distribution up to $p_T = 1$ TeV are shown in figure 28. All PDF sets are chosen to consistently have $\alpha_s(m_Z) = 0.118$ at NLO and NNLO to eliminate large differences caused by different central values of α_s . Results are presented in similar fashion to the plots in the last section, where each column is normalized by the PDF4LHC central value. The left two columns show results with NNLO PDFs but using LO and NLO matrix elements, respectively, whereas the rightmost column shows consistent NLO results. The left two columns indicate small differences of at most two percent between PDF uncertainties estimated using LO and NLO matrix elements.

In the right column, for the consistent NLO calculation, the ABMP16 set shows a peculiar behavior: the predictions are consistently larger than all other sets and reach values that are about 16% larger than the PDF4LHC central value. Uncertainties can in no way account for this behavior. This is despite the fact that using NNLO PDFs (the middle column) the central predictions of ABMP16 are very close to those of the NNPDF31 set and the two are mutually consistent within uncertainties. In that light, only the uncertainty of the ABMP16 NLO fit seems to be somewhat underestimated.

We conclude that with current modern PDF sets one observes predictions for central values that differ by more than ten percent at the largest considered transverse momenta approaching 1 TeV, the region that is the most relevant for new physics analyses. The




Figure 28. PDF uncertainties for the Higgs transverse momentum distribution normalized to PDF4LHC with varying matrix element orders and PDF set orders. The left column uses LO matrix elements with NNLO PDF sets, the middle column NLO matrix elements with NNLO PDF sets, and the right column uses NLO matrix elements with NLO PDF sets. All PDF sets use $\alpha_s(m_Z) = 0.118$.

differences between the sets that take into account the newest LHC data (ABMP16 and NNPDF31) and the older sets can not be explained by PDF uncertainties at the one sigma level. As long as these differences are not resolved and are competing with scale uncertainties at the same level, new physics will be very difficult to constrain. This version of MCFM provides a robust framework for carefully studying such issues in the light of new data.

6 Precision studies of W and Z production

Theory predictions for W and Z production are available at NNLO QCD in MCFM [21] and include one-loop electroweak corrections for Z production [22]. Remaining perturbative truncation uncertainties, estimated by scale variation, are at about the 1% level. In contrast, experimental uncertainties reach the level of a few per mille. To achieve a similar accuracy in the theoretical predictions is highly difficult and even the purely numerical integration uncertainty can easily surpass the expected systematic theory uncertainties.

These numerical difficulties are illustrated, for example, in ref. [94] where the *numerical* differences between the codes FEWZ and DYNNLO, implementing the same NNLO corrections for Drell-Yan production, are found to be 1.2% for fiducial W^+ cross sections. This is similar to the estimated truncation uncertainty and double the size of the experimental uncertainty of 0.6%. This has significant phenomenological impact. The extracted strange quark density, based on HERA and ATLAS W and Z data, is affected at the 8% level by this difference, again larger than the 6% experimental uncertainty [94]. The measurement of $|V_{cs}|$ is similarly impacted by the FEWZ/DYNNLO difference, requiring an additional theoretical uncertainty of 1%, which is again large compared to the theory truncation uncertainty of less than half a percent.

In these cases the problem is exacerbated by a choice of experimental cuts that is pathological for at least one of the calculations, as explained in ref. [95]. Nevertheless, this highlights the importance of precise control over the theoretical calculations, where different techniques for computing higher-order corrections can result in small but important systematic differences. In particular, the effect of any type of technical cutoff or numerical integration artifact must be able to be studied at a level of precision below any other uncertainties, typically around a few per mille. These are topics that we have partially addressed in general already, but in this section we illustrate them in detail by performing a dedicated study of these processes with MCFM-9.0. This follows closely the benchmark comparison study performed in ref. [95], which collects predictions from a number of codes including FEWZ and DYNNLO.

NLO results. As a first point of comparison we have computed the W^+ cross section with the settings in ref. [95].¹⁰ First, for the *tuned comparison* cuts and parameters, we find a value of 2900.1 ± 0.1 pb at NLO compatible with the values in table 3 of ref. [95]; for reference, these range from 2899.0 ± 0.1 pb to 2899.9 ± 0.3 pb, with values in-between

¹⁰To do this we had to use the MSTW2008 NLO PDF set even with the LO matrix element, which is not clear from the text.

with larger uncertainties. These values seem to indicate a discrepancy, but the choice of symmetric cuts could potentially lead to underestimated uncertainties even if all input parameters match.

The cuts used in ref. [95] have an equal minimum transverse momentum requirement on the lepton and neutrino (symmetric cuts). As has been noted in ref. [95], and much longer ago in ref. [96], such symmetric cuts are pathological, since one is highly sensitive to the cancellation of collinear singularities between virtual and real emission corrections. While being exposed to the edge of the singularity is problematic for an efficient integration in general, calculations with an intrinsic slicing or cutoff parameter are affected much worse. Since the jettiness slicing method implemented in MCFM is just such a calculation it is important for us to study the limitation that this presents in practice.

In the following we use the slightly altered *benchmark* cuts and parameters from ref. [95], used for all their NNLO results. Most importantly, the p_T cuts on lepton and neutrino are kept symmetric. In figure 29 we show the NLO cross section with varying minimum neutrino transverse momentum between 25 GeV and 26 GeV in steps of 0.01 GeV. To generate this plot we have computed the neutrino transverse momentum distribution, and subtracted the accumulated p_T bins up to a certain value from the total result. The darkest line is the exact result from a dipole subtraction (exact) calculation. The top panel shows the absolute cross section, while the bottom panel shows the ratio to the exact result. The horizontal black line at 3091 pb denotes the cross section for fully symmetric cuts obtained using the exact calculation. The exact line does not meet the black line at 3091 pb, because the first point with a minimum neutrino transverse momentum of 25.01 GeV, a difference of 0.01 GeV to the symmetric cuts, already leads to a shift of about one per mille. As noted in ref. [96], the approach to the symmetric case has an infinite slope. The other curves are obtained using jettiness slicing with a finite value of τ_{cut} , or correspond to the result of a fit to these values. The figure clearly shows that the dependence on τ_{cut} worsens dramatically as the case of symmetric cuts is approached.

It is clear that the case of fully symmetric cuts requires a very small value of τ_{cut} in order to obtain results that are not affected by the slicing procedure. Of course, for such a value of τ_{cut} it is already computationally very expensive to obtain small numerical uncertainties even at NLO. As we can see, only with a τ_{cut} value of 0.002 is the prediction reliable within one per mille, while for a 1 GeV asymmetry in the cuts one obtains the same precision for a τ_{cut} value of 0.05 GeV. This significantly improves when the asymmetry is increased further. In practice, the symmetric cuts lead to order of magnitude increases in the computational resources that are required. Although this alone might not be a reason to abandon symmetric cuts for experimental studies, it emphasizes the importance of the choice of these cuts for precise numerical comparisons, even at NLO. Referring to figure 29 again, we do however notice that — even for the exact calculation — the NLO cross section displays an odd behavior below 25.3 GeV: the cross section decreases as the cut is reduced. This is simply the remnant of the problematic cancellation of singularities, and reflects the fact that the perturbative calculation is not reliable in this region. This represents a strong motivation to abandon fully symmetric cuts, as has been done some time ago (for the same reasons) in the case of diphoton [97] and dijet [96] production.




Figure 29. NLO W^+ cross section with cuts as in ref. [95], except that the cut on the neutrino transverse momentum is varied (while keeping $p_T^\nu > 25$ GeV). The leftmost point corresponds to $p_T^\nu > 25.01$ GeV. The darkest line displays the result obtained using dipole subtraction, the other curves show various choices of τ_{cut} (in GeV). The upper panel shows the absolute NLO result and the lower panel the result normalized to the smallest displayed value of $\tau_{\text{cut}} = 0.002$ GeV.

NNLO results. The equivalent plot for the NNLO calculation is shown in figure 30. In this case we do not include our τ_{cut} fit at NNLO, since for our reasonably reached numerical precision the uncertainties for the fitted corrections are 50–100% due to the presence of the symmetric cuts. For τ_{cut} values below 0.01 GeV it appears that the residual effect of a finite τ_{cut} is at the few per mille level. On the other hand the τ_{cut} dependence for the fully symmetric cuts is so large that no apparent convergence towards $\tau_{\text{cut}} \rightarrow 0$ is visible. This is in contrast to the τ_{cut} dependence for $p_T^\nu > 26$ GeV (not shown in the plot) that




Figure 30. NNLO W^+ cross section with cuts as in ref. [95], $p_T^\nu > 25$ GeV in dependence of a minimum neutrino transverse momentum cut. The leftmost point corresponds to $p_T^\nu > 25.01$ GeV. Note the uncertainty of 10 pb on the results. The upper panel shows the absolute NNLO result and the lower panel the result normalized to the smallest displayed value of $\tau_{\text{cut}} = 0.001$ GeV. The black lines denote the results of 3207 ± 2 pb (FEWZ) and 3191 ± 7 pb (DYNNLO) from ref. [95], solid for the central values and dotted for the uncertainties.

is below one percent for $\tau_{\text{cut}} = 0.2$ GeV. Since this is still above our expectation from the benchmark cuts (cf. figure 10) it suggests that an asymmetry of at least a few GeV should be used in order to more easily obtain a precision prediction with this method.

The absolute position of our set of curves has an uncertainty of about 10 pb, corresponding to 3 per mille in relative terms, similar to the result reported from DYNNLO in ref. [95]. Within this uncertainty, our $\tau_{\text{cut}} = 0.001$ GeV result is compatible with the one from DYNNLO but slightly lower than the one from FEWZ. In order to press the

comparison further, say at the 1 per mille level, would require running at even lower values of τ_{cut} to better probe the asymptotic logarithmic dependence; however, we are already limited by the numerical precision that we are able to reasonably achieve.

Overall we find that it is computationally very expensive to obtain numerically precise values for fully symmetric cuts. In the presence of these cuts the power corrections are sizeable and force one to use tiny values of τ_{cut} , which amplify cancellations between contributions and challenge all parts of the numerical integration. In addition to this, with our binning size of 0.01 GeV and $\tau_{\text{cut}} = 0.001$ GeV, we already see a difference of ~ 10 pb when a cut is made at 25.01 GeV instead of 25.00 GeV. This indicates that for the NNLO coefficient the approach towards fully symmetric cuts is at least as steep as that at NLO. In light of these findings we believe that not just experimental studies should avoid the symmetric cuts, but also comparison benchmarks, since no precise comparisons at the per mille level are possible for calculations with intrinsic cutoff parameters. We note that all observations in this section transfer equivalently to Z production with symmetric cuts on the two leptons.

In figure 31 we show the NNLO positron transverse momentum distribution in W^+ production normalized to the NLO distribution for the *benchmark* cuts and parameters in ref. [95]. This is to be compared with the corresponding results in figure 17 of ref. [95].¹¹ The distributions for the neutrino as well as for the positron are virtually indistinguishable, so we only include the positron distribution here. Again, the jettiness cutoff effects become sizable towards fully symmetric cuts, but also in the region around $m_W/2$, which is again sensitive to soft radiation effects and problematic in fixed order perturbation theory.

Summary. In ref. [95] the authors of various codes, including FEWZ and DYNNLO, provide benchmark comparisons for W and Z production at NNLO accuracy. At NNLO the difference between FEWZ and DYNNLO inclusively is at the level of 0.5%, compatible within their respective numerical uncertainties. Considering that the experimental uncertainties are at that level, one now needs to aim for higher precision. Since the benchmark setup uses a choice of cuts that cause numerical instabilities, this is difficult to achieve in practice. We therefore advocate the use of a set of cuts that allows for one per mille numerical precision, so that numerical errors no longer constitute a substantial uncertainty, and the implementations and methods can be compared at a more useful level. After having established a common precision baseline, numerically difficult cuts could be studied.

To achieve a more useful comparison, we propose the use of cuts that are slightly asymmetric, differing by at least a few GeV. Although highly asymmetric cuts decrease the cross section and remove part of the phase space that is calculated at higher orders, the benefits of a small asymmetry are twofold. First, it removes a region of phase space that

¹¹We found that in ref. [95] the NLO-normalized NNLO neutrino and lepton distributions use opposite NLO-normalization factors. That is, we can fully reproduce their differential distributions for W^\pm when using lepton NLO distributions for the normalization of the NNLO neutrino distributions, and vice versa. For example in figure 20 the normalized l^+ transverse momentum distribution in the phase space region predominated by W +jets shows a 10% difference between the fixed order NNLO result and POWHEG+PYTHIA, while agreement is found with SHERPA NLO+PS. With the fixed normalization this situation is exactly reversed.

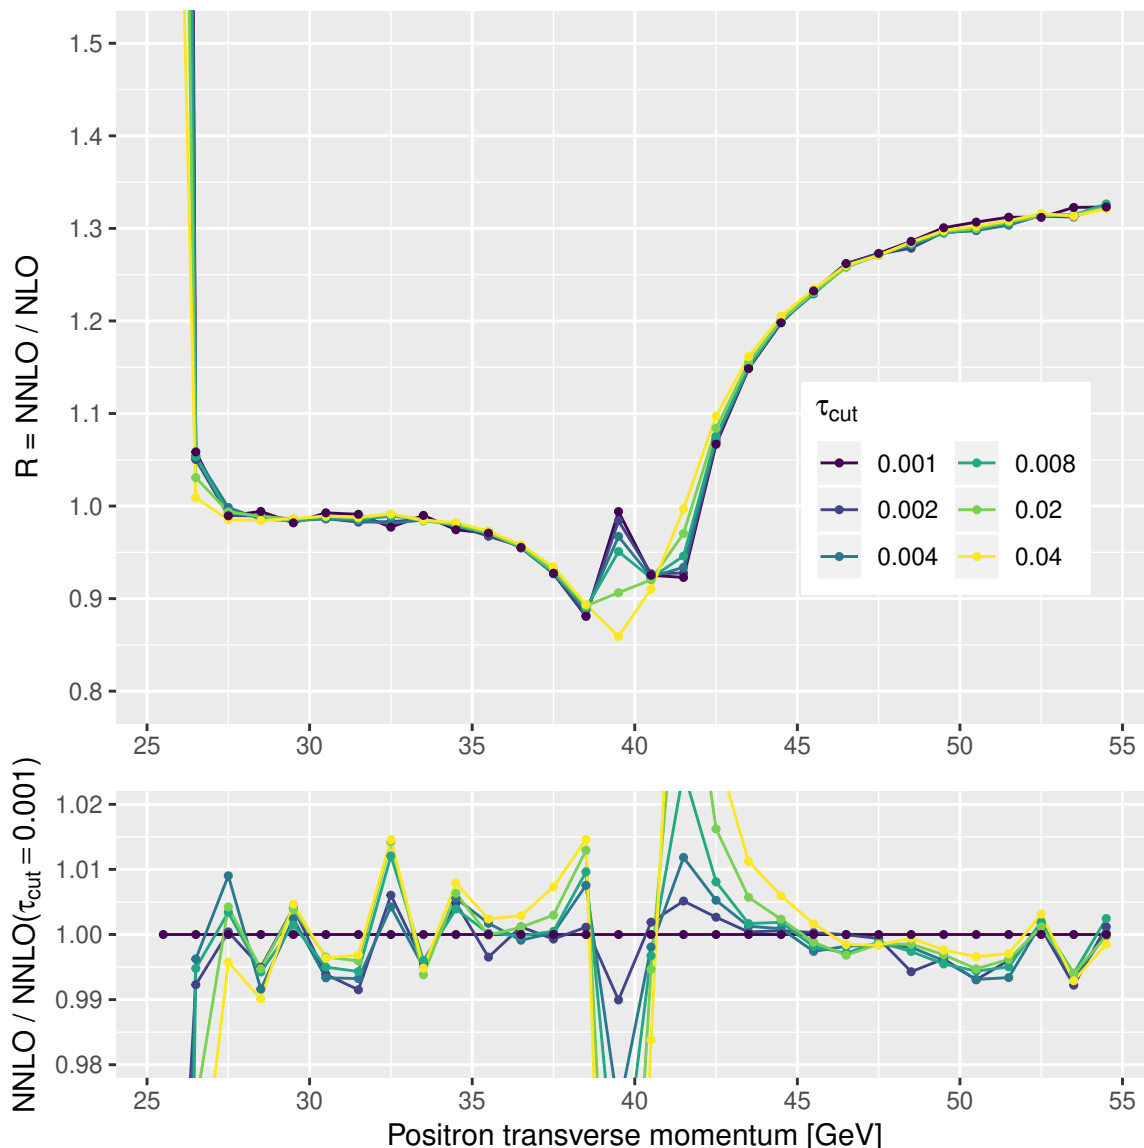


Figure 31. NNLO positron transverse momentum distribution in W^+ production. The upper panel shows the ratio to the NLO result. The lower panel shows the ratio to the smallest displayed τ_{cut} value of 0.001.

is problematic in a fixed order calculation and results in unphysical predictions (negative cross sections in the transverse momentum distribution). Second, it enables precise NNLO predictions with a variety of NNLO codes that in principle should agree.

7 Conclusions

With the onset of key hadron-collider measurements at the per mille level, interpretation of the results — and thus our understanding of nature — should be limited by uncertainties inherent in the theoretical predictions. Even with current higher-order predictions, that in some cases have percent-level scale uncertainties, control of numerical and methodological

errors at the per mille level is required to demonstrate their reliability. This then allows for the study of input parameters, and their impact, at the necessary level of precision. Unfortunately practical resource limitations, set by local workstations or even expensive computing clusters, are easily reached by precise calculations at NNLO. These limitations are even easier to saturate when including scans over additional parameters in the predictions. More often than not these limitations result in the introduction of errors, or the use of uncontrolled approximations, that may lead to a loss of the required precision. This can have a direct phenomenological impact that, for instance, can decide between the advent of a signal for new physics or the continued success of the SM.

In this paper we have addressed this issue by demonstrating that control at the per-mille level can be achieved with a new version of the code MCFM. A key component of the theoretical prediction is the numerical integration over the available phase space, where any type of technical cutoff or artifact must be able to be controlled below that level of precision. We first ensured that the raw numerical predictions can reach these levels of precision and that their errors are reliably estimated. This has been achieved through our newly implemented fully parallelized MPI+OMP Vegas integration that adaptively selects contributions with the largest uncertainties and is fully resumable through automatically written snapshots. Our approach allows reliable error estimates for precision predictions because it can use a huge number of calls per single integral estimate, i.e. per iteration. We have compared to a naïve parallelization, obtained by combining many independent low-statistics calls, and find that it requires statistical analysis methods to obtain trustworthy error estimates. For such a situation we recommend the use of the well-known bootstrap/jackknife technique.

A further consideration is that a number of today’s NNLO calculations, including those implemented in MCFM, depend on a slicing cutoff to regularize IR divergences (a jettiness cutoff, τ_{cut} , in this paper). Results can only be obtained as an extrapolation $\tau_{\text{cut}} \rightarrow 0$, otherwise residual finite τ_{cut} effects enter as a systematic error. To estimate the slicing cutoff uncertainties for a finite τ_{cut} , additional integrations must be performed for a range of τ_{cut} values and the dependence assessed. This is especially important differentially, where the residual dependence can be highly non-uniform and large compared to inclusive results. This extrapolation is extraordinarily computationally expensive since smaller values of τ_{cut} lead to larger numerical cancellation effects. Reaching either the required precision or a small enough τ_{cut} for these independent runs is not always possible. To address this we have implemented an automatic correlated sampling of multiple τ_{cut} values within one single integration run. This saves orders of magnitude in resources or, equivalently, improves results with equal resources by orders of magnitude. Furthermore we have implemented a boosted jettiness definition for all processes and included leading power corrections, which in combination lead to further order of magnitude performance improvements.

Taken together, these improvements ensure that the τ_{cut} dependence can be controlled at the per mille level on small computing clusters. To assess the reliability and give concrete error estimates, we have presented a detailed scheme for estimating the residual τ_{cut} error of NNLO results based on our automatic sampling of additional τ_{cut} values and their fitting. Since our fully differential fitting is based on the expected behavior in the asymptotic

regime, we can reliably exploit it for both improvements and error or reliability estimates. We have shown examples where the differential fit improves results by an order of magnitude and furthermore makes the differential τ_{cut} dependence uniform. We have also considered the case where the fit is no longer reported to be reliable, illustrating the identification of regions that need to be scrutinized further for a valid error estimate.

To avoid introducing the approximation of using NLO matrix elements for the calculation of PDF uncertainties, our new implementation allows the use of multiple PDF sets and PDF set members simultaneously at NNLO. They are computed simultaneously in a correlated way, saving many orders of magnitude of computational resources compared to uncorrelated integrations. We have used these improvements to study cases where lower-order matrix elements are used to approximate full NNLO PDF uncertainties and shown that our correlated implementation allows for per mille level comparisons between different PDF sets. Studies that discern the impact of different data sets and methods in the fits become directly tractable at NNLO at a high precision. We have demonstrated this feature for all NNLO processes in MCFM and also for the high- p_T tail of the Higgs boson transverse momentum spectrum at NLO. In addition we have used our code to attempt to reproduce results contained in the benchmarking exercise of ref. [95] but find that the cuts used there prohibit a comparison at the per-mille level.

Modern calculations at the level of NNLO QCD and beyond require the assembly and availability of dozens of components that each represents years of work. It is therefore mandatory that such components can be easily and reliably reused for further studies. MCFM provides a repository of such work, and has been used several times as a basis for fixed order NNLO and NLO SM calculations, resummed calculations, as well as implementations of physics beyond the SM. Its library of amplitudes has found use in dozens of projects and studies. Given this track record, the overhaul of all core components of the code described here is an important step to increase its usability and reliability, and to keep it an important tool for both experimentalists and theorists. In particular, the features and efficiency gains documented here will enable a public distribution of NNLO W^\pm , Z and H production processes in association with a jet in the near future.

Acknowledgments

This work was supported by the U.S. Department of Energy under award No. DE-SC0008347. This document was prepared using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359.

A Detailed description of new features

In this section we present the new and modified features in MCFM and describe how to use them on a technical level, complementing the new MCFM manual. With all re-implemented and newly implemented components we strive for Fortran 2008 compliance, making explicit use of its features. Following the Fortran standard furthermore allows us to achieve compatibility with not just the GNU compiler. In previous versions of MCFM the licensing was unclear, since none was specified. We now license all code under the GNU GPL 3 license.¹²

Improved input file mechanism. We have implemented a new input file mechanism based on the configuration file parser `config_fortran` [98]. This INI-like file format no longer depends on a strict ordering of configuration elements, allows easy access to configuration elements through a single global configuration object, and makes it easy to add new configuration options of scalar and array numerical and string types. Using the parser package also allows one to override or specify all configuration options as command line arguments to MCFM, for example running MCFM like `./mcfm_omp input.ini -general%nproc=200 -general%part=nlo`. This is useful for batch parameter run scripts. Settings can also be overridden with additional input files that specify just a subset of options.

New histogramming. We replaced the previous Fortran77 implementation of histograms, that used routines from 1988 by M. Mangano, with a new suite of routines. The new histogram implementation allows for any number of histograms with any number of bins, each of which is dynamically allocated. Furthermore, everything is also handled in a fully multi-threaded approach with the integration. For each OMP thread temporary histograms are allocated which are then reduced to a single one after each integration iteration, so that no OMP locks (critical regions) are required.

New Vegas integration, part-adaptive and resumable. The previous implementation of the Vegas routine was based on Numerical Recipes code. We have re-implemented Vegas and the surrounding integration routines. All parts of a NLO or NNLO calculation are now chosen adaptively based on the largest absolute numerical uncertainty. A precision goal can be set in the input file as well as a χ^2 /it goal and a precision goal for the warmup run. If the goals for the warmup are not reached, the warmup repeats with twice the number of calls. With the setting `writeintermediate` one can control whether histograms are written in intermediate stages during the integration. Enabling the setting `readin` allows one to resume the integration from any point from a previous run. Snapshots saving the whole integration state are saved automatically. When resuming, the only parameter that the user can safely officially change is the `precisiongoal`. Further tweak configuration options to control the stages of the integration have been introduced, which can provide benefits over the default settings in certain situations.

¹²See <https://www.gnu.org/licenses/gpl-3.0.en.html>.

The section `integration` in the configuration file allows for tweaks in the following way. The precision goal can be adjusted by setting `precisiongoal` to a relative precision that should be reached. Similarly, the settings `warmupprecisiongoal` and `warmupchisqgoal` control the minimum relative precision and χ^2/it for the warmup phase of `iterbatchwarmup` (default 5) iterations. If the warmup criterion fails, the number of calls is increased by a factor of two. The calls per iteration get increased by a factor of `callboost` (default 4) after the warmup. From then on the number of calls per iteration is increased by a factor of `itercallmult` (default 1.4) for a total of `iterbatch1` iterations. After these first `iterbatch1` iterations, the increase happens for every `iterbatch2` iterations. The setting `maxcallsperiter` controls the cap for the number of calls per iteration. The number of Vegas grid subdivisions can be controlled with `ndmx` (default 100).

The purpose of these settings is a fine control in certain situations. For example to compute expensive PDF uncertainties, one wants a relatively precise warmup run (where additional PDF sets are not sampled) and as few calls as necessary afterwards: for the plots in this paper we thus chose a relative warmup precision goal of 10%, and set `callboost` to 0.25. This means that the first `iterbatch1` iterations after the warmup run only with a quarter of the calls than during the warmup. This precision is sufficient to compute precise PDF uncertainties, when making use of the strong correlations as in MCFM-9.0. Any further iterations come in batches of `iterbatch2`, which we set to 1. It allows for a quick switching to parts of the NNLO cross section that have the largest uncertainty. For normal applications one wants to boost the number of calls after the warmup significantly, so a default value of `callboost=4` is chosen.

We provide default settings for the initial number of calls per iteration for all components of a NNLO calculation. They can be overridden with the following settings in the `integration` section: `initcallslord`, `initcallsnlovirt`, `initcallsnloreal`, `initcallsnlofrag` for parts of a NLO calculations, `initcallssnlobelow`, `initcallssnloabove` for parts of a SCET based NLO calculation, and `initcallsnnlobelow`, `initcallsnnlovirtabove`, as well as `initcallssnlorealabove` for the parts of the NNLO coefficient.

Low discrepancy sequence. MCFM-8.0 and prior relied on a linear congruential generator implementation from Numerical Recipes for the generation of a pseudo-random sequence. With newer versions the MT19937 implementation of the C++ standard library is used, and with this version of MCFM we include an implementation of the Sobol low discrepancy sequence based on the code `sobseq` [52] with initialization numbers from ref. [53]. The Sobol sequence is used by default and can be toggled using the flag `usesobol = .true.` in the `integration` section of the input file, see also section 3. When running in MPI mode, the number of nodes has to be a power of two for the Sobol sequence, because we use it in a strided manner. Otherwise the code will automatically fall back to using the MT19937 sequence with seed value `seed` in the integration section of the input file. A `seed` value of 0 denotes a randomly initialized seed.

Fully parallelized OMP+MPI use of LHAPDF. In previous versions of MCFM calls to LHAPDF were forced to access from only a single OMP thread through a lock. This is because the interface was based on the old LHAGlue interface, part of LHAPDF. We have written an interface to LHAPDF from scratch based on the new object oriented treatment of PDFs in LHAPDF 6. For each OMP thread we initialize a copy of the used PDF members which can be called fully concurrently. The amount of PDF sets with or without PDF uncertainties is only limited by the available system memory. The memory usage of MCFM can then range from roughly 20MB when only one central PDF grid is being used, to ~ 7.4 GB when 32 OMP threads fully load all members of the PDF sets CT14nnlo, MMHT2014nnlo68c1, ABMP16als118.5_nnlo, NNPDF30_nnlo_as_0118, NNPDF31_nnlo_as_0118 and PDF4LHC15_nnlo_30 for PDF uncertainties. The total number of members for these grids is 371, each loaded for every of the 32 OMP threads.

Since each OMP thread allocates its own copy of PDF members and histograms we have no need to introduce any OMP locks. On the other hand the memory usage increases and one runs into being CPU cache or DRAM bandwidth bound earlier. In practice, we find that this is still faster than having OMP locks, which directly decrease the speedup in the spirit of Amdahl's law. Ideally the LHAPDF library should be improved to allow for thread-safe calls with just one memory allocation.

Histograms for additional values of τ_{cut} , μ_R , μ_F and multiple PDFs. When using the automatic scale variation, in addition to the normal histograms, additional histograms with filenames `_scale_XY_` are generated, where X is a placeholder for the renormalization scale variation and Y for the factorization scale variation. X and Y can either be u for an upwards variation by a factor of two, d for a downwards variation by a factor of two, or just - if no change of that scale was made. The envelope of maximum and minimum can then easily be obtained.

For the sampling of additional values of τ_{cut} for NLO and NNLO calculations using jet-tiness subtractions, additional histograms with filenames `_taucut_XXX_` are written. Here XXX is a placeholder for the chosen τ_{cut} values in the optional array `taucutarray`, if specified, or one of the five automatically chosen values. These additional files only contain the *differences* to the nominal choice of τ_{cut} , so that $\Delta\sigma(\tau_{\text{cut,nominal}}) - \Delta\sigma(\tau_{\text{cut},i})$ is stored. If `taucutarray` has not been specified, the automatic choice of additional τ_{cut} values is enabled based on the default nominal τ_{cut} for the process or the users choice of the nominal τ_{cut} value as specified in `taucut`. In addition a file with `_taucutfit_` is generated, which in addition to the fitted corrections and its uncertainty includes columns for the maximum relative integration uncertainty for the additionally sampled τ_{cut} values and the reduced χ^2 of the fit. With the procedure in section 4, the fit together with the individual τ_{cut} histograms allows the user to assess the systematic τ_{cut} error and possibly improve results.

When multiple PDF sets are chosen, additional files with the names of the PDF sets are generated. In case PDF uncertainties are enabled, the histograms also include the upper and lower bounds of the PDF uncertainties.

User cuts, histograms and re-weighting. Modifying the plotting routines in the files `src/User/nplotter*.f` allows for modification of the pre-defined histograms and addition of any number of arbitrary observables. The routine `gencuts_user` can be adjusted in the file `src/User/gencuts_user.f90` for additional cuts after the jet algorithm has performed the clustering. In the same file the routine `reweight_user` can be modified to include a manual re-weighting for all integral contributions. This can be used to obtain improved uncertainties in, for example, tails of distributions. One example is included in the subdirectory `examples`, where the `reweight_user` function approximately flattens the Higgs transverse momentum distribution, leading to equal relative uncertainties even in the tail at 1 TeV.

A.1 Compatibility with the Intel compiler and benchmarks

Previous versions of MCFM were developed using `gfortran` as a compiler. MCFM contained code that did not follow a specific Fortran standard, and was only compatible with using `gfortran`. We fixed code that did not compile or work with the recent Intel Fortran compiler `ifort` 19.0.1. This does not mean that we claim to be strictly standards compliant with a specific Fortran version, but we aim to be compliant with Fortran 2008. We now fully support GCC versions newer than 7 and Intel compilers newer than 19. There might still be compatibility issues with other Fortran compilers, but we are happy to receive bug reports for any issues regarding compilation, that are not due to a lack of modern Fortran 2008 features. To use the Intel compiler one has to change the `USEINTEL` flag in the files `Install` and `makefile` to `YES`.

To see whether MCFM can make use of potential Intel compiler improvements over the GNU compiler collection (GCC) we benchmarked the double real emission component of Higgs production at NNLO. We perform tests on our cluster with Intel Xeon 64-bit X5650 2.67 GHz Westmere CPUs, where two six-core CPUs are run in a dual-socket mode with a total of twelve cores. Similarly, we have an AMD 6128 HE Opteron 2GHz quad-socket eight-core setup, thus each having 32 cores per node.

We benchmark both the Intel and GCC compilers on both the Intel and AMD systems. On the Intel system we use 16 MPI processes each with 12 OMP threads, and on the AMD system we have 8 MPI processes using 32 OMP threads. With this we have the same total clockrate of 512 GHz for each setup. For all benchmarks we find that the scaling is perfect up to this size, that is if we use half the number of MPI or OMP threads we double our run-time.

We first try both the Intel fortran compiler 19.0.1 and GCC 9.1.0 on the Intel system with the highest generic optimization flags `-O3 -xsse4.2` and `-O3 -march=westmere`, respectively. Furthermore, we lower the optimizations to `-O2` each and remove the processor specific optimization flags `-xsse4.2` and `-march=westmere`, respectively. All our benchmark run-times in the following are consistent within ± 0.5 s.

We do not support enabling unsafe math operations with `-ffast-math`, since the code relies on the knowledge of NaN values and checks on those. Such checks would be skipped with the meta flag `-ffast-math` which sets `-ffinite-math-only`.

Compiler/flags	wall time \pm 0.5s
ifort -O3 -xsse4.2	90s
ifort -O2 -xsse4.2	86s
ifort -O2	90s
ifort -O1	103s
gfortran -O3 -march=westmere	101s
gfortran -O2 -march=westmere	105s
gfortran -O2	105s
gfortran -O1	110s

Table 2. Benchmark results on the Intel system with $10 \cdot 25\text{M}$ calls distributed over 16 MPI processes, each using 12 OMP threads. The GCC version is 9.1.0 and the Intel Fortran compiler 19.0.1.

The benchmark results in table 2 show that using the Intel compiler, performance benefits of $\simeq 10 - 20\%$ can be achieved. Our goal here is not to go beyond this and check whether exactly equivalent optimization flags have been used in both cases. Enabling optimizations beyond `-O2` have little impact, but come with a penalty for the Intel compiler and with a slight benefit for gfortran. We also notice that processor specific optimizations play no significant role. This might also be in part due to the fact that MCFM does not offer much space for (automatic) vectorization optimizations. To summarize, the default optimization flags of `-O2` should be sufficient in most cases. We do not expect that the conclusions from these benchmarks change for different processes. On the other hand if computing PDF uncertainties, the majority of time is used by LHAPDF and different optimization flags for LHAPDF might play a role then. We performed the same benchmark with an older version of GCC, version 7.1.0 using `-O2` optimizations, and found that the run-times are the same as for the newer version.

Finally, we performed some benchmarks on our AMD setup and found that it is $\simeq 2.5$ times slower for the same total clockrate. Using the Intel compiler for the AMD setup decreased the performance by another $\simeq 30\%$. This is likely due to the fact that the Intel compiler already optimizes for the general Intel architecture.

These benchmarks try to give a general impression and might depend in detail on the process, the number of histograms and whether to compute PDF uncertainties, for example. Especially when computing PDF uncertainties the perfect scaling we tested here might break down since the computation can become memory bound. We discuss this caveat in more detail in appendix A.2.

A.2 Remarks on memory bound performance issues

To get numerically precise predictions at the per mille level for NNLO cross sections, already hundreds of million of calls are necessary. Obtaining PDF uncertainties using those NNLO matrix elements significantly increases the computational time. In a simplified

view the total computational time composes as $N_{\text{calls}} * (T + N_{\text{PDF}} \cdot T_{\text{PDF}})$, where T is the computational effort for the matrix element piece, and the PDF part is proportional to the time calling the PDF evolution N_{PDF} times and code related to performing the convolutions. For tree level matrix element evaluations, usually also $T \lll T_{\text{PDF}}$ holds, so the computational cost grows linearly with the number of PDFs.

This naive picture breaks down in practice when a lot of PDFs are sampled together with a lot of histograms or histogram bins. The total memory necessary to store all the histogram information grows like $N_{\text{PDF}} \cdot N_{\text{bins}} \cdot N_{\text{thr.}}$, where N_{PDF} is the number of PDF members, N_{bins} the number of histogram bins summed over all histograms and $N_{\text{thr.}}$ is the number of OMP threads. The factor $N_{\text{thr.}}$ enters since we have thread-local storage to avoid OMP locks. The values are stored in double precision, so the total memory used is $N_{\text{PDF}} \cdot N_{\text{bins}} \cdot N_{\text{thr.}} \cdot 8$ bytes.

Assuming for example, 300 PDF members, 10 histograms with each 20 bins and 12 threads, this sums up to 720 kb of memory. For the virtual corrections and LO pieces, one has to update this amount of memory once for each call. For the real emission matrix elements one has to accumulate all dipole contributions, so this number additionally scales with the number of dipole contributions. All the histogram updates are usually fully vectorized for modern superscalar processors with SSE and/or AVX extensions. But if this used memory is too large and does not easily fit into the CPU core caches anymore, a transfer to and from DRAM happens, which now is the limiting factor and significantly slows down the computation. Because for that reason, one should work with a minimal number of necessary histograms when working with a lot of PDF members. This is especially important for cluster setups that are not optimized towards memory bound applications, non-NUMA systems. For example in our cluster we have relatively old AMD Opteron quad-socket eight-core nodes with little CPU cache, and with above numbers we are already limited in wall-time improvements with using ~ 16 cores. Then reducing the number of histograms will *significantly* improve the performance. In principle one can reduce the histogram precision to single precision and cut memory transfer and storage in half, while doubling the computational speed. This might lead to problems with accumulated rounding errors though, and we have not investigated this further, since in practice one can sufficiently limit the number of histograms or PDF sets.

B Supporting plots for the jackknife-after-bootstrap procedure

In figure 32 and figure 33 the lower panel displays a visualization of the jackknife after bootstrap technique. Each point represents one of the N data points that is being left out. The dashed lines represent quantiles of the bootstrap distribution. figure 32 shows that leaving out the single point number 557 would significantly shrink the percentiles and make the Gaussian distribution symmetric. After removing the outlier figure 33 is obtained, where now no single point would significantly modify the bootstrap distribution. For more details we refer to [61].

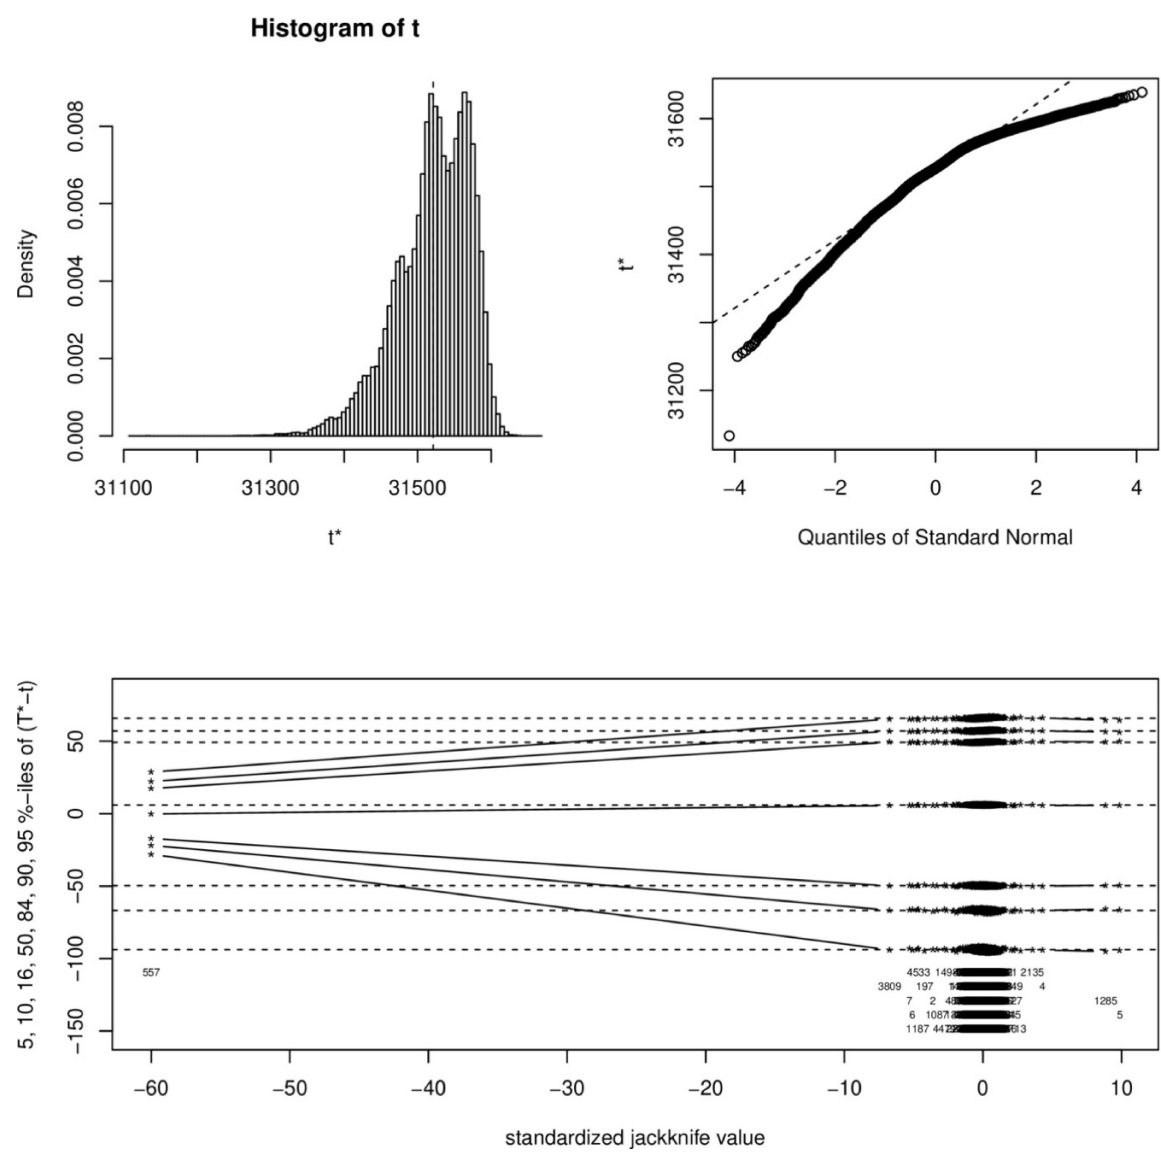


Figure 32. Result of applying the bootstrap technique to our MT19937 data set of about 4500 data points with 10 million calls each. The sample size is not Gaussian due to one significant outlier.

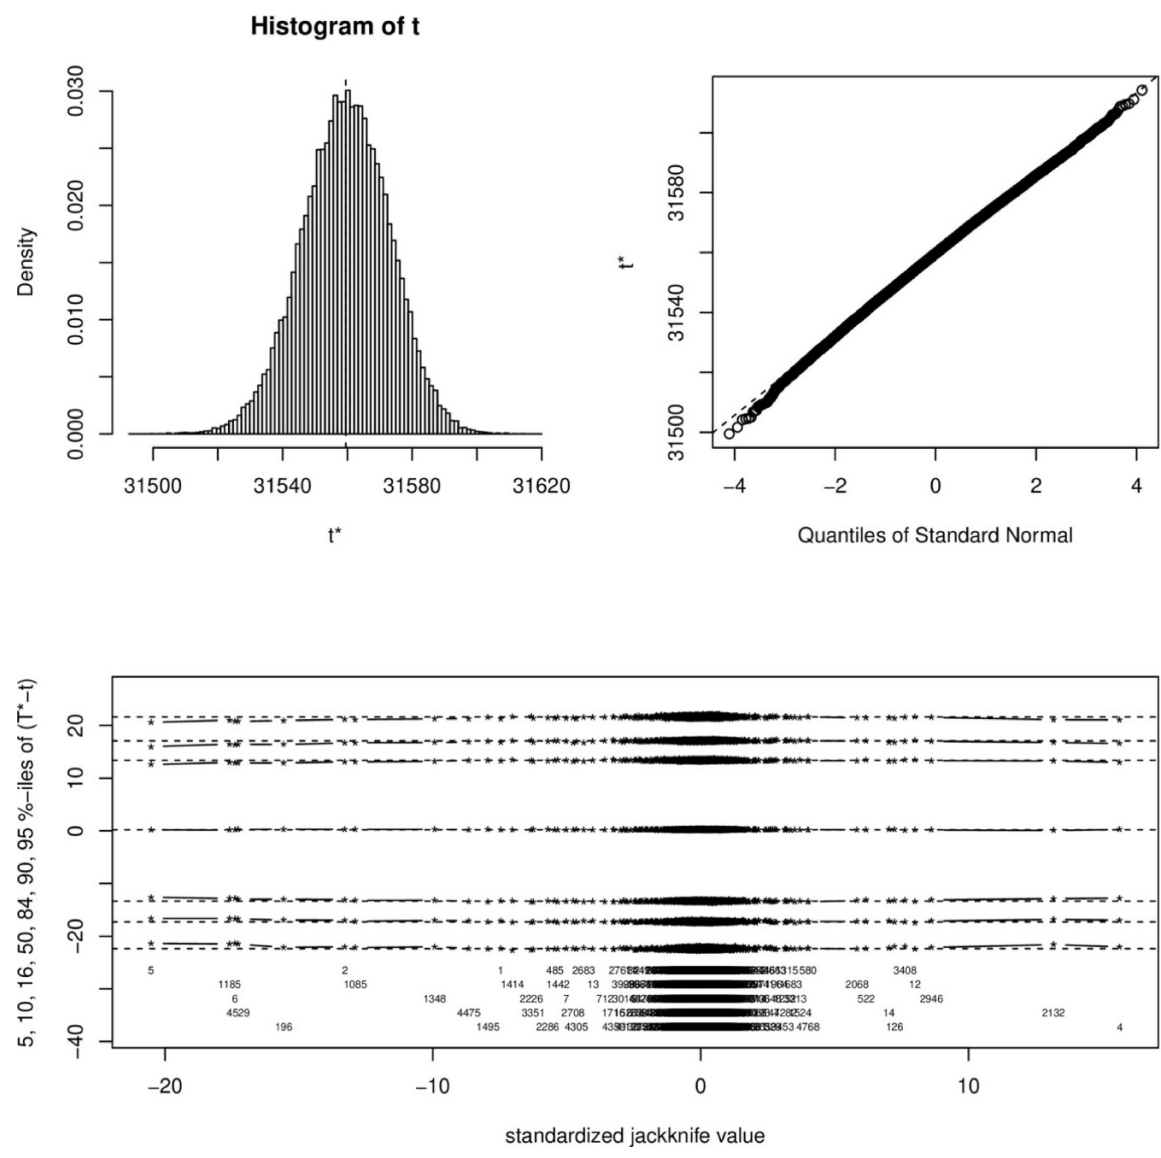


Figure 33. Result of the applying the bootstrap technique to our MT19937 data set of about 4500 data points with 10 million calls each. The worst four outliers as shown in the jackknife plot in figure 32 have been removed. The result is 31559 ± 13 .

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] S. Catani, L. Cieri, G. Ferrera, D. de Florian and M. Grazzini, *Vector boson production at hadron colliders: a fully exclusive QCD calculation at NNLO*, *Phys. Rev. Lett.* **103** (2009) 082001 [[arXiv:0903.2120](https://arxiv.org/abs/0903.2120)] [[INSPIRE](#)].
- [2] S. Catani and M. Grazzini, *An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC*, *Phys. Rev. Lett.* **98** (2007) 222002 [[hep-ph/0703012](https://arxiv.org/abs/hep-ph/0703012)] [[INSPIRE](#)].
- [3] M. Grazzini, *NNLO predictions for the Higgs boson signal in the $H \rightarrow WW \rightarrow l\nu l\nu$ and $H \rightarrow ZZ \rightarrow 4l$ decay channels*, *JHEP* **02** (2008) 043 [[arXiv:0801.3232](https://arxiv.org/abs/0801.3232)] [[INSPIRE](#)].
- [4] M. Grazzini and H. Sargsyan, *Heavy-quark mass effects in Higgs boson production at the LHC*, *JHEP* **09** (2013) 129 [[arXiv:1306.4581](https://arxiv.org/abs/1306.4581)] [[INSPIRE](#)].
- [5] R.V. Harlander, S. Liebler and H. Mantler, *SusHi: A program for the calculation of Higgs production in gluon fusion and bottom-quark annihilation in the Standard Model and the MSSM*, *Comput. Phys. Commun.* **184** (2013) 1605 [[arXiv:1212.3249](https://arxiv.org/abs/1212.3249)] [[INSPIRE](#)].
- [6] R.V. Harlander, S. Liebler and H. Mantler, *SusHi Bento: Beyond NNLO and the heavy-top limit*, *Comput. Phys. Commun.* **212** (2017) 239 [[arXiv:1605.03190](https://arxiv.org/abs/1605.03190)] [[INSPIRE](#)].
- [7] M. Bonvini, R.D. Ball, S. Forte, S. Marzani and G. Ridolfi, *Updated Higgs cross section at approximate N^3LO* , *J. Phys. G* **41** (2014) 095002 [[arXiv:1404.3204](https://arxiv.org/abs/1404.3204)] [[INSPIRE](#)].
- [8] C. Anastasiou, K. Melnikov and F. Petriello, *Fully differential Higgs boson production and the di-photon signal through next-to-next-to-leading order*, *Nucl. Phys. B* **724** (2005) 197 [[hep-ph/0501130](https://arxiv.org/abs/hep-ph/0501130)] [[INSPIRE](#)].
- [9] C. Anastasiou, S. Bucherer and Z. Kunszt, *HPro: A NLO Monte-Carlo for Higgs production via gluon fusion with finite heavy quark masses*, *JHEP* **10** (2009) 068 [[arXiv:0907.2362](https://arxiv.org/abs/0907.2362)] [[INSPIRE](#)].
- [10] M. Cacciari, F.A. Dreyer, A. Karlberg, G.P. Salam and G. Zanderighi, *Fully Differential Vector-Boson-Fusion Higgs Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **115** (2015) 082002 [*Erratum ibid.* **120** (2018) 139901] [[arXiv:1506.02660](https://arxiv.org/abs/1506.02660)] [[INSPIRE](#)].
- [11] F.A. Dreyer and A. Karlberg, *Fully differential Vector-Boson Fusion Higgs Pair Production at Next-to-Next-to-Leading Order*, *Phys. Rev. D* **99** (2019) 074028 [[arXiv:1811.07918](https://arxiv.org/abs/1811.07918)] [[INSPIRE](#)].
- [12] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Diphoton production at hadron colliders: a fully-differential QCD calculation at NNLO*, *Phys. Rev. Lett.* **108** (2012) 072001 [*Erratum ibid.* **117** (2016) 089901] [[arXiv:1110.2375](https://arxiv.org/abs/1110.2375)] [[INSPIRE](#)].
- [13] R. Gavin, Y. Li, F. Petriello and S. Quackenbush, *FEWZ 2.0: A code for hadronic Z production at next-to-next-to-leading order*, *Comput. Phys. Commun.* **182** (2011) 2388 [[arXiv:1011.3540](https://arxiv.org/abs/1011.3540)] [[INSPIRE](#)].
- [14] R. Gavin, Y. Li, F. Petriello and S. Quackenbush, *W Physics at the LHC with FEWZ 2.1*, *Comput. Phys. Commun.* **184** (2013) 208 [[arXiv:1201.5896](https://arxiv.org/abs/1201.5896)] [[INSPIRE](#)].

- [15] Y. Li and F. Petriello, *Combining QCD and electroweak corrections to dilepton production in FEWZ*, *Phys. Rev. D* **86** (2012) 094034 [[arXiv:1208.5967](#)] [[INSPIRE](#)].
- [16] S. Alioli, C.W. Bauer, C. Berggren, F.J. Tackmann and J.R. Walsh, *Drell-Yan production at NNLL'+NNLO matched to parton showers*, *Phys. Rev. D* **92** (2015) 094020 [[arXiv:1508.01475](#)] [[INSPIRE](#)].
- [17] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, *Eur. Phys. J. C* **78** (2018) 537 [[arXiv:1711.06631](#)] [[INSPIRE](#)].
- [18] J. Bellm et al., *Jet Cross Sections at the LHC and the Quest for Higher Precision*, [arXiv:1903.12563](#) [[INSPIRE](#)].
- [19] J.M. Campbell and R.K. Ellis, *An Update on vector boson pair production at hadron colliders*, *Phys. Rev. D* **60** (1999) 113006 [[hep-ph/9905386](#)] [[INSPIRE](#)].
- [20] J.M. Campbell, R.K. Ellis and W.T. Giele, *A Multi-Threaded Version of MCFM*, *Eur. Phys. J. C* **75** (2015) 246 [[arXiv:1503.06182](#)] [[INSPIRE](#)].
- [21] R. Boughezal et al., *Color singlet production at NNLO in MCFM*, *Eur. Phys. J. C* **77** (2017) 7 [[arXiv:1605.08011](#)] [[INSPIRE](#)].
- [22] J.M. Campbell, D. Wackerroth and J. Zhou, *Study of weak corrections to Drell-Yan, top-quark pair and dijet production at high energies with MCFM*, *Phys. Rev. D* **94** (2016) 093009 [[arXiv:1608.03356](#)] [[INSPIRE](#)].
- [23] T. Neumann and Z.E. Sullivan, *Off-Shell Single-Top-Quark Production in the Standard Model Effective Field Theory*, *JHEP* **06** (2019) 022 [[arXiv:1903.11023](#)] [[INSPIRE](#)].
- [24] S. Catani, D. de Florian, G. Ferrera and M. Grazzini, *Vector boson production at hadron colliders: transverse-momentum resummation and leptonic decay*, *JHEP* **12** (2015) 047 [[arXiv:1507.06937](#)] [[INSPIRE](#)].
- [25] G. Bozzi, S. Catani, G. Ferrera, D. de Florian and M. Grazzini, *Production of Drell-Yan lepton pairs in hadron collisions: Transverse-momentum resummation at next-to-next-to-leading logarithmic accuracy*, *Phys. Lett. B* **696** (2011) 207 [[arXiv:1007.2351](#)] [[INSPIRE](#)].
- [26] D. de Florian, G. Ferrera, M. Grazzini and D. Tommasini, *Higgs boson production at the LHC: transverse momentum resummation effects in the $H \rightarrow 2\gamma$, $H \rightarrow WW \rightarrow \nu\nu$ and $H \rightarrow ZZ \rightarrow 4l$ decay modes*, *JHEP* **06** (2012) 132 [[arXiv:1203.6321](#)] [[INSPIRE](#)].
- [27] L. Arpino, A. Banfi, S. Jäger and N. Kauer, *BSM WW production with a jet veto*, *JHEP* **08** (2019) 076 [[arXiv:1905.06646](#)] [[INSPIRE](#)].
- [28] T. Carli et al., *A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: The APPLGRID Project*, *Eur. Phys. J. C* **66** (2010) 503 [[arXiv:0911.2985](#)] [[INSPIRE](#)].
- [29] J. Gao, C.S. Li and H.X. Zhu, *Top Quark Decay at Next-to-Next-to Leading Order in QCD*, *Phys. Rev. Lett.* **110** (2013) 042001 [[arXiv:1210.2808](#)] [[INSPIRE](#)].
- [30] T. Melia, K. Melnikov, R. Rontsch, M. Schulze and G. Zanderighi, *Gluon fusion contribution to $W+W-$ + jet production*, *JHEP* **08** (2012) 115 [[arXiv:1205.6987](#)] [[INSPIRE](#)].
- [31] I. Anderson et al., *Constraining Anomalous HVV Interactions at Proton and Lepton Colliders*, *Phys. Rev. D* **89** (2014) 035007 [[arXiv:1309.4819](#)] [[INSPIRE](#)].
- [32] G.P. Lepage, *Vegas: an adaptive multidimensional integration program*, (1980) [[INSPIRE](#)].

- [33] R.K. Ellis and G. Zanderighi, *Scalar one-loop integrals for QCD*, *JHEP* **02** (2008) 002 [[arXiv:0712.1851](#)] [[INSPIRE](#)].
- [34] S. Carrazza, R.K. Ellis and G. Zanderighi, *QCDLoop: a comprehensive framework for one-loop scalar integrals*, *Comput. Phys. Commun.* **209** (2016) 134 [[arXiv:1605.03181](#)] [[INSPIRE](#)].
- [35] T. Gehrmann and E. Remiddi, *Numerical evaluation of two-dimensional harmonic polylogarithms*, *Comput. Phys. Commun.* **144** (2002) 200 [[hep-ph/0111255](#)] [[INSPIRE](#)].
- [36] A. Buckley et al., *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J. C* **75** (2015) 132 [[arXiv:1412.7420](#)] [[INSPIRE](#)].
- [37] G.P. Lepage, *A New Algorithm for Adaptive Multidimensional Integration*, *J. Comput. Phys.* **27** (1978) 192 [[INSPIRE](#)].
- [38] P. Bratley and B. Fox, *Algorithm 659: Implementing Sobol's Quasirandom Sequence Generator*, *ACM Trans. Math. Software* **14** (1988) 88.
- [39] B. Fox, *Algorithm 647: Implementation and Relative Efficiency of Quasirandom Sequence Generators*, *ACM Trans. Math. Software* **12** (1986) 362.
- [40] Antonov and Saleev, *An economic method of computing LP_τ -sequences*, *USSR Comput. Math. Math. Phys.* **19** (1980) 252.
- [41] I. Sobol, *Uniformly distributed sequences with an additional uniform property*, *USSR Comput. Math. Math. Phys.* **16** (1977) 236.
- [42] I.M. Sobol and Y.L. Levitan, *The Production of Points Uniformly Distributed in a Multidimensional Cube* (in Russian), *Preprint IPM Akad. Nauk SSSR* (1976).
- [43] J.M. Campbell, R.K. Ellis and C. Williams, *Driving missing data at the LHC: NNLO predictions for the ratio of $\gamma + j$ and $Z + j$* , *Phys. Rev. D* **96** (2017) 014037 [[arXiv:1703.10109](#)] [[INSPIRE](#)].
- [44] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and D.M. Walker, *Vector Boson Production in Association with a Jet at Forward Rapidities*, *Eur. Phys. J. C* **79** (2019) 526 [[arXiv:1901.11041](#)] [[INSPIRE](#)].
- [45] T. Kluge, K. Rabbertz and M. Wobisch, *FastNLO: Fast pQCD calculations for PDF fits*, in *Deep inelastic scattering. Proceedings, 14th International Workshop, DIS 2006*, Tsukuba, Japan, 20–24 April 2006, pp. 483–486 (2006) [[DOI:10.1142/9789812706706_0110](#)] [[hep-ph/0609285](#)] [[INSPIRE](#)].
- [46] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, *J. Phys. G* **43** (2016) 023001 [[arXiv:1510.03865](#)] [[INSPIRE](#)].
- [47] A. Accardi et al., *A Critical Appraisal and Evaluation of Modern PDFs*, *Eur. Phys. J. C* **76** (2016) 471 [[arXiv:1603.08906](#)] [[INSPIRE](#)].
- [48] NNPDF collaboration, *A first determination of parton distributions with theoretical uncertainties*, *Eur. Phys. J. C* (2019) 79:838 [[arXiv:1905.04311](#)] [[INSPIRE](#)].
- [49] T. Hahn, *CUBA: A Library for multidimensional numerical integration*, *Comput. Phys. Commun.* **168** (2005) 78 [[hep-ph/0404043](#)] [[INSPIRE](#)].
- [50] J.F. Koksma, *A general theorem from the theory of uniform distribution modulo 1*, *Mathematica, Zutphen. B.* **11** (1942) 7.

- [51] E. Hlawka, *Funktionen von beschränkter Variation in der Theorie der Gleichverteilung*, *Ann. Mat. Pura Appl.* **54** (1961) 325.
- [52] D. van Vugt and K. Beljaars, *Modern fortran implementation of a Sobol sequence*, <https://github.com/Exteris/sobseq> (2016).
- [53] S. Joe and F.Y. Kuo, *Sobol sequence direction numbers*, <https://web.maths.unsw.edu.au/~fkuo/sobol/> (2010).
- [54] S. Joe and F. Kuo, *Constructing sobol sequences with better two-dimensional projections*, *SIAM J. Sci. Comput.* **30** (2008) 2635.
- [55] S. Joe and F.Y. Kuo, *Remark on algorithm 659: Implementing sobol's quasirandom sequence generator*, *ACM Trans. Math. Softw.* **29** (2003) 49.
- [56] Z. Nagy and Z. Trócsányi, *Next-to-leading order calculation of four jet observables in electron positron annihilation*, *Phys. Rev. D* **59** (1999) 014020 [Erratum *ibid.* **D 62** (2000) 099902] [[hep-ph/9806317](#)] [[INSPIRE](#)].
- [57] Z. Nagy, *Next-to-leading order calculation of three jet observables in hadron hadron collision*, *Phys. Rev. D* **68** (2003) 094002 [[hep-ph/0307268](#)] [[INSPIRE](#)].
- [58] A. Gehrmann-De Ridder, T. Gehrmann, N. Glover, A. Huss and T.A. Morgan, *NNLO QCD corrections for Z boson plus jet production*, *PoS(RADCOR2015)075* (2016) [[arXiv:1601.04569](#)] [[INSPIRE](#)].
- [59] A. Canty and B. Ripley, *Functions and datasets for bootstrapping from the book "Bootstrap Methods and Their Application"*, <https://cran.r-project.org/web/packages/boot/index.html> (2019).
- [60] A.C. Davison and D.V. Hinkley, *Bootstrap Methods and Their Application*, Cambridge University Press (1997).
- [61] B. Efron, *Jackknife-after-bootstrap standard errors and influence functions (with Discussion)*, *J. Roy. Statist. Soc. B* (1992) 83.
- [62] R. Boughezal, C. Focke, X. Liu and F. Petriello, *W-boson production in association with a jet at next-to-next-to-leading order in perturbative QCD*, *Phys. Rev. Lett.* **115** (2015) 062002 [[arXiv:1504.02131](#)] [[INSPIRE](#)].
- [63] J. Gaunt, M. Stahlhofen, F.J. Tackmann and J.R. Walsh, *N-jettiness Subtractions for NNLO QCD Calculations*, *JHEP* **09** (2015) 058 [[arXiv:1505.04794](#)] [[INSPIRE](#)].
- [64] I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *N-Jettiness: An Inclusive Event Shape to Veto Jets*, *Phys. Rev. Lett.* **105** (2010) 092002 [[arXiv:1004.2489](#)] [[INSPIRE](#)].
- [65] I. Moulton, L. Rothen, I.W. Stewart, F.J. Tackmann and H.X. Zhu, *Subleading Power Corrections for N-Jettiness Subtractions*, *Phys. Rev. D* **95** (2017) 074023 [[arXiv:1612.00450](#)] [[INSPIRE](#)].
- [66] J.M. Campbell, T. Neumann and C. Williams, *Z γ Production at NNLO Including Anomalous Couplings*, *JHEP* **11** (2017) 150 [[arXiv:1708.02925](#)] [[INSPIRE](#)].
- [67] M.A. Ebert, I. Moulton, I.W. Stewart, F.J. Tackmann, G. Vita and H.X. Zhu, *Power Corrections for N-Jettiness Subtractions at $\mathcal{O}(\alpha_s)$* , *JHEP* **12** (2018) 084 [[arXiv:1807.10764](#)] [[INSPIRE](#)].
- [68] J.M. Campbell, R.K. Ellis, Y. Li and C. Williams, *Predictions for diphoton production at the LHC through NNLO in QCD*, *JHEP* **07** (2016) 148 [[arXiv:1603.02663](#)] [[INSPIRE](#)].

- [69] J. Burkardt, *MINPACK, a Fortran90 implementation*, https://people.sc.fsu.edu/~jburkardt/f_src/minpack/minpack.html (2010).
- [70] J. More, B. Garbow and K. Hillstrome, *User Guide for MINPACK-1*, Technical Report ANL-80-74, Argonne National Laboratory (1980).
- [71] J. More, D. Sorenson, B. Garbow and K. Hillstrome, *The MINPACK Project*, in *Sources and Development of Mathematical Software*, Prentice-Hall (1984).
- [72] S. Catani and M.H. Seymour, *A General algorithm for calculating jet cross-sections in NLO QCD*, *Nucl. Phys. B* **485** (1997) 291 [Erratum *ibid.* **B 510** (1998) 503] [[hep-ph/9605323](#)] [[INSPIRE](#)].
- [73] C. Anastasiou et al., *High precision determination of the gluon fusion Higgs boson cross-section at the LHC*, *JHEP* **05** (2016) 058 [[arXiv:1602.00695](#)] [[INSPIRE](#)].
- [74] LHCb collaboration, *Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **01** (2016) 155 [[arXiv:1511.08039](#)] [[INSPIRE](#)].
- [75] LHCb collaboration, *Measurement of forward $W \rightarrow e\nu$ production in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **10** (2016) 030 [[arXiv:1608.01484](#)] [[INSPIRE](#)].
- [76] LHCb collaboration, *Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **09** (2016) 136 [[arXiv:1607.06495](#)] [[INSPIRE](#)].
- [77] LHCb collaboration, *Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **08** (2015) 039 [[arXiv:1505.07024](#)] [[INSPIRE](#)].
- [78] S. Alekhin, J. Blümlein, S. Moch and R. Placakyte, *Parton distribution functions, α_s and heavy-quark masses for LHC Run II*, *Phys. Rev. D* **96** (2017) 014011 [[arXiv:1701.05838](#)] [[INSPIRE](#)].
- [79] S. Dulat et al., *New parton distribution functions from a global analysis of quantum chromodynamics*, *Phys. Rev. D* **93** (2016) 033006 [[arXiv:1506.07443](#)] [[INSPIRE](#)].
- [80] L.A. Harland-Lang, A.D. Martin, P. Motylinski and R.S. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J. C* **75** (2015) 204 [[arXiv:1412.3989](#)] [[INSPIRE](#)].
- [81] NNPDF collaboration, *Parton distributions for the LHC Run II*, *JHEP* **04** (2015) 040 [[arXiv:1410.8849](#)] [[INSPIRE](#)].
- [82] NNPDF collaboration, *Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 663 [[arXiv:1706.00428](#)] [[INSPIRE](#)].
- [83] CMS collaboration, *Inclusive search for a highly boosted Higgs boson decaying to a bottom quark-antiquark pair*, *Phys. Rev. Lett.* **120** (2018) 071802 [[arXiv:1709.05543](#)] [[INSPIRE](#)].
- [84] R. Boughezal, F. Caola, K. Melnikov, F. Petriello and M. Schulze, *Higgs boson production in association with a jet at next-to-next-to-leading order in perturbative QCD*, *JHEP* **06** (2013) 072 [[arXiv:1302.6216](#)] [[INSPIRE](#)].
- [85] X. Chen, T. Gehrmann, E.W.N. Glover and M. Jaquier, *Precise QCD predictions for the production of Higgs + jet final states*, *Phys. Lett. B* **740** (2015) 147 [[arXiv:1408.5325](#)] [[INSPIRE](#)].
- [86] R. Boughezal, F. Caola, K. Melnikov, F. Petriello and M. Schulze, *Higgs boson production in association with a jet at next-to-next-to-leading order*, *Phys. Rev. Lett.* **115** (2015) 082003 [[arXiv:1504.07922](#)] [[INSPIRE](#)].

- [87] R. Boughezal, C. Focke, W. Giele, X. Liu and F. Petriello, *Higgs boson production in association with a jet at NNLO using jetiness subtraction*, *Phys. Lett. B* **748** (2015) 5 [[arXiv:1505.03893](#)] [[INSPIRE](#)].
- [88] F. Caola, K. Melnikov and M. Schulze, *Fiducial cross sections for Higgs boson production in association with a jet at next-to-next-to-leading order in QCD*, *Phys. Rev. D* **92** (2015) 074032 [[arXiv:1508.02684](#)] [[INSPIRE](#)].
- [89] X. Chen, J. Cruz-Martinez, T. Gehrmann, E.W.N. Glover and M. Jaquier, *NNLO QCD corrections to Higgs boson production at large transverse momentum*, *JHEP* **10** (2016) 066 [[arXiv:1607.08817](#)] [[INSPIRE](#)].
- [90] J.M. Campbell, R.K. Ellis and S. Seth, *H + 1 jet production revisited*, *JHEP* **10** (2019) 136 [[arXiv:1906.01020](#)] [[INSPIRE](#)].
- [91] S.P. Jones, M. Kerner and G. Luisoni, *Next-to-Leading-Order QCD Corrections to Higgs Boson Plus Jet Production with Full Top-Quark Mass Dependence*, *Phys. Rev. Lett.* **120** (2018) 162001 [[arXiv:1802.00349](#)] [[INSPIRE](#)].
- [92] T. Neumann, *NLO Higgs+jet production at large transverse momenta including top quark mass effects*, *J. Phys. Comm.* **2** (2018) 095017 [[arXiv:1802.02981](#)] [[INSPIRE](#)].
- [93] T. Neumann and C. Williams, *The Higgs boson at high p_T* , *Phys. Rev. D* **95** (2017) 014004 [[arXiv:1609.00367](#)] [[INSPIRE](#)].
- [94] ATLAS collaboration, *Precision measurement and interpretation of inclusive W^+ , W^- and Z/γ^* production cross sections with the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 367 [[arXiv:1612.03016](#)] [[INSPIRE](#)].
- [95] S. Alioli et al., *Precision studies of observables in $pp \rightarrow W \rightarrow l\nu_l$ and $pp \rightarrow \gamma, Z \rightarrow l^+l^-$ processes at the LHC*, *Eur. Phys. J. C* **77** (2017) 280 [[arXiv:1606.02330](#)] [[INSPIRE](#)].
- [96] S. Frixione and G. Ridolfi, *Jet photoproduction at HERA*, *Nucl. Phys. B* **507** (1997) 315 [[hep-ph/9707345](#)] [[INSPIRE](#)].
- [97] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Diphoton production at the LHC: a QCD study up to NNLO*, *JHEP* **04** (2018) 142 [[arXiv:1802.02095](#)] [[INSPIRE](#)].
- [98] J. Teunissen, *config_fortran — A configuration file parser for Fortran*, https://github.com/jannisteunissen/config_fortran/ (2019).