

The Structure of the Proton in the LHC Precision Era

Jun Gao^a, Lucian Harland-Lang^b, Juan Rojo^{c,d}

^a*Institute of Nuclear and Particle Physics,*

Shanghai Key Laboratory for Particle Physics and Cosmology,

School of Physics and Astronomy, Shanghai Jiao Tong University, Shanghai, China

^b*Department of Physics and Astronomy, University College London, WC1E 6BT, United Kingdom*

^c*Department of Physics and Astronomy, VU University, De Boelelaan 1081, 1081HV Amsterdam, The Netherlands*

^d*Nikhef, Science Park 105, NL-1098 XG Amsterdam, The Netherlands*

Abstract

We review recent progress in the determination of the parton distribution functions (PDFs) of the proton, with emphasis on the applications for precision phenomenology at the Large Hadron Collider (LHC). First of all, we introduce the general theoretical framework underlying the global QCD analysis of the quark and gluon internal structure of protons. We then present a detailed overview of the hard-scattering measurements, and the corresponding theory predictions, that are used in state-of-the-art PDF fits. We emphasize here the role that higher-order QCD and electroweak corrections play in the description of recent high-precision collider data. We present the methodology used to extract PDFs in global analyses, including the PDF parametrization strategy and the definition and propagation of PDF uncertainties. Then we review and compare the most recent releases from the various PDF fitting collaborations, highlighting their differences and similarities. We discuss the role that QED corrections and photon-initiated contributions play in modern PDF analysis. We provide representative examples of the implications of PDF fits for high-precision LHC phenomenological applications, such as Higgs coupling measurements and searches for high-mass New Physics resonances. We conclude this report by discussing some selected topics relevant for the future of PDF determinations, including the treatment of theoretical uncertainties, the connection with lattice QCD calculations, and the role of PDFs at future high-energy colliders beyond the LHC.

Keywords: Parton Distributions, Quantum Chromodynamics, Large Hadron Collider, Higgs boson, Standard Model, Electroweak theory

Email addresses: jung49@sjtu.edu.cn (Jun Gao), l.harland-lang@ucl.ac.uk (Lucian Harland-Lang), j.rojo@vu.nl (Juan Rojo)

Contents

1	Introduction	4
2	The global QCD analysis framework	6
2.1	A brief history of PDF fits	6
2.2	QCD factorization in deep-inelastic scattering	9
2.3	QCD factorization in hadronic collisions	11
2.4	The DGLAP evolution equations	12
2.5	Heavy quark structure functions	15
3	Experimental data and theoretical calculations	17
3.1	Overview	18
3.2	Deep-inelastic scattering	19
3.3	Inclusive jets	25
3.4	Inclusive gauge boson production	30
3.5	The p_T distribution of Z bosons	36
3.6	Direct photon production	42
3.7	Top quark production	45
3.8	Charm production in pp collisions	49
3.9	W production in association with charm quarks	52
3.10	Central exclusive production	53
3.11	Fast interfaces to (N)NLO calculations	54
4	Fitting methodology	57
4.1	PDF parametrization	57
4.1.1	Choice of functional form	58
4.1.2	Sum rules	59
4.1.3	Quark flavour assumptions	60
4.2	Fit quality and minimization strategies	61
4.2.1	Fit quality and χ^2 definition	62
4.2.2	Minimization strategies	63
4.3	PDF uncertainties	64
4.3.1	The Hessian method	64
4.3.2	The Monte Carlo method	65
4.3.3	The Lagrange multiplier method	68
4.4	Combined and reduced PDF sets	70
4.5	Treatment of theoretical parametric uncertainties	74
4.5.1	The strong coupling constant α_S	75
4.5.2	Heavy quark masses	76
4.6	Approximate methods	77
4.6.1	Bayesian Monte Carlo reweighting	78
4.6.2	Hessian profiling	79
4.7	Public delivery: LHAPDF	81

5	PDF analyses: state of the art	82
5.1	CT	82
5.2	MMHT	86
5.3	NNPDF	89
5.4	ABM	93
5.5	CTEQ-JLab (CJ)	96
5.6	HERAFitter/xFitter	98
5.7	PDF efforts by the LHC collaborations	100
5.7.1	ATLAS	102
5.7.2	CMS	102
6	The proton structure	103
6.1	The gluon PDF	104
6.2	Quark flavour separation	107
6.3	Strangeness	109
6.4	The charm content of the proton	111
7	Electroweak corrections and the photon PDF	115
7.1	Photon-induced processes	116
7.2	Electroweak corrections	123
8	Implications for LHC phenomenology	124
8.1	Higgs production cross-sections	124
8.2	PDF uncertainties and searches for new massive particles	127
8.3	Precision measurements of SM parameters	129
9	The future of PDF determinations	132
9.1	PDFs with theoretical uncertainties	132
9.1.1	MHOU in matrix element calculations	132
9.1.2	MHOU in PDF determination	134
9.2	Lattice QCD calculations of the proton structure	135
9.3	Parton distributions at future high-energy colliders	139
9.3.1	PDFs at high-energy lepton-hadron colliders	141
9.3.2	PDFs at a 100 TeV hadron collider	143
10	Conclusions	145

1. Introduction

The determination of the quark and gluon structure of the proton is a central component of the precision phenomenology program at the Large Hadron Collider (LHC). This internal structure is quantified in the collinear QCD factorisation framework by the Parton Distribution Functions (PDFs), which encode the information related to the momentum distribution of quarks and gluons within the proton. Being driven by low-scale non-perturbative dynamics, PDFs cannot be computed from first principles, at least with current technology, and therefore they need to be determined using experimental data from a variety of hard-scattering cross sections in lepton-proton and proton-proton collisions. This global QCD analysis program involves combining the most PDF-sensitive data and the highest precision QCD and electroweak calculations available within a statistically robust fitting methodology. See Refs. [1–9] for recent reviews on PDF determinations.

One of the main motivations to improve our understanding of the internal structure of the proton is provided by the fact that PDFs and their associated uncertainties play a decisive role in several LHC applications. To begin with, they represent one of the dominant theoretical uncertainties for the determination of the Higgs boson couplings [10], where any deviation from the tightly fixed SM predictions would indicate a smoking gun for new physics. PDF uncertainties also affect the production of new high-mass resonances, as predicted by many Beyond the Standard Model (BSM) scenarios [11], since they probe PDFs at large values of the momentum fraction x , for which current experimental constraints are scarce. A third example is provided by the measurement of precision SM parameters at hadron colliders, such as the W mass [12–15] or the strong coupling constant $\alpha_s(Q)$ [16–20]. These can be sensitive to BSM effects (for instance via virtual effects of new particles too heavy to be produced directly [21, 22]) and in many cases PDF uncertainties represent one of the limiting factors of the measurements.

Beyond the LHC, there are also several other instances where PDFs play an important role, for instance in astroparticle physics, such as for the accurate predictions for signal [23] and background [24–27] events at ultra-high energy neutrino telescopes. And needless to say, parton distributions will keep playing an important role for any future higher-energy collider involving hadrons in the initial state [28, 29], and therefore improving our understanding of PDFs also strengthens the physics potential of such future colliders.

A number of collaborations provide regular updates of their PDF sets, see [30–36] and references therein. There are a range of differences between these analyses, arising for example at the level of the selection of the input fitted dataset, the theoretical calculations of cross sections, methodological choices for the parametrisation of PDFs, the estimate and propagation of PDF uncertainties, and the treatment of external parameters. For instance, while some PDF fits are based on a global dataset, including the widest possible variety of experimental constraints, some others are based instead on reduced datasets (for example, without jet data) or even on a single dataset, as the HERAPDF2.0 set [34] which is based entirely on the HERA inclusive structure functions.

Despite these differences, it has been shown that, under some well-specified conditions, PDF sets can be statistically combined among them into a unified set. The most popular realisation of this paradigm are the PDF4LHC15 sets [2], which combine the CT14, MMHT14, and NNPDF3.0 sets using the Monte Carlo (MC) method [37], and are subsequently reduced to small number of Hessian eigenvectors [38, 39] or MC replicas [40] to facilitate phenomenological applications.

This Report is motivated by the fact that the recent years have seen a number of important breakthroughs in our understanding of the quark and gluon structure of the proton. To begin with, the impressive recent progress in NNLO QCD calculations has now made it possible to include essentially all relevant collider cross sections consistently into a NNLO global analysis, from top quark differential distributions [41] to inclusive jets [42] and dijets [43], isolated photons [44], and the p_T distribution of Z bosons [45, 46],

among others. These theoretical developments have been matched by the availability of high-precision measurements from ATLAS, CMS, and LHCb at different centre-of-mass energies \sqrt{s} , in several cases with statistical uncertainties at the per-mile level and systematic errors at the few-percent level. The combination of these state-of-the-art calculations and high-precision data provides an unprecedented opportunity to constrain the PDFs, but it also represents a challenge to verify if the global QCD analysis framework can satisfactorily accommodate them. Indeed, the validity of the QCD factorisation theorems is being pushed to a level never reached before.

In addition, there are a number of PDF-related topics that have attracted a lot of attention recently. One of these is the role that QED and electroweak effects, and specifically the photon PDF $\gamma(x, Q^2)$, play in global fits of parton distributions. Recent progress has demonstrated that the photon PDF can be computed with percent level accuracy [47, 48], improving on existing model [49–51] and data-driven determinations [52–54], with direct implications for LHC cross sections. Another important development is related to more sophisticated theoretical treatments of the charm quark PDF, moving beyond the assumption that charm is purely generated by perturbative DGLAP evolution [55, 56]. From the methodological point of view, there have been several improvements in the way that PDFs are parametrized and the various associated sources of uncertainty are estimated among the PDF fitting groups. In addition, there has also been a recent explosion in the number of tools available for PDF studies, from the open-source fitting framework `xFitter` [57], to new fast (N)NLO interfaces and public codes for the PDF evolution [58–60] and the efficient calculation of hadronic cross sections [61–65]. It is therefore the goal of this Report to present a detailed overview of these various recent developments, and how they have modified our present understanding of the quark and gluon structure of the proton, with emphasis on the resulting phenomenological applications.

This Report focuses only on one of the main aspects of the internal structure of nucleons, namely collinear unpolarized PDFs, which are its most relevant feature for the exploration of the high-energy frontier at the LHC. There are however many other aspects of the inner life of protons that would each deserve their own separate review, and that due to space limitations cannot be covered here. These include, among others, the determination of its spin structure by means of the polarized PDFs [66–68], the nuclear modifications of the free-proton PDFs [69–71], relevant for the understanding of cold nuclear matter effects at the RHIC and LHC heavy-ion program, and the three-dimensional imaging of nucleons in terms of transverse-momentum-dependent PDFs (TMD-PDFs) [72]. We only note here that these fields do not exist in isolation, and progress in these topics can often directly impact unpolarized PDF fits, and vice versa.

The structure of this Report is as follows. First, in Sect. 2 we review the theoretical foundations of the global PDF analysis framework, specifically the QCD factorization theorems of lepton-hadron and hadron-hadron collisions, as well as the scale dependence of the PDFs as encoded in the DGLAP equations. In Sect. 3 we discuss the hard-scattering experimental data, as well as the corresponding state-of-the-art theoretical calculations, that are used to constrain the PDFs in modern global analyses. In Sect. 4 we present the methodological framework of PDF fits, including the various approaches to parametrizing the PDFs and estimating and propagating the uncertainties from theory and data to physical cross sections. In Sect. 5 we summarise the main features of the different PDF collaborations that provide regular updates of their PDF fits, and then in Sect. 6 we compare them, assessing their differences and similarities for different aspects of the proton structure such as the gluon PDF, quark flavour separation, and the strange and charm content of the proton.

In Sect. 7 we then discuss the role that QED and electroweak corrections play in modern PDF fits, with emphasis on the photon content of the proton. In Sect. 8 we highlight a number of representative

examples of the role of PDFs and their uncertainties for the LHC precision physics program, in particular the characterization of the Higgs sector, searches for massive New Physics resonances, and the measurement of precision parameters such as M_W . In the last part of this Report, Sect. 9, we discuss some of the topics that are likely to play an important role for the future of PDF determinations, such as the quantification of theoretical uncertainties, the interplay with lattice QCD calculations, and the application of PDFs for future higher energy lepton–proton and proton–proton colliders. Finally, we conclude in Sect. 10.

2. The global QCD analysis framework

In this section we first present a brief historical account of PDF determinations. We then introduce the basic foundations of the global PDF analysis program, namely the QCD factorization theorems for lepton-hadron and hadron-hadron collisions. We also discuss the scale dependence of parton distributions as encoded in the DGLAP evolution equations. In the final part of this section we briefly address the topic of mass effects in deep-inelastic structure functions.

2.1. A brief history of PDF fits

The first indications of a non-trivial proton internal structure were obtained in the the pioneering experiments of Hofstadter *et al.* on elastic electron–nucleon scattering [73, 74]. In these experiments, by examining the deviations from the simple Mott scattering formulae for point–like particles, the finite extension of the proton could be resolved. There, the charge radius of the proton was determined to be $\simeq 0.7$ fm to within a few percent precision.

Although this result already hinted at an underlying substructure, the serious possibility that the proton is composite originated only later, with the ideas put forward independently by Zweig [75] and Gell–Mann [76] in 1964. By postulating the existence of three ‘aces’ (Zweig’s term) or ‘quarks’ (Gell–Mann’s) with fractional electric charge and baryon number, and spin–1/2, the complex structure of the hadrons and meson multiplets could be explained in a rather compact way. However, Zweig and Gell–Mann were understandably cautious about interpreting these objects as physical particles of finite mass, rather than simply convenient mathematical structures, specially given that the mechanism for binding such quarks together was not understood, and stable quarks had never been seen experimentally.

This situation changed dramatically in 1967 with the new experimental data on deep inelastic scattering (DIS) provided by the SLAC 20 GeV linear accelerator. The physicists of the SLAC–MIT collaboration were surprised to find that, in contrast to the case of elastic lepton–proton scattering, the two form factors associated with the DIS cross section, the so–called structure functions, were roughly independent of Q^2 [77, 78]. Moreover, these appeared to exhibit the scaling behaviour predicted by Bjorken in 1969 [79], namely that the structure functions should depend only on the ratio of Q^2 to the lepton energy loss ν in the proton rest frame¹.

These observations led Feynman to introduce the parton model [81], in which the incident lepton scatters incoherently and instantaneously from the point–like ‘partonic’ constituents of the proton. This concept, developed further in [82], naturally explains the observed Bjorken scaling behaviour, with the point–like partons in this simple picture providing no additional scale through which Bjorken scaling could be broken. At the same time Callan and Gross [83] showed that the DIS structure functions obey a simple relation for the case of spin– $\frac{1}{2}$ quark constituents, a finding that was also supported by the data [84]. These partons

¹Indications of this scaling were also observed at the DESY electron synchrotron in the same year [80].

were therefore naturally associated with the “constituent” quarks of Gell–Mann and Zweig. The subsequent demonstration of asymptotic freedom in 1973 in strongly interacting non–abelian gauge theories [85, 86] provided a simple explanation for the observed absence of free quarks, through the process of confinement, and the QCD parton model became the established approach to describe scattering processes in the strong interactions.

A central ingredient of this QCD parton model are the probability distributions of the momentum carried by these partons, known as the parton distribution functions, or PDFs. The first studies concentrated on developing simple models for these objects based on the limited experimental input available, for example: in [82] phase space considerations were used to conclude that the PDFs must also include a contribution from the now well-known sea of quark–antiquark pairs in addition to the valence quarks; in [87] a gluon PDF was introduced to account for the observed quark momentum fractions in a physically reasonable way consistent with energy conservation, and simple x dependencies of the PDFs were predicted according to general Regge theory and phase space expectations.

The idea of fitting a freely parameterised set of PDF followed soon after these initial studies. In [88] the approach of [87] was extended to a more general phenomenological form, and a 4–parameter fit to the quark PDFs was performed to the available data on proton and neutron structure functions. As the amount and type of data increased, the shape of the PDFs became increasingly general, see for example [89, 90]. Although the momentum fraction carried by the gluon could be determined by the missing contribution to DIS appearing in the momentum sum rule, it was only possible to fit its shape following the observation of scaling violations in the structure functions, first observed at FNAL [91] in 1974. Such Q^2 dependent deviations from simple Bjorken scaling occur due to higher–order QCD corrections to DIS and were directly connected through the DGLAP equation [92–95] to the Q^2 evolution of the PDFs. This allowed the first determinations of the shape of the gluon PDF to be made in [96] (see also [97]).

The subsequent LO fits of [98] (based on [99]) to fixed target structure function and neutrino DIS data, and [100], which also included J/ψ meson and muon pair hadroproduction, were widely used for a range of phenomenological applications. By the late 1980s, PDF fits at NLO in the strong coupling were standard, with the earlier analyses of [101, 102] fitting to fixed target DIS and the subsequent fits of [103–105] including prompt photon and Drell–Yan hadroproduction. By this time, the powerful strategy of the ‘global QCD analysis’ [105] had emerged, which emphasized fitting multiple experimental data sets from a diverse range of processes to disentangle PDFs of various flavors. Around the same period, the ‘dynamical’ PDF set of [106] was produced with the assumption that at low scale the quark sea vanished and the gluon becomes proportional to the valence quark distributions, themselves determined from DIS data.

Up to the early 90s, all DIS measurements was made with fixed target experiments and hence limited to the higher $x \gtrsim 0.01$ region. This changed in 1992 when the HERA high-energy collider at DESY started taking data. HERA collided 920 GeV protons with ≈ 27.5 GeV electrons for most of the run period, allowing the previously unexplored region down to $x \sim 10^{-4}$ region to be probed at high Q^2 . By 1994, data from HERA were included for the first in the MRS(A) [107] and CTEQ3 [108] global fits. These were also the first fits to include data from the Tevatron $p\bar{p}$ collider, with in particular the W asymmetry data providing new information on the quark flavour decomposition, as originally suggested in [109]. In the years that followed, further public releases within these approaches were produced, including in particular the increasingly precise HERA measurements, as well as Tevatron data on jet production, which placed new and important constraints on the poorly-constrained medium and large- x gluon.

These PDF sets corresponded to the best fit only, that is, no estimate of the uncertainty on the PDFs due to the errors on the data in the fit were included, beyond relatively simple studies where a range of fits under different input assumptions might be performed to give some estimate of the spread. This was an

acceptable situation when the uncertainties on the hadron collider data were sufficiently large, however as the data precision increased, the lack of a reliable estimate of PDF uncertainties rapidly became an issue. In the 1996 CDF measurement [110] of inclusive jet production, for example, there was an apparent excess of events at high jet E_{\perp} that was interpreted at the time as a possible sign for new physics. In the subsequent study of [111] it was shown that the gluon PDF could be modified in a way that still fit all available data, including the CDF jets. Clearly, a robust evaluation of the PDF uncertainties was needed to avoid such a situation.

The first attempts to produce such uncertainties, based on linear propagation of the experimental systematic and statistical errors through to the PDFs, considered a restricted set of DIS data [112–115]. The extension of these methods to the wider data set included in a global PDF fit was a complicated problem, both from a purely technical point of view, but also more conceptually. In particular, more conventional statistical approaches to evaluating the uncertainty on the fitted PDF parameters, such as a standard ‘ $\Delta\chi^2 = 1$ ’ variation, are only appropriate when fitting perfectly consistent data sets with purely Gaussian errors against a well-defined theory. For PDF fits based instead on a wide range of experimental measurements, none of these criteria are fulfilled: individual data sets are often found to have a low associated likelihood, with a χ^2 per degree of freedom well above one, the experimental systematic uncertainties will not generally be Gaussian in nature, the fixed-order perturbative theory calculation will carry its own (usually omitted) uncertainties, and there will be in general many different possible choices of the PDF parametrization leading to comparable values of the χ^2 . These issues were addressed in the CTEQ [116, 117] and MRST [118] PDF releases in 2002, with the basic idea being to allow the χ^2 to vary from the minimum by a larger degree, or ‘tolerance’, to properly account for these various effects.

In the mid to late 90s important advances in the theoretical treatment of heavy quark PDFs were also made, via the development [119, 120] of Variable Flavour Number Schemes (VFNS), which provided a unified treatment of initial-state heavy quarks in the both the $Q^2 \lesssim m_q^2$ and $Q^2 \gg m_q^2$ regimes. The proof of factorization for hard scattering with heavy quarks, which underpins this, was presented in 1998 [121]. The use of such schemes is widespread in modern PDF sets.

The calculation of the NNLO splitting functions in 2004 [122, 123] provided the necessary tools to go to NNLO in PDF fits, and by the time of the release of the MSTW08 [124] and CT10 [125] sets (the successors to the MRST and CTEQ sets, respectively) NNLO was the standard for global PDFs. At the same time the ABKM09 [126] NNLO PDFs were released. These were based on the earlier studies of [112, 127, 128], and fit to a reduced data set of DIS and fixed target Drell–Yan and dimuon production, with a classical ‘ $\Delta\chi^2 = 1$ ’ error treatment applied. A further set to consider a reduced data sample to appear at this time was the HERAPDF1.0 [129] PDFs. These included only the combined H1 and ZEUS measurements from the HERA Run I phase, with the aim of determining the PDFs from a completely consistent DIS data sample. This allowed the PDF uncertainty to again be described without the introduction of a larger tolerance factor, while the uncertainties due to model assumptions and choice of parameterisation were included in addition. This NLO set was extended to NNLO in the HERAPDF2.0 [34] PDFs, which used the final combined HERA I + II data sample. The NNLO JR09 [130] set included a range of DIS and fixed target data, applying both a ‘standard’ fitting approach and the ‘dynamical’ approach of [106]. The subsequent JR14 [35] set included a range of data updates, including jet production from the Tevatron.

The approaches described above differ greatly in many respects, both in the choice of input data sets, and the treatment of the corresponding theory predictions. However, while there are significant differences in the precise choice of parameterisation, in all cases these rely on parameterising the PDFs in terms of reasonably contained, $O(20 - 40)$, number of free variables. Moreover, while the precise prescription may vary, these are again all based on the ‘Hessian’ linear error propagation procedure. A different approach, first discussed

in [131], was taken by the NNPDF collaboration. There, the PDF functional forms are based on artificial neural networks, in particular feed-forward multi-layer perceptrons, allowing many more, $O(200 - 300)$, free parameters. In addition, rather than constructing the PDF error from the χ^2 variation about the best fit values, a ‘Monte Carlo’ (MC) approach is taken, with a large enough sample of PDF ‘replica’ sets each fit to randomly distributed pseudo-data generated according to the measured data values and their uncertainties. The first NNPDF1.0 fit was reported in [132], at NLO and to a range of DIS and fixed target data. Subsequently, NNPDF2.1 [133, 134] provided the first NNLO PDF set within this approach, and included Tevatron data for the first time.

More recently, data from the LHC has played an increasingly important role in PDF determination. The CT14 [31], MMHT14 [32] and NNPDF2.3 [135] sets included LHC data on jets, W and Z boson production, and top pair production for the first time. In addition, ABM12 [136] was the first set from this group to include constraints from the LHC, with data on W and Z boson production fit. As we will discuss in some detail in this review, these LHC measurements, which are being produced with increasingly high precision, are now providing some of the most stringent constraints on the PDFs.

In parallel to these developments, there has been increasing focus on the use of PDFs as precise tools for LHC phenomenology, emphasising the need for clear benchmarking exercises between sets and PDF combinations, to provide an overall PDF uncertainty. The PDF4LHC Working Group, created in 2008, has played a significant role in this, with the benchmarking described in [4] leading to first so-called PDF4LHC recommendation [137] for the use of PDFs and their uncertainties at the LHC. This has subsequently been updated in [2].

2.2. QCD factorization in deep-inelastic scattering

The importance of deep-inelastic scattering for PDF fits cannot be overemphasised. As discussed in the previous section, this process was instrumental in the discovery of quarks, and has since then represented the backbone of global PDF fits. The DIS mechanism is schematically represented in Fig. 1. Here, an energetic lepton, which can be either charged (electron or muon) or neutral (a neutrino) scatters off a proton (or some other hadron) by means of the interchange of a virtual photon γ^* or a W^\pm or Z boson. The large virtuality Q of the gauge boson, $Q \gg \Lambda_{\text{QCD}}$, ensures that the process can be described within the perturbative QCD factorization framework in terms of coefficient functions and parton distributions, as we show below.

The kinematics of the DIS process can be described in terms of a few invariant quantities, namely

$$x \equiv \frac{Q^2}{2P \cdot q}, \quad Q^2 \equiv -q^2, \quad y \equiv \frac{q \cdot P}{k \cdot P}, \quad (1)$$

where k and k' are the four-momenta of the incoming and outgoing leptons, q is the four-momentum of the exchanged gauge boson, and P is incoming proton’s momentum. Here x is known as the Bjorken variable, and although it is defined purely in terms of the kinematics of the initial and final-state particles, it can be shown that in the parton model it corresponds to the proton momentum fraction carried by the struck parton in the centre-of-mass frame of the high-energy lepton-proton collision. Recall that by momentum conservation $q = k' - k$, and thus all the variables in Eq. (1) can be determined by the knowledge of the incoming momenta of the lepton k and of the proton P as well as the outgoing momentum of the lepton k' without any reference to the final hadronic state X . In Eq. (1), $Q^2 = -q^2$ is the virtuality of the electroweak gauge boson and the variable y is known as the inelasticity. The centre of mass energy W of the proton-photon collision is given by

$$W^2 = (P + q)^2 = Q^2 \frac{1-x}{x} + m_p^2. \quad (2)$$

The value $x = 1$ corresponds to the elastic limit, where the proton remains intact after the collision.

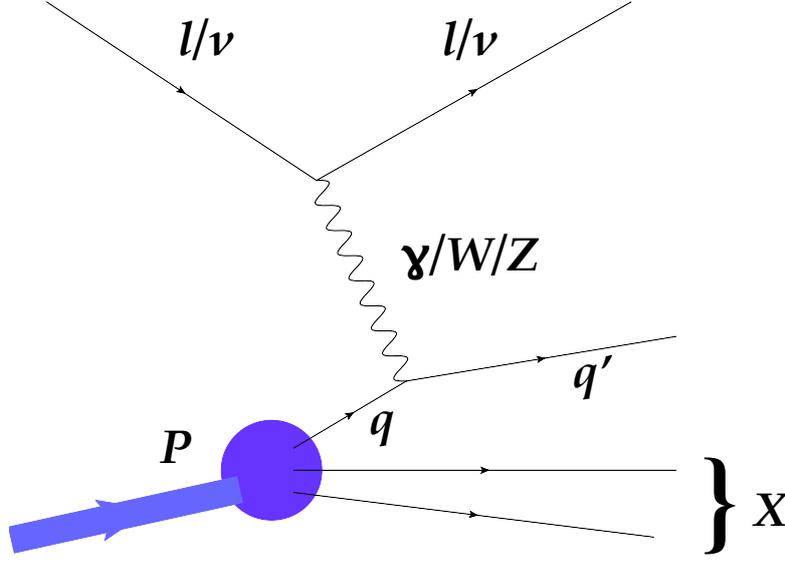


Figure 1: Schematic representation of the deep-inelastic scattering process. An energetic lepton (electron, muon or neutrino) scatters off one of the quarks in the proton by means of the interchange of an electroweak gauge boson (γ , W^\pm or Z). The large virtuality carried by the gauge boson, $Q \gg \Lambda_{\text{QCD}}$, ensures that the process can be described within the perturbative QCD factorization framework in terms of coefficient functions and PDFs.

It can be shown that only two of the three variables in Eq. (1) are independent, and therefore differential cross sections in DIS are measured for instance as a function of (x, Q^2) or (x, y) . Using Lorentz invariance and kinematic arguments, it can be shown that the DIS cross sections can be expressed in terms of a series of independent “structure functions” that describe the dynamics of the interaction between the gauge boson and the incoming hadron. In the neutral current (NC) case, that is, where either a virtual photon γ^* or a Z boson is exchanged, the DIS differential cross section for a charged lepton ℓ^\pm scattering off a proton can be decomposed in terms of structure functions as follows:

$$\frac{d^2\sigma^{\text{NC},\ell^\pm}}{dx dQ^2}(x, y, Q^2) = \frac{2\pi\alpha^2}{xQ^4} \left[Y_+ F_2^{\text{NC}}(x, Q^2) \mp Y_- x F_3^{\text{NC}}(x, Q^2) - y^2 F_L^{\text{NC}}(x, Q^2) \right], \quad (3)$$

where we have defined

$$Y_\pm = 1 \pm (1 - y)^2. \quad (4)$$

In most cases, experimental measurements are given in terms of a reduced cross section, defined as

$$\tilde{\sigma}^{\text{NC},\ell^\pm}(x, y, Q^2) = \left[\frac{2\pi\alpha^2}{xQ^4} Y_+ \right]^{-1} \frac{d^2\sigma^{\text{NC},\ell^\pm}}{dx dQ^2}(x, y, Q^2), \quad (5)$$

which is more closely related to the dominant structure function $F_2(x, Q^2)$, and thus to the underlying PDFs of the proton.

In the case of charged current (CC) DIS, when neutrinos are used as projectiles or when the incoming charged leptons interact with the proton by means of the exchange of a charged weak gauge boson W^\pm , the

differential cross sections are given by:

$$\begin{aligned} \frac{d^2\sigma^{\text{CC},\ell^\pm}}{dx dQ^2}(x, y, Q^2) &= \frac{G_F^2}{4\pi x} \left(\frac{M_W^2}{M_W^2 + Q^2} \right)^2 \\ &\times \frac{1}{2} \left[Y_+ F_2^{\text{CC},\ell^\pm}(x, Q^2) \mp Y_- x F_3^{\text{CC},\ell^\pm}(x, Q^2) - y^2 F_L^{\text{CC},\ell^\pm}(x, Q^2) \right]. \end{aligned} \quad (6)$$

which is generally rescaled to define a reduced cross section

$$\tilde{\sigma}^{\text{CC},\ell^\pm}(x, y, Q^2) = \left[\frac{G_F^2}{4\pi x} \left(\frac{M_W^2}{M_W^2 + Q^2} \right)^2 \right]^{-1} \frac{d^2\sigma^{\text{CC},\ell^\pm}}{dx dQ^2}(x, y, Q^2), \quad (7)$$

similarly to the NC case. In Eqns. (6) and (7), ℓ^\pm labels either the incoming or outgoing charged lepton.

According to the QCD factorization theorem, the general expression for the DIS structure functions can be written schematically as

$$F(x, Q^2) = x \int_x^1 \frac{dy}{y} \sum_i C_i \left(\frac{x}{y}, \alpha_s(\mu_R), \mu_F, Q \right) f_i(y, \mu_F), \quad (8)$$

where the C_i are known as the coefficient functions, $f_i(y, \mu_F)$ are the PDFs, and μ_F (μ_R) are the factorization (renormalization) scales, typically set to $\mu_R = \mu_F = Q$. The coefficient functions represent the cross section for the partonic scattering process $q_i + \gamma^* \rightarrow X$, and can be computed in perturbation theory as a series expansion in the strong coupling α_s , as well as in the electroweak coupling α_W if these corrections are included.

While the coefficient functions encode the short distance dynamics of the parton–boson collision, the PDFs are instead determined by long distance non–perturbative QCD dynamics, and so cannot be computed using perturbative methods. On the other hand, the crucial factorization property of Eq. (8) is that while the coefficient functions (or in general the partonic cross sections) are process dependent, the PDFs themselves are universal. This allows us to parameterise and extract the PDFs from a global analysis of hard scattering measurements. These can then be used to make predictions for other PDF–dependent processes.

2.3. QCD factorization in hadronic collisions

In a similar way to the DIS structure functions for electron–proton collisions, the production cross sections in proton–proton collisions can be factorized in terms of the convolution between two universal PDFs and a process–dependent partonic cross section. For example, the Drell–Yan production cross section, $\sigma^{\text{DY}}(pp \rightarrow l^+ l^- + X)$, can be expressed [138, 139] as

$$\frac{d^2\sigma^{\text{DY}}}{dy dQ^2}(y, Q^2, \mu_R^2, \mu_F^2) = \sum_{a,b=q,\bar{q},g} \int_{\tau_1}^1 dx_1 f_a(x_1, \mu_F^2) \int_{\tau_2}^1 dx_2 f_b(x_2, \mu_F^2) \frac{d^2\hat{\sigma}_{ab}^{\text{DY}}}{dy dQ^2}(x_1, x_2, y, Q^2, \mu_R^2, \mu_F^2), \quad (9)$$

where y and Q^2 are the rapidity and invariant mass squared of the lepton pair, and s is the centre–of–mass energy of the two incoming protons, while μ_F (μ_R) are the factorization (renormalization) scales. The lower integration limits are $\tau_{1,2} = \sqrt{Q^2/s} e^{\pm y}$. The partonic cross sections that appear in Eq. (9) can be computed as a perturbative expansion in α_s :

$$\frac{d^2\hat{\sigma}_{ab}^{\text{DY}}}{dy dQ^2}(x_1, x_2, y, Q^2, \mu_R^2, \mu_F^2) = \sum_{n=0}^{\infty} \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^n \frac{d^2\hat{\sigma}_{ab}^{(n)\text{DY}}}{dy dQ^2}. \quad (10)$$

From Eq. (9) we observe that the definition of the PDFs, once perturbative QCD corrections are accounted for, requires the introduction of a factorization scale μ_F , below which additional collinear emissions are absorbed into a PDF redefinition. To all orders, the physical cross section, as a product of the PDFs and partonic cross section, is independent of the choice of the factorization scale. However, at any fixed order in the perturbative series, there will be some sensitivity due to the missing higher orders, which can be minimised by choosing a suitable value of μ_F so as to maintain a better convergence of the series. In Drell–Yan production, the conventional scale choice is $\mu_F^2 = Q^2$, namely the invariant mass of the dilepton pair.

In full generality, for the case of the total inclusive cross section for a narrow resonance production with mass M , the cross section can be factorized as

$$\sigma = \sum_{a,b=q,\bar{q},g} \int_{M^2}^s \frac{d\hat{s}}{\hat{s}} \mathcal{L}_{ab}(\hat{s}, \mu_F^2) \hat{\sigma}_{ab}(\hat{s}, M^2, \mu_R^2, \mu_F^2), \quad (11)$$

where \hat{s} is the centre-of-mass energy of the two incoming partons, the parton–parton luminosity can be defined as [140]

$$\mathcal{L}_{ab}(\tau, \mu_F^2) = \frac{1}{s} \int_{\tau/s}^1 \frac{dx}{x} f_a(\tau/sx, \mu_F^2) f_b(x, \mu_F^2). \quad (12)$$

The partonic cross section depends only on the kinematic variable $z \equiv M^2/\hat{s}$ and $\mu_{F,R}$

$$\hat{\sigma}_{ab}(\hat{s}, M^2, \mu_R^2, \mu_F^2) = \sum_{n=0}^{\infty} \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^n C_{ab}^{(n)}(z, \mu_R^2, \mu_F^2). \quad (13)$$

The usefulness of this factorized form Eq. (11) is that the complete PDF dependence of the hadronic cross section is now encoded in the partonic luminosities \mathcal{L}_{ab} . As will be discussed in Sect 3, there has been significant recent progress in the higher-order calculation of partonic cross sections. In the case of inclusive processes, for instance, the coefficient functions $C_{ab}^{(n)}(z, \mu_R^2, \mu_F^2)$ are known to NNLO for Drell–Yan production [141] and top quark pair production [142], and to N³LO for Higgs boson production via gluon fusion in the limit of infinite top quark mass [143].

2.4. The DGLAP evolution equations

As discussed above, the PDFs depend on two variables: the Bjorken variable x , which at leading order can be identified with the momentum fraction carried by the considered parton, and the scale Q^2 , which in DIS corresponds to the virtuality of the exchanged gauge boson. While the dependence of the PDFs on x is determined by non-perturbative dynamics, and therefore cannot be computed perturbatively, the situation is different for the Q^2 variable. Here, the Q^2 dependence of the PDFs is introduced when higher-order initial-state collinear singularities of the partonic cross section are regularised by means of a PDF redefinition. Such singularities arise from universal long-distance QCD dynamics, and therefore are process-independent.

For this reason, the Q^2 dependence of the PDFs can be computed in QCD perturbation theory up to any given order. This dependence is determined by a series of integro-differential equations, known as the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) evolution equations [92–95], which have the generic form

$$Q^2 \frac{\partial}{\partial Q^2} f_i(x, Q^2) = \sum_j P_{ij}(x, \alpha_s(Q^2)) \otimes f_j(x, Q^2), \quad (14)$$

where $P_{ij}(x, \alpha_s(Q^2))$ are the Altarelli–Parisi splitting functions, which can be computed in perturbation theory

$$P_{ij}(x, \alpha_s(Q^2)) = \sum_{n=0} \left(\frac{\alpha_s(Q^2)}{2\pi} \right)^{n+1} P_{ij}^{(n)}(x), \quad (15)$$

and where \otimes denotes the convolution

$$f(x) \otimes g(x) \equiv \int_x^1 \frac{dy}{y} f(y) g\left(\frac{x}{y}\right), \quad (16)$$

which appears ubiquitously in QCD calculations. The splitting functions Eq. (15) depend on the type of initial and final state parton that is involved in the splitting. At leading order, the DGLAP splitting functions are given by

$$P_{qq} = \frac{4}{3} \left[\frac{1+x^2}{(1-x)_+} + \frac{3}{2} \delta(1-x) \right], \quad (17)$$

$$P_{qg} = \frac{1}{2} [x^2 + (1-x^2)], \quad (18)$$

$$P_{gq} = \frac{4}{3} \left[\frac{1+(1-x)^2}{x} \right], \quad (19)$$

$$P_{gg} = 6 \left[\frac{1-x}{x} + x(x-1) + \frac{x}{(1-x)_+} \right] + \frac{33-2n_f}{6} \delta(1-x). \quad (20)$$

Note that both P_{gg} and P_{qg} have a singularity at $x=0$: this fact is responsible for the rapid growth at small x of the gluons and consequently of the sea quarks in this region. The overall coefficients of the splitting functions are related to the QCD color factors. Some splitting functions exhibit a singularity at $x=1$, which is regularized by means of the plus prescriptions, defined as

$$\int_0^1 dx f(x) \left[\frac{1}{1-x} \right]_+ \equiv \int_0^1 dx (f(x) - f(1)) \left[\frac{1}{1-x} \right]. \quad (21)$$

The structure of the DGLAP evolution equations is significantly simplified if we use specific linear combinations of PDFs. For instance, below the charm threshold, where there are only $n_f = 3$ active quarks, the following combination

$$\begin{aligned} \Sigma(x, Q^2) &\equiv \sum_{i=1}^{n_f} (q_i + \bar{q}_i)(x, Q^2), \\ T_3(x, Q^2) &\equiv (u + \bar{u} - d - \bar{d})(x, Q^2), \\ T_8(x, Q^2) &\equiv (u + \bar{u} + d + \bar{d} - 2(s + \bar{s}))(x, Q^2), \\ V(x, Q^2) &\equiv \sum_{i=1}^{n_f} (q_i - \bar{q}_i)(x, Q^2), \\ V_3(x, Q^2) &\equiv (u - \bar{u} - d + \bar{d})(x, Q^2), \\ V_8(x, Q^2) &\equiv (u - \bar{u} + d - \bar{d} - 2(s - \bar{s}))(x, Q^2), \end{aligned} \quad (22)$$

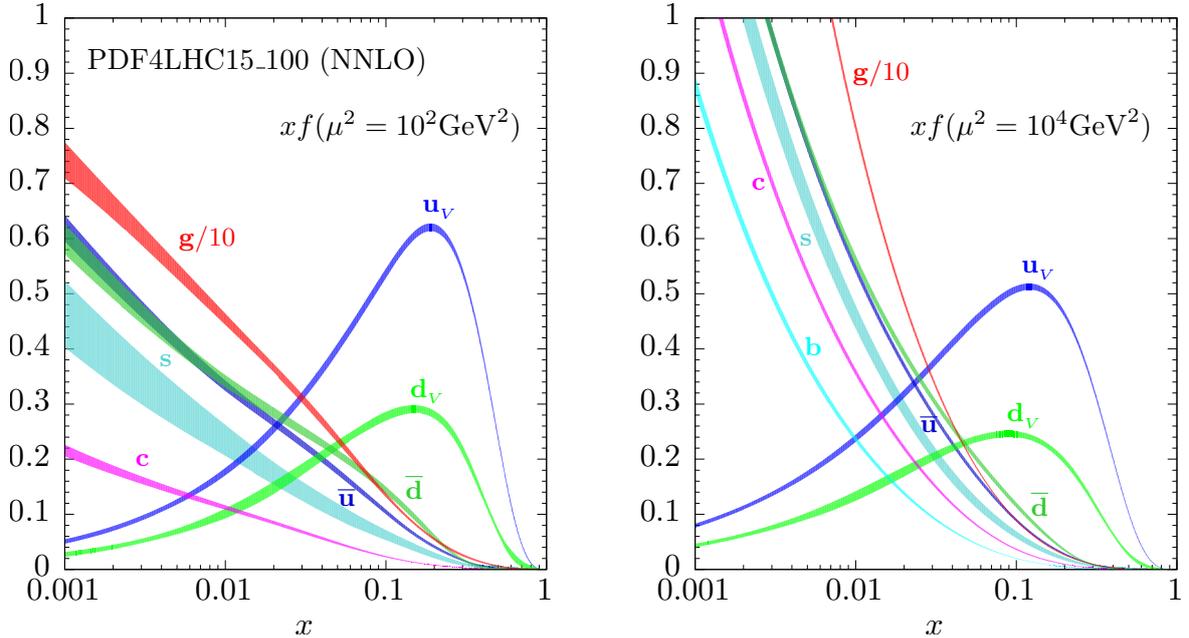


Figure 2: The effect of the DGLAP evolution in the PDF4LHC15 NNLO Hessian set (with 100 eigenvectors). We compare the PDFs at a low scale of $Q^2 = 10 \text{ GeV}^2$ (left) with the same PDFs evolved up to a typical LHC scale of $Q^2 = 10^4 \text{ GeV}^2$ (right plot). In this plot, the PDFs include the corresponding one-sigma uncertainty band.

has the property that all the quark PDF combinations except for Σ , known as the total quark singlet, evolve independently using their own specific splitting functions. These ‘non-singlet’ flavour combinations therefore obey a particularly simple evolution equation. As the $g \rightarrow q\bar{q}$ splitting can only generate an overall $q + \bar{q}$, only the singlet PDF evolution is explicitly coupled to the gluon.

The splitting functions Eq. (15) are known up to $\mathcal{O}(\alpha_s^3)$ (NNLO) [122, 123],² and thus PDF evolution can be performed up to this order. Several public codes implement the numerical solution of the DGLAP equations, with the HOPPET [59], APFEL [58] and QCDNUM [60] codes using x -space methods, while the PEGASUS [146] code performs the evolution in Mellin (moment) space. These codes have undergone detailed benchmarking studies, with agreement at the level of $\mathcal{O}(10^{-5})$ or better being found [147, 148].

In order to illustrate the impact of the DGLAP evolution on the PDFs, in Fig. 2 we show the PDF4LHC15 NNLO Hessian set with 100 eigenvectors, comparing the PDFs at a low scale of $Q^2 = 10 \text{ GeV}^2$ (left) with the same PDFs evolved up to a typical LHC scale of $Q^2 = 10^4 \text{ GeV}^2$ (right plot). In this plot, the PDFs are shown together with the corresponding one-sigma PDF uncertainty band. From this comparison we see that while the effects of the evolution are relatively mild on the non-singlet combinations $u_V = u - \bar{u}$ and $d_V = d - \bar{d}$, they are dramatic on the gluon and the sea quarks, where they induce a very steep growth at small x . This steep growth is driven by the small- x structure of the splitting functions, see Eq. (20). We also observe from Fig. 2 is that the valence PDFs xu_V and xd_V are integrable, with a similar shape but with $u_V \simeq 2d_V$. This behaviour is a consequence of the valence sum rules, see Sect. 4.

²Recently, the first results towards N³LO splitting functions have been presented [144, 145].

2.5. Heavy quark structure functions

The contribution of the charm structure function $F_2^c(x, Q^2)$ to the total inclusive neutral current structure function $F_2^p(x, Q^2)$ at HERA can be as high as 25% at small x and Q^2 , and therefore it is crucial to compute it with high accuracy. In particular, the finite value of heavy quark mass must be taken into account. There are various theoretical schemes that have been proposed for the computation of heavy quark production in DIS structure functions:

- The Zero–Mass Variable Flavour Number scheme (ZM–VFNS), where all heavy quark mass effects are ignored but potentially large logarithms, $\ln Q/m$, are resummed into the heavy quark parton distribution. This is also known as the massless scheme.
- The Fixed–Flavour Number Scheme (FFNS), where the heavy quark is always treated as a massive particle and never as a massless parton irrespective of the value of the scale Q . In this scheme, the heavy quark PDF does not exist and the number of active flavours is always kept fixed. This scheme takes into account heavy quark mass effects in the coefficient functions, but does not resum logarithmically enhanced terms of the form $\ln Q/m$ that become numerically relevant at high scales. This is also known as the massive scheme.
- The General–Mass Variable Flavour Number scheme (GM–VFNS) combines the advantage of the massive and massless calculations by means of an interpolated scheme which is valid for any value of the scale Q , and that matches the FFN and ZM–VFNS schemes at small and large values of Q , respectively.

Here we review the basic steps that enter into the construction of the GM–VFNS calculation of heavy quark DIS structure functions, using the FONLL derivation from Ref. [149] for illustration purposes. Note however that from the phenomenological point of view the resulting construction turns out to be rather similar to that of related GM–VFNS schemes such as ACOT [119], S–ACOT [150] and TR [120, 151]. Moreover, residual differences can be traced back to the treatment of subleading corrections, as explicitly demonstrated in the Les Houches benchmark studies of heavy quark structure functions [152, 153].

We start by the expression of a generic DIS structure function $F(x, Q^2)$, in a kinematic regime where one has n_l light flavours and a single heavy flavour of mass m . In the massless scheme, accurate when $W \gg 4m^2$, the expression of F in terms of PDFs and coefficient functions is the following

$$F^{(n_l+1)}(x, Q^2) = x \int_x^1 \frac{dy}{y} \sum_{i=q, \bar{q}, h, \bar{h}, g} C_i^{(n_l+1)}\left(\frac{x}{y}, \alpha_s^{(n_l+1)}(Q^2)\right) f_i^{(n_l+1)}(y, Q^2), \quad (23)$$

where q are the light quarks and h is the heavy quark. As indicated from the sum, in this scheme the heavy quark is treated as a massless parton, with all finite mass effects therefore neglected.

Now, in the massive (or decoupling [154]) scheme, which is most suitable when $W \approx 4m^2$ and thus heavy quark mass effects must be accounted for, this structure function reads

$$F^{(n_l)}(x, Q^2) = x \int_x^1 \frac{dy}{y} \sum_{i=q, \bar{q}, g} C_i^{(n_l)}\left(\frac{x}{y}, \frac{Q^2}{m^2}, \alpha_s^{(n_l)}(Q^2)\right) f_i^{(n_l)}(y, Q^2). \quad (24)$$

where now the massive coefficient functions $C_i^{(n_l)}$ includes the full mass dependence, and the heavy quark is no longer treated as a massless initial–state parton. In this scheme, the PDFs and α_s satisfy evolution

equations with n_l active quarks. Massive coefficient functions are known up to $\mathcal{O}(\alpha_s^2)$ both for neutral-current [155, 156] and charged-current [157, 158] scattering, with a number of partial results available for the massive neutral-current $\mathcal{O}(\alpha_s^3)$ coefficient functions, see [159] and references therein.

The construction of the GM–VFNS structure functions is then based on two steps. First of all, one needs to express the PDFs and α_s in the massless scheme by means of the matching conditions

$$\alpha_s^{(n_l+1)}(Q^2) = \alpha_s^{(n_l)}(Q^2) + \sum_{i=2}^{\infty} c_i(L) \times (\alpha_s^{(n_l)}(m^2))^i, \quad (25)$$

$$f_i^{(n_l+1)}(x, Q^2) = \int_x^1 \frac{dy}{y} \sum_{j=q, \bar{q}, g} K_{ij} \left(\frac{x}{y}, L, \alpha_s^{(n_l)}(Q^2) \right) f_j^{(n_l)}(y, Q^2), \quad (26)$$

where $L \equiv \log Q^2/m^2$. Then one uses these transformed expressions to write down $F^{(n_l)}$ in terms of PDFs and α_s in the massless scheme,

$$F^{(n_l)}(x, Q^2) = x \int_x^1 \frac{dy}{y} \sum_{i=q, \bar{q}, g} B_i \left(\frac{x}{y}, \frac{Q^2}{m^2}, \alpha_s^{(n_l+1)}(Q^2) \right) f_i^{(n_l+1)}(y, Q^2), \quad (27)$$

Once we have expressed both $F^{(n_l)}$ and $F^{(n_l+1)}$ in terms of PDFs and α_s in the massless scheme, the second step is to match the two expressions while removing any double counting. In this way we will maintain the main advantages of the two schemes (heavy quark mass effects in $F^{(n_l)}$, resummation of large $\ln Q^2/m^2$ logarithms in $F^{(n_l+1)}$) within a single scheme that is valid for any scale Q . To achieve this, one defines the massless limit of the massive scheme structure function as follows

$$F^{(n_l, 0)}(x, Q^2) = x \int_x^1 \frac{dy}{y} \sum_{i=q, \bar{q}, g} B_i^{(0)} \left(\frac{x}{y}, \frac{Q^2}{m^2}, \alpha_s^{(n_l+1)}(Q^2) \right) f_i^{(n_l+1)}(y, Q^2), \quad (28)$$

where in the coefficient functions $B_i^{(0)}$ all the terms which are power suppressed of the form m/Q are neglected, and the only dependence on the heavy quark mass m is on logarithms of the form $\ln Q/m$,

The FONLL approximation for F is then given by

$$F^{\text{FONLL}}(x, Q^2) = F^{(d)}(x, Q^2) + F^{(n_l)}(x, Q^2), \quad (29)$$

$$F^{(d)}(x, Q^2) \equiv [F^{(n_l+1)}(x, Q^2) - F^{(n_l, 0)}(x, Q^2)], \quad (30)$$

where Eq. (30) is constructed out of the massless scheme expression $F^{(n_l+1)}$, and the massless limit $F^{(n_l, 0)}$ of the massive scheme expression as in Eq. (28). It is thus easy to see that in the limit where $Q \gg m$, the FONLL structure function reduces to the massless calculation, while for $Q \sim m$ the FONLL result coincides with the massive calculation up to subleading (higher order) terms. Since close to threshold the difference term $F^{(d)}$ is formally subleading but numerically non-negligible, it is customary to further suppress it using for instance a damping factor [160] or χ rescaling [161], see also the discussion in [162].

In the specific case of FONLL, the GM–VFNS formalism can also be generalised to include mass effects for heavy quark initiated contribution [163], such as those required *i.e.* in the presence of a non-perturbative charm content in the proton. See [55] and references therein for the corresponding discussion of massive heavy quark initiated contribution in the ACOT framework, as well as [164] for a explicit comparison between the FONLL and ACOT derivations.

To illustrate the numerical impact that the heavy quark mass effects have on the DIS functions, in Fig. 3 we show the inclusive proton structure function $F_2(x, Q^2)$ at NNLO [165] as a function of Q^2 for

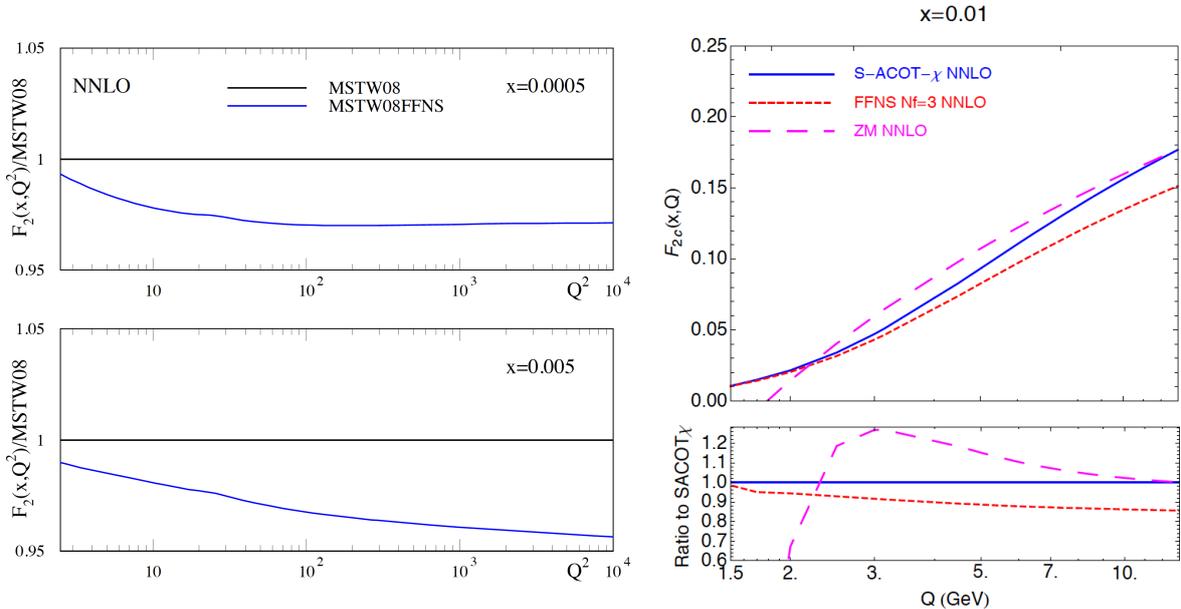


Figure 3: Left plot: the inclusive proton structure function $F_2(x, Q^2)$ at NNLO as a function of Q^2 for two different values of x in the TR' GM-VFNS as compared to the FFNS calculation [165]. Right plot: the NNLO charm structure function $F_2^c(x, Q)$ as a function of Q for $x = 0.01$ comparing the S-ACOT- χ GM-VFNS with the corresponding ZM and FFN scheme calculations [167].

two different values of x in the TR' GM-VFNS [166], compared to the FFNS calculation. We see that differences can be as large as a few percent, comparable to or larger than the uncertainties on the available DIS data. In the same figure (right) we show the NNLO charm structure function $F_2^c(x, Q)$ as a function of Q for $x = 0.01$ comparing the S-ACOT- χ [167] GM-VFNS with the corresponding ZM and FFN scheme calculations. We can see that the S-ACOT- χ calculation smoothly interpolates between the FFN scheme at low values of Q and the massless result at high Q .

Finally, it is important to note that there are subtleties associated with the definition of $F_2^c(x, Q^2)$ starting at $\mathcal{O}(\alpha_s^2)$ [149, 167]. From the experimental point of view, all processes leading to a charm quark in the final state enter the definition of F_2^c . On the other hand, from the theory point of view it is more natural to define F_2^c as including only those terms where the charm quark couples to the virtual photon γ^* , and to include other diagrams with final-state charm quarks, such as those arising from gluon radiation off light quarks, in the 'light' part of the structure function. Fortunately, the mismatch between the theoretical and experimental definitions is small and limited to the large- x and large- Q^2 region, where experimental uncertainties on F_2^c are in any case rather large. Moreover, this issue does not affect the inclusive structure function measurements and corresponding theory calculations as it arises only when the charm quark part of F_2 is explicitly separated from the light quark part.

3. Experimental data and theoretical calculations

In this section, we discuss the experimental data that are used in modern global PDF analyses, as well as the status of the corresponding theoretical calculations and fast interfaces for their inclusion in the PDF fits.

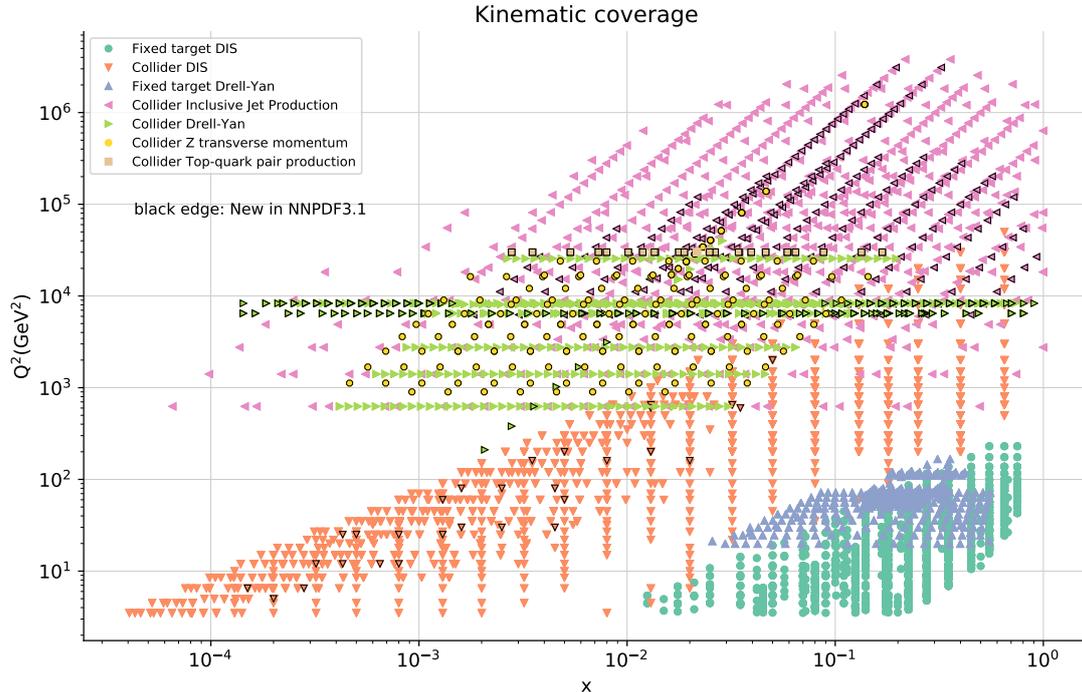


Figure 4: Typical kinematical coverage in the (x, Q^2) plane for the datasets included in a global analysis, in this case NNPDF3.1. For hadronic observables, leading order kinematics are assumed to map each data bin to a pair of (x, Q^2) values. The various datasets are clustered into families of related processes.

For each process, we discuss first the PDF sensitivity, then the available data and state-of-the-art theory calculations, and finally illustrate the impact on PDFs. We start with a general overview of the datasets that are available for PDF studies and then we move to discuss each process separately, beginning with DIS and then considering inclusive jet and weak boson production, the p_T of Z bosons, direct photon, top-quark pair production, open charm and W boson in association with charm production, and central exclusive production. In the last part of this section we discuss the technical but important topic of fast (N)NLO interfaces, allowing the efficient inclusion of higher-order calculations in PDF fits.

While the list of processes described in this section focuses on those that have so far been demonstrated to provide the most important constraints for PDF determinations, it is certainly not intended to be exhaustive. There are certainly other processes that we do not discuss here, such as W and Z boson production in association with jets and their ratios [168], which can provide PDF constraints.

3.1. Overview

In Fig. 4 we show the kinematic coverage in the (x, Q^2) plane of the datasets included in a representative global analysis, in this case the recent NNPDF3.1 fit [169]. For the hadronic observables, leading order kinematics are assumed to map each data bin to a pair of (x, Q^2) values, while the various datasets are clustered into families of related processes.

We can see that a global dataset provides a rather wide coverage in the (x, Q^2) plane. The low- x and Q^2 region is dominated by the inclusive HERA structure function measurements, which provide information down to $x \sim 3 \cdot 10^{-5}$. The high- x region is covered by various processes, from fixed-target DIS structure

functions at low Q^2 to collider jet, Drell–Yan and top quark pair production at large Q^2 . The very high Q^2 region, up to a few TeV^2 , is only covered by inclusive jet production data from ATLAS and CMS. Until relatively recently, most PDF fits were only based on DIS and fixed–target data, with some data from the Tevatron included. The breath of experimental information that is now included in the latest PDF fits is therefore quite impressive, with data from processes such as the $Z p_T$ and the $t\bar{t}$ differential distribution only recently being considered for the first time.

In Table 1 (an extended version of Table 1 from [124]) we present another overview of the data entering a modern global PDF analysis. Here, we summarise the various hard scattering processes which are used to constrain PDFs in a global analysis. In each case we indicate the hadron–level process, the corresponding dominant parton–level process, as well as the partons which are constrained in each case and the corresponding range of x . Note that the latter are necessarily approximate, and only indicate in a qualitative way the x region that dominates the PDF sensitivity of each measurement. The necessity to include as broad a set of input datasets as possible, in order to constrain all quark flavour combinations and the gluon in the phenomenologically relevant region of x , is clear. We also note that the medium to small– x region, $x \lesssim 0.01$, is mostly covered by the HERA collider structure functions as well as by some LHC data. The very small– x region, below the coverage of the HERA data, $x \lesssim 5 \cdot 10^{-5}$, can only be accessed via D meson production and exclusive J/ψ production.

In the rest of this section, we discuss in turn the various processes that can be used to constrain the parton distributions in a global analysis, listed in Table 1. We restrict the discussion to theoretical calculations based on fixed–order perturbative QCD, see Refs. [170–172] and [173] for studies of the impact of the PDF fit of theory calculations based on all–order resummations of logarithmically enhanced terms at small– x and large– x respectively.

3.2. Deep-inelastic scattering

PDF sensitivity

Before the establishment of QCD as the renormalizable quantum field theory of the strong interaction, the results of DIS experiments were interpreted in the context of the so–called quark parton model. In this model, the proton was composed by non–interacting, co–moving quarks each of them carrying a given fraction x of its total momentum, and then the DIS structure functions have particularly simple expressions in terms of the PDFs. Moreover, in this model the PDFs have a simple probabilistic interpretation, with $q_i(x)\Delta x$ giving the probability of finding a quark of flavour i inside the proton carrying out a momentum fraction in the range $[x, x + \Delta x]$. The expressions of the DIS structure functions in the quark parton model therefore provide a useful way to illustrate the PDF sensitivity of this process.

For the NC DIS structure functions F_2 and F_3 , as defined in (3), the quark parton model expressions are given by

$$\left[F_2^\gamma, F_2^{\gamma Z}, F_2^Z \right] = x \sum_{i=1}^{n_f} \left[e_i^2, 2e_i g_V^i, g_V^{i2} + g_A^{i2} \right] (q_i + \bar{q}_i), \quad (31)$$

$$\left[F_3^\gamma, F_3^{\gamma Z}, F_3^Z \right] = x \sum_{i=1}^{n_f} \left[0, 2e_i g_A^i, 2g_V^i g_A^i \right] (q_i - \bar{q}_i), \quad (32)$$

while the longitudinal structure function vanishes in this model, $F_L = 0$, and the superscripts on the LHS indicate the gauge boson which is being interchanged, as well as the contribution from the γZ interference term. e_i is the electric charge of the quark of flavour i and the weak couplings are given by $g_V^i = \pm \frac{1}{2} - 2e_i \sin^2 \theta_W^2$ and $g_A^i = \pm \frac{1}{2}$, where the \pm corresponds to a u or d type quark. The sum runs over all the n_f quarks that are active for the specific scale at which the scattering takes place. From Eqns. (31) and (32) we

	Process	Subprocess	Partons	x range
Fixed Target	$\ell^\pm \{p, n\} \rightarrow \ell^\pm + X$	$\gamma^* q \rightarrow q$	q, \bar{q}, g	$x \gtrsim 0.01$
	$\ell^\pm n/p \rightarrow \ell^\pm + X$	$\gamma^* d/u \rightarrow d/u$	d/u	$x \gtrsim 0.01$
	$pp \rightarrow \mu^+ \mu^- + X$	$u\bar{u}, d\bar{d} \rightarrow \gamma^*$	\bar{q}	$0.015 \lesssim x \lesssim 0.35$
	$pn/pp \rightarrow \mu^+ \mu^- + X$	$(u\bar{d})/(u\bar{u}) \rightarrow \gamma^*$	\bar{d}/\bar{u}	$0.015 \lesssim x \lesssim 0.35$
	$\nu(\bar{\nu})N \rightarrow \mu^-(\mu^+) + X$	$W^* q \rightarrow q'$	q, \bar{q}	$0.01 \lesssim x \lesssim 0.5$
	$\nu N \rightarrow \mu^- \mu^+ + X$	$W^* s \rightarrow c$	s	$0.01 \lesssim x \lesssim 0.2$
	$\bar{\nu} N \rightarrow \mu^+ \mu^- + X$	$W^* \bar{s} \rightarrow \bar{c}$	\bar{s}	$0.01 \lesssim x \lesssim 0.2$
Collider DIS	$e^\pm p \rightarrow e^\pm + X$	$\gamma^* q \rightarrow q$	g, q, \bar{q}	$0.0001 \lesssim x \lesssim 0.1$
	$e^+ p \rightarrow \bar{\nu} + X$	$W^+ \{d, s\} \rightarrow \{u, c\}$	d, s	$x \gtrsim 0.01$
	$e^\pm p \rightarrow e^\pm c\bar{c} + X$	$\gamma^* c \rightarrow c, \gamma^* g \rightarrow c\bar{c}$	c, g	$10^{-4} \lesssim x \lesssim 0.01$
	$e^\pm p \rightarrow e^\pm b\bar{b} + X$	$\gamma^* b \rightarrow b, \gamma^* g \rightarrow b\bar{b}$	b, g	$10^{-4} \lesssim x \lesssim 0.01$
	$e^\pm p \rightarrow \text{jet} + X$	$\gamma^* g \rightarrow q\bar{q}$	g	$0.01 \lesssim x \lesssim 0.1$
Tevatron	$p\bar{p} \rightarrow \text{jet} + X$	$gg, qg, q\bar{q} \rightarrow 2j$	g, q	$0.01 \lesssim x \lesssim 0.5$
	$p\bar{p} \rightarrow (W^\pm \rightarrow \ell^\pm \nu) + X$	$ud \rightarrow W^+, \bar{u}\bar{d} \rightarrow W^-$	u, d, \bar{u}, \bar{d}	$x \gtrsim 0.05$
	$p\bar{p} \rightarrow (Z \rightarrow \ell^+ \ell^-) + X$	$uu, dd \rightarrow Z$	u, d	$x \gtrsim 0.05$
	$p\bar{p} \rightarrow t\bar{t} + X$	$q\bar{q} \rightarrow t\bar{t}$	q	$x \gtrsim 0.1$
LHC	$pp \rightarrow \text{jet} + X$	$gg, qg, q\bar{q} \rightarrow 2j$	g, q	$0.001 \lesssim x \lesssim 0.5$
	$pp \rightarrow (W^\pm \rightarrow \ell^\pm \nu) + X$	$u\bar{d} \rightarrow W^+, d\bar{u} \rightarrow W^-$	$u, d, \bar{u}, \bar{d}, g$	$x \gtrsim 10^{-3}$
	$pp \rightarrow (Z \rightarrow \ell^+ \ell^-) + X$	$q\bar{q} \rightarrow Z$	q, \bar{q}, g	$x \gtrsim 10^{-3}$
	$pp \rightarrow (Z \rightarrow \ell^+ \ell^-) + X, p_\perp$	$gq(\bar{q}) \rightarrow Zq(\bar{q})$	g, q, \bar{q}	$x \gtrsim 0.01$
	$pp \rightarrow (\gamma^* \rightarrow \ell^+ \ell^-) + X, \text{Low mass}$	$q\bar{q} \rightarrow \gamma^*$	q, \bar{q}, g	$x \gtrsim 10^{-4}$
	$pp \rightarrow (\gamma^* \rightarrow \ell^+ \ell^-) + X, \text{High mass}$	$q\bar{q} \rightarrow \gamma^*$	\bar{q}	$x \gtrsim 0.1$
	$pp \rightarrow W^+ \bar{c}, W^- c$	$sg \rightarrow W^+ c, \bar{s}g \rightarrow W^- \bar{c}$	s, \bar{s}	$x \sim 0.01$
	$pp \rightarrow t\bar{t} + X$	$gg \rightarrow t\bar{t}$	g	$x \gtrsim 0.01$
	$pp \rightarrow D, B + X$	$gg \rightarrow c\bar{c}, b\bar{b}$	g	$x \gtrsim 10^{-6}, 10^{-5}$
	$pp \rightarrow J/\psi, \Upsilon + pp$	$\gamma^*(gg) \rightarrow c\bar{c}, b\bar{b}$	g	$x \gtrsim 10^{-6}, 10^{-5}$
$pp \rightarrow \gamma + X$	$gq(\bar{q}) \rightarrow \gamma q(\bar{q})$	g	$x \gtrsim 0.005$	

Table 1: Overview of the various hard-scattering processes which are used to constrain PDFs in a global analysis. In each case we indicate the hadron-level process and the corresponding dominant parton-level process, as well as the partons which are constrained by each specific process in a given range of x . This table is an extended version of Table 1 of [124]. The x ranges are merely indicative and based on approximate leading-order kinematics.

see that the main limitation of the NC structure functions is that they provide limited access to quark flavour separation and in particular they cannot separate quarks from antiquarks, unless one goes to very high Q^2 values where the suppression induced by the Z boson propagator can be ignored.

In the case of CC DIS, the corresponding expressions for the structure functions in the parton model, assuming that we are above the charm threshold but below the top quark threshold, and the CKM suppressed

transitions can be neglected, are given by

$$\begin{aligned}
F_2^{W^-} &= 2x(u + \bar{d} + \bar{s} + c), \\
F_3^{W^-} &= 2x(u - \bar{d} - \bar{s} + c), \\
F_2^{W^+} &= 2x(d + \bar{u} + \bar{c} + s), \\
F_3^{W^+} &= 2x(d - \bar{u} - \bar{c} + s),
\end{aligned} \tag{33}$$

where again the longitudinal structure function $F_L^{W^\pm} = 0$ vanishes in this model. By comparing the NC and CC expressions, we can see that the main difference between them is that in the latter case the F_3^W structure function, which provides information on the difference between quark flavours, is not suppressed with respect to F_2^W . For this reason, CC structure functions, both from HERA and from neutrino fixed-target experiments, are generally included in global fits in order to improve the discrimination between quarks and anti-quarks.

These quark parton model expressions are also valid at LO in perturbative QCD, once the effects of the DGLAP evolution are accounted for as described in Sect. 2.4. It is only at NLO that the contribution from the gluon PDF must also be included, and therefore the inclusive DIS structure functions will only be sensitive to the gluon PDF either through scaling violations (that is, the effect on the quark DGLAP evolution) or via the small $\mathcal{O}(\alpha_s)$ contribution to the coefficient functions. While scaling violations represent a small effect at medium and large- x , and thus the gluon in this region is poorly constrained in a DIS-only fit, scaling violations are a significantly more important effect at small- x , in particular in the region covered by HERA. Therefore, scaling violations from the HERA F_2^p data provide direct information on the small- x gluon.

Additional information on the gluon PDF can be obtained by means of the longitudinal structure function F_L , which vanishes at LO, while at NLO it is non-zero and directly sensitive to the gluon PDF. Indeed, it can be shown that this structure function is given by

$$F_L(x, Q^2) = \frac{\alpha_s(Q^2)}{\pi} \left[\frac{4}{3} \int_x^1 \frac{dy}{y} \left(\frac{x}{y} \right)^2 F_2(y, Q^2) + 2 \sum_i e_i^2 \int_x^1 \frac{dy}{y} \left(\frac{x}{y} \right)^2 \left(1 - \frac{x}{y} \right) g(y, Q^2) \right], \tag{34}$$

which is known as the Altarelli–Martinelli relation [174]. For this reason, F_L measurements can, in principle, provide direct constraints on the gluon in particular at small x , provided experimental uncertainties are competitive.

Finally, in addition to the inclusive structure functions, it is also possible to determine the heavy quark structure functions experimentally, by selecting DIS events with charm or bottom mesons in the final state. The LO process proceeds via the photon–gluon fusion mechanism $\gamma g \rightarrow q\bar{q}$, see Fig. 5 (right), and therefore heavy quark structure functions offer direct information on the gluon PDF, as well as on the treatment of heavy quark mass effects in the theoretical calculation. Charm structure functions in addition are an important ingredient for the determination of the charm mass m_c together with the PDFs [175–177]. On the other hand, while data on F_2^b are known to have a small impact in the global fit, it is relevant for specific applications, for instance the determination of the bottom quark mass m_b from the PDF fit.

Experimental data

Since the pioneering DIS experiments at SLAC in the late 60s and early 70s, there have been many other measurements of the DIS structure functions. These have been performed using either electrons, positrons, muons or neutrinos as projectile, and scattering off protons, deuterons and neutrons, either for fixed-target or for collider kinematics. We now discuss this various measurements in turn.

The fixed–target DIS measurements available for PDF fits can be divided into NC and CC datasets. In the NC case this includes:

- Proton and deuteron structure function data by the BCDMS collaboration [178, 179], using muons as projectiles.
- Proton and deuteron structure function data by the NMC collaboration [180, 181], as well as measurements of the ratio between deuteron to proton structure functions, F_2^d/F_2^p .
- SLAC measurements of the proton and deuteron NC structure functions [182].
- Proton, neutron and deuteron structure function data at high x and low Q^2 performed by JLAB experiments such as CLAS [183]. While these are excluded from most PDF fits by the typical DIS cuts in x and Q^2 , these are included in the CJ fits [36, 184, 185].
- Older structure function data from the EMC collaboration [186].

In the specific case of the EMC charm structure functions F_2^c , despite their age this measurement has never been repeated, and thus it provides unique information on the charm content of the proton at high x . Several groups have however experienced trouble in fitting this data, which is known to be affected by a number of problems such as underestimated systematic errors and an incorrect charm branching fraction into muons. The impact of this data on the PDFs, and in particular on the charm content of the proton, is further discussed in Sects. 5.3 and 6.4.

In the CC case we have:

- Inclusive structure function measurements due to neutrino beams on nuclear targets, by the CDHSW, CCFR [187, 188] and CHORUS [189] and NuTeV [190] collaborations.
- Charm production in neutrino–induced DIS. This process is often referred to as dimuon production, since the charm quark hadronizes into a D meson which then decays semi–leptonically, see Fig. 5. Data has been taken by the CCFR and NuTeV [191, 192] and CHORUS [193] collaborations on the same nuclear targets as the corresponding inclusive measurements, and also by the NOMAD collaboration [194].

For the DIS measurements from the HERA lepton–proton collider we have:

- The final measurements of the NC and CC differential cross sections using electron and positron projectiles from the combination of the Run I and Run II data–taking periods [34].
- The latest heavy flavour measurements from HERA include the combined NC cross sections of charm production in DIS, $\tilde{\sigma}_c$ [195] and the H1 and ZEUS data on the bottom structure function $F_2^b(x, Q^2)$ [196, 197].

This HERA legacy combination of DIS inclusive structure functions supersedes all previous inclusive measurements from H1 and ZEUS, including the Run I combined dataset [129] as well as the separate measurements by the two experiments from Run II [198–201]. The impact of replacing these individual datasets by the final HERA combination of inclusive structure functions has been studied by different groups [202–204], and is found to be quite moderate in general. We also note that previous measurements of the longitudinal structure function F_L by the H1 and ZEUS collaborations [199, 205, 206] are now superseded by the final inclusive HERA combination.

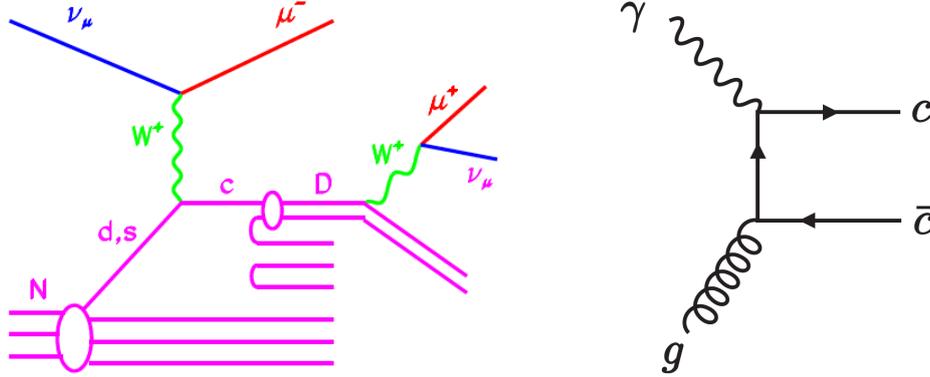


Figure 5: Left plot: D meson production in CC neutrino-induced DIS. This is known as the ‘dimuon’ process, since events are tagged when the D meson decays semi-leptonically, with the pair of oppositely-charged muons providing a clean signature. Right plot: charm production in neutral current DIS at leading order proceeds via the photon-gluon fusion process, highlighting its sensitivity to the gluon PDF.

Theoretical calculations and tools

The coefficient functions of the DIS structure functions in the NC case are available up to $\mathcal{O}(\alpha_s^3)$ in the massless limit [207, 208], and up to $\mathcal{O}(\alpha_s^2)$ taking into account heavy quark mass effects [155, 156], though there has been considerable recent progress towards the completion of the $\mathcal{O}(\alpha_s^3)$ calculation of massive DIS structure functions [209, 210], in particular of the terms that dominate in the $Q^2 \gg m^2$ limit. For charged current structure functions, massless coefficients are available up to $\mathcal{O}(\alpha_s^3)$ and massive coefficient functions up to $\mathcal{O}(\alpha_s^2)$ [157]. For heavy-quark initiated processes, massive coefficient functions are available only up to $\mathcal{O}(\alpha_s)$ [211].

These coefficient functions have been implemented in a number of private and public codes, which allow the efficient calculation of DIS structure functions using state-of-the-art theoretical information, such as QCDNUM [60], APFEL [58],³ and OpenQCDrad [213]. The lengthy exact expressions for the NNLO DIS coefficient functions are also available in the form of more compact interpolated expressions, which reduce the computational burden of their evaluation and allows for efficient evaluation of DIS cross sections. Moreover, DIS structure functions can be evaluated either in terms of the heavy pole mass or in terms of the running $\overline{\text{MS}}$ mass, as discussed in [213]. This statement is valid both in the FFNS as well as in any GM-VFNS, see for instance the discussion of the FONLL case in [175].

Impact on PDFs

From purely theoretical arguments, we expect that DIS structure functions in charged-lepton scattering will be constraining for some specific PDF combinations, in particular for the total quark singlet Σ , the quark isotriplet T_3 from deuteron measurements, and the gluon from scaling violations at HERA. The separation between quarks and anti-quarks can then be improved by using data taken with neutrino beams, which as discussed above probe different quark combinations than in charged-lepton DIS. Specifically, the strange PDF can be disentangled by exploiting the information contained in the neutrino-induced dimuon

³The APFEL program is currently being rewritten into C++ [212].

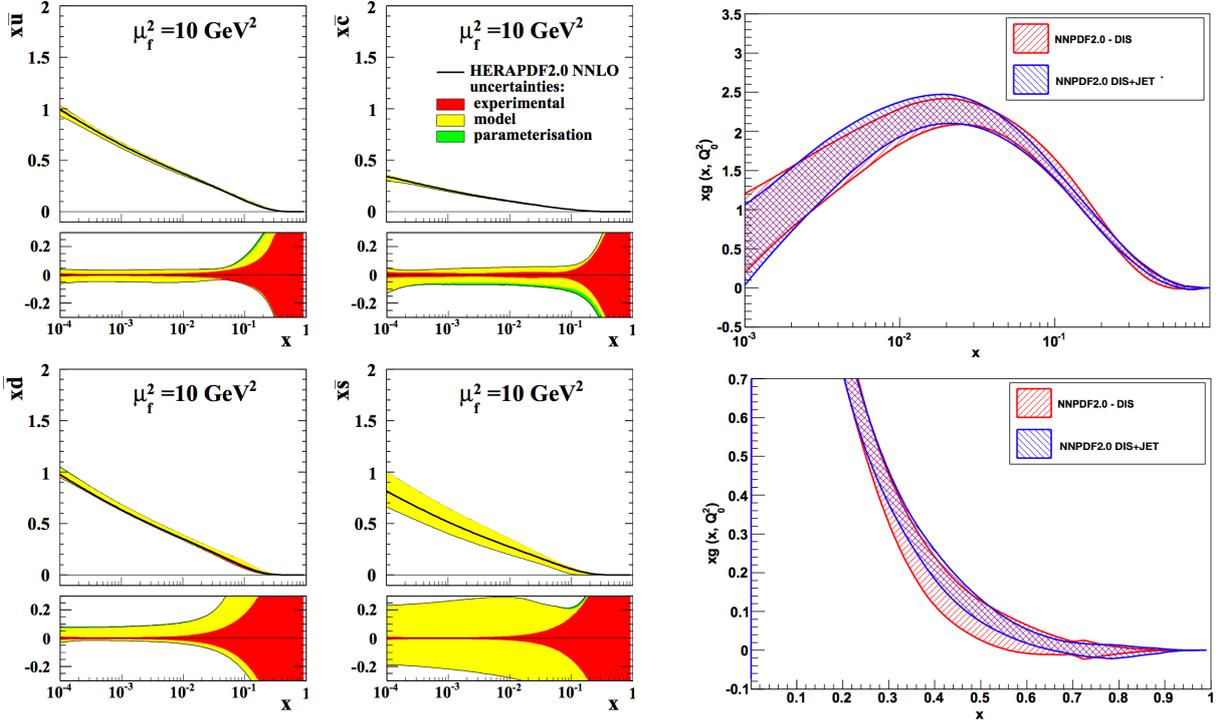


Figure 6: Left: the relative PDF uncertainties in the HERAPDF2.0NNLO fit, separated into experimental, model, and parametrization uncertainties, at $Q^2 = 10 \text{ GeV}^2$, for the various antiquark PDFs. Right: the gluon PDF in the NNPDF2.0 analysis at small x (top) and large x (bottom), comparing the results of a DIS-only fit with those of a DIS+jet fit. See text for more details.

cross sections. The same arguments indicate that other PDF combinations would be poorly constrained in a DIS-only fit, such as the gluon at large x .

Given that DIS structure functions constitute the backbone of any PDF fit, it is not straightforward to gauge its impact by adding say DIS data on top of a non-DIS dataset, as in the latter case the PDFs will be very poorly constrained to begin with. Here, in order to illustrate the role that DIS structure functions have in a PDF analysis, we have chosen instead two other related examples.

First, in Fig. 6 we show the relative PDF uncertainties in the HERAPDF2.0NNLO analysis, separated into experimental, model, and parametrization uncertainties, at $Q^2 = 10 \text{ GeV}^2$, for the various anti-quarks. This PDF fit is based only on the inclusive HERA structure function data, and therefore it can be used to illustrate the information from quark flavour separation that can (and that cannot) be obtained from DIS measurements. We observe that while at medium and small x the experimental PDF uncertainties are small, the contribution from model choices is larger, adding up to between 5% and 10%. The exception is the strange PDF, which is only poorly constrained by the HERA CC cross sections, and where errors are at $\sim 20\%$ or larger. Moreover, at large x the PDF uncertainties for all the antiquarks increase rapidly, reflecting the fact that the HERA data has no sensitivity to the quark flavour separation in this region.

While DIS data provide only loose information on quark flavour separation, unless the information from neutrino-nucleus scattering is used, they give some of the most constraining information on the gluon PDF in the medium and small- x region, in particular in the HERA region from scaling violations. To quantify this, in the right plot of Fig. 6 we show the gluon in the NNPDF2.0 analysis at small x (top) and large x (bottom),

comparing the results of a DIS-only fit with those of a DIS+jet fit. At small and medium x , the impact of jet data is minimal, showing that in a global fit the DIS data provides the most important constraints on the gluon in this kinematic region. On the other hand, at large x we observe that the information on the gluon from DIS data is only partial, and adding jet production to the fit helps to significantly reduce the PDF uncertainties there.

We also emphasise that neutrino-induced CC structure functions, in particular the dimuon cross sections, provide valuable constraints on the strange content of the proton [214–218]. The impact of various datasets into the strange PDF will be discussed in more detail in Sect. 6.3.

3.3. Inclusive jets

Since the first run of the Tevatron at Fermilab, inclusive jet production at hadron colliders has provided the dominant constraint on the gluon PDF at large x . The definition of jet cross sections starts from a well-defined jet algorithm [219], which is usually chosen to be infrared and collinear safe so that the corresponding parton-level cross section for hard scattering at high energies can be calculated in perturbative QCD. The most commonly used jet algorithm at the LHC is the anti- k_T algorithm [220], provided with the 4-vector recombination scheme. Other common choices include the k_T algorithm [221, 222], the Cambridge–Aachen algorithm [223], the Seedless Infrared-Safe cone algorithm (SIScone) [224], as well as the Midpoint algorithm [225], which was sometimes used at the Tevatron. The main parameter entering a jet algorithm is the jet radius R , which roughly defines how large the jet is in the (η, ϕ) plane, and whose optimal value depends on the specific application [226].

When comparing the predicted parton-level cross section to the experimentally measured jet cross section, it is essential to correct these to the hadron level. That is, additional non-perturbative corrections due, for example, to the underlying event and hadronization effects, must be accounted for. These are usually provided by the experimental collaborations as multiplicative factors derived from leading-order event generators. The size of such corrections can be significant at low p_T , as high as $\sim 20\%$, while at high p_T they are generally small, at the percent level [227]. Variations of these non-perturbative corrections, by considering for example difference generator predictions, are then treated as an additional source of correlated systematic error. Although PDF fits typically use parton-level predictions, results also exist that include the matching of NLO calculations to parton shower and hadronization [228], and which can then be directly compared with the data at hadron-level.

PDF sensitivity

At leading order, jet production at hadron colliders includes the following subprocesses

$$\begin{aligned} gg \rightarrow gg, \quad gg \rightarrow q\bar{q}, \quad gq \rightarrow gq, \quad q\bar{q} \rightarrow gg, \\ q\bar{q} \rightarrow q\bar{q}, \quad q\bar{q} \rightarrow q'\bar{q}', \quad q\bar{q}' \rightarrow q\bar{q}', \quad qq \rightarrow qq, \quad qq' \rightarrow qq', \end{aligned} \quad (35)$$

along with the corresponding charge conjugate processes. Therefore, jet production is sensitive to both the gluon and quark PDFs, with the dominant partonic subprocess depending on the specific jet p_T . The kinematics of the two leading jets in the final state can be characterised by their rapidities $y_{(1,2)}$ and their transverse momenta $p_{T,(1,2)}$. At LO we have $p_{T,1} = p_{T,2} = p_T$, and the momentum fractions carried by the two incoming partons are given by

$$x_1 = \frac{p_T}{\sqrt{s}}(e^{y_1} + e^{y_2}), \quad x_2 = \frac{p_T}{\sqrt{s}}(e^{-y_1} + e^{-y_2}), \quad (36)$$

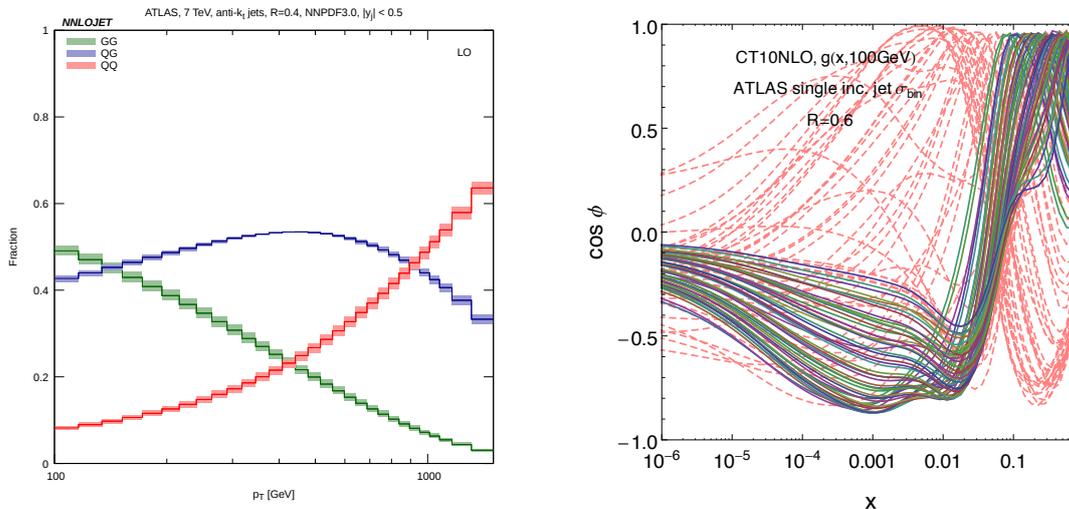


Figure 7: Left: Fractional contributions from different partonic channels to the single-inclusive jet production at the LHC 7 TeV at LO in the central rapidity region, computed with NNPDF3.0 [229]. Right: Correlations between binning cross sections from ATLAS on the single-inclusive jet production at the LHC 7 TeV and the gluon PDF; the dashed curves correspond to experimental bins at low p_T .

where \sqrt{s} is the centre of mass energy of the two incoming hadrons. If we instead consider the rapidity of the jet in the centre-of-mass frame of the dijet system, $y^* \equiv (y_1 - y_2)/2$, and the boost of the dijet $y_b \equiv (y_1 + y_2)/2$, we have

$$x_1 x_2 = \frac{4p_T^2 \cosh^2 y^*}{s}, \quad x_1/x_2 = e^{2y_b}. \quad (37)$$

Note that beyond LO there can be multiple jets in the final state from additional QCD radiation, so that in general the p_T balance of the two leading jets will be lost.

Experimentally, jet production can be measured in various ways. The most common observable for PDF fits is the single-inclusive jet cross section, double-differential in the jet p_T and rapidity y . Here, one counts all jets in a single event and includes them in the same distribution. Such a double-differential cross section is sensitive to different flavour combinations, depending on the kinematic region considered. In Fig. 7 (left) the fractional contributions from the different parton-level subprocesses to the inclusive jet cross section in the central rapidity region at the LHC is shown, as a function of the jet p_T . We can see that at low p_T the channels involving initial-state gluons are dominant, while at higher p_T the qq and $q\bar{q}$ contributions increase with the former one being dominant, but nonetheless with a sizeable gluon-induced fraction. As the quark PDFs are generally already well constrained by DIS data in these kinematic regions, jet data is therefore dominantly sensitive to the gluon PDF. This is illustrated in Fig. 7, which show the correlation coefficients (see Sect. 4.3.1) between the inclusive jet cross section and the gluon PDF at various x values. This follows the ATLAS binning [230], with each curve corresponding to one bin. From this we can see that the inclusive jet production can potentially constrain the gluon PDF in a wide range of x between $x \simeq 10^{-3}$ and $x \simeq 0.7$.

In addition to the single-inclusive case, there are also measurements of the double-differential cross sections for inclusive dijet production, that is with respect to y^* and the invariant mass of the two leading jets, or even triple-differential cross sections, e.g., with respect to y_b , y^* , and the average p_T of the two

leading jets. Through such refined binning one can probe different initial states more efficiently. The large y_b region usually receives more contributions from gluon initial states, while at large y^* and p_T initial states with two valence quarks dominate, allowing the d valence PDF at high x to be further constrained. We should mention that jets production in DIS at HERA can also provide complementary informations on the gluon PDF at large- x e.g., see Ref. [34] for most recent analysis. There are also theoretical progresses have been made on the NNLO QCD corrections to dijet production in DIS [231].

Experimental data

The currently available measurements on jet production at hadron collider which are relevant for constraining the PDFs are as follows:

- The double–differential single–inclusive jet production cross section data from the CDF [232, 233] and D0 [234, 235] collaborations, at Tevatron Run II (1.96 TeV).
- The double–differential single–inclusive jet production cross section data from the ATLAS [230, 236–238] and CMS [227, 239, 240] collaborations at LHC Run I (7 and 8 TeV).
- The double–differential inclusive dijet production cross section data from the ATLAS [230, 236, 241] and CMS [239, 242] collaborations at LHC Run I (7 and 8 TeV).
- The measurements of the ratio of double–differential cross sections in single–inclusive jet production at different centre–of–mass energies, 2.76, 7 and 8 TeV, from the ATLAS [237] and CMS [227] collaborations, at LHC Run I.
- The double–differential single–inclusive jet production cross section data from the CMS collaboration [243] at LHC Run II (13 TeV).
- More recently, measurements of triple differential dijet cross sections are becoming available, see e.g. the recent CMS analysis [244] at 8 TeV.

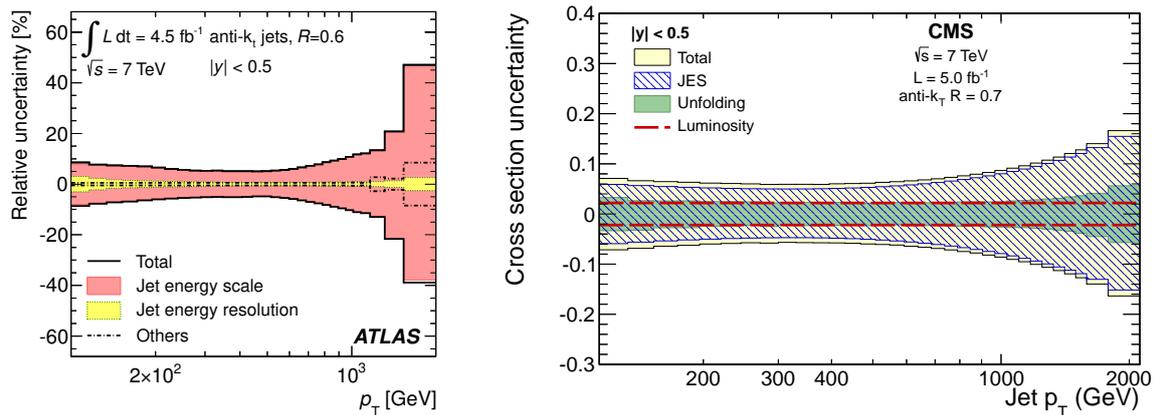


Figure 8: Representative systematic errors in single–inclusive jet measurement at LHC 7 TeV in the central rapidity region, from ATLAS (left) [238] and CMS (right) [240]. The luminosity error is not included in the ATLAS plot. We note that in both cases the jet energy scale (JES) is the dominant source of experimental uncertainties.

Due to the complexity of jet reconstruction and calibration there are a large number of experimental systematic uncertainties, with typically $N_{\text{sys}} \sim 50\text{--}100$ correlated systematic errors for $N_{\text{dat}} \sim 100\text{--}200$ data points, in the case of both ATLAS and CMS. In the most recent ATLAS and CMS 7 TeV measurements [238, 240], the total correlated experimental uncertainties are at the level of about $5 \sim 20\%$ in most regions. On the other hand, the uncorrelated systematic errors and statistical errors are at one percent level or less in general, and therefore the uncertainty on such data is generally completely systematics dominated. The typical experimental systematics from both ATLAS and CMS are shown in Fig. 8, and are seen to be dominated by the jet energy scale [238, 240]. The increasing precision of the LHC jet data, and the generally small uncorrelated errors, makes it rather challenging to fit the jet data well across the entire kinematic region in *e.g.* the case of the ATLAS 7 TeV measurement from the 2011 dataset [238]. A full account of these issues will almost certainly require a better understanding of both the experimental systematics and sources of theoretical errors that have not generally been included in PDFs fits previously.

In addition, for measurements of the ratios of the double-differential inclusive jet production cross sections at different centre-of-mass the experimental systematic errors largely cancel out, although the statistical uncertainties are somewhat larger, see *e.g.* the CMS 2.76, 7 and 8 TeV [227] and ATLAS 2.76 and 7 TeV [237] measurements. Such measurements of cross section ratios between different centre-of-mass energies \sqrt{s} are potentially useful for PDF fits [245], due to the cancellation not only of experimental systematics but also of theory errors such as scale variations.

Theoretical calculations and tools

The NLO QCD corrections to single-inclusive jet and inclusive dijet production were first calculated in the early 90's [246, 247], and has been implemented in two numerical programs, NLOjet++ [248, 249] and MEKS [250]. Recently, the NNLO QCD corrections to the same process have been completed for all partonic channels [42, 43, 251, 252], with the exception of some sub-leading colour contributions. The calculation is based on the Antenna subtraction method [253, 254] for isolating the infrared singularities in QCD real radiations, and is part of the NNLOJET program for the evaluation of NNLO hadron collider cross sections.

Fig. 9 (left) shows the NNLO QCD corrections to inclusive jet production at the 7 TeV LHC, with the anti- k_T algorithm and a central scale choice of the leading jet p_T . The NNLO QCD corrections are seen to be significant at low p_T , leading to a 10% increase with respect to NLO, while at high p_T the NNLO corrections are small. The NLO scale variations bands are asymmetric at low p_T and, interestingly, largely underestimate the perturbative uncertainties. Electroweak (EW) corrections can be significant at high p_T for central rapidities due to the presence of large EW Sudakov logarithms, but are well below 1% for a rapidity greater than unity [255].

There are ambiguities in choosing the appropriate QCD scale even in the simplest case of single-inclusive jet production [3, 256]. In particular, one can take either the p_T of the individual jet or the leading jet in the event. While these variables are the same at LO, where the two jets are produced back-to-back, at higher orders there can exist more than two jets which may have large differences in p_T . The NNLO predictions using these two choices for the central scales are studied in [229], and are found to lead to non-negligible differences specially at low values of p_T . This is shown in Fig. 9 (right), where at high p_T the two predictions converge as expected, but at low and intermediate p_T , there are significant differences of the central values in comparison to the size of scale variations. Indeed, the two error bands (from scale variations) do not even overlap. Although it seems that the NNLO predictions using the individual jet p_T as the central scale tend to follow the trends of ATLAS data better, clearly further investigations are needed to resolve the ambiguity of scale choice in the NNLO predictions, including specifically PDF fits based on exact NNLO theory for the jet data.

As well as fixed-order predictions, there are various theoretical calculations including analytic QCD

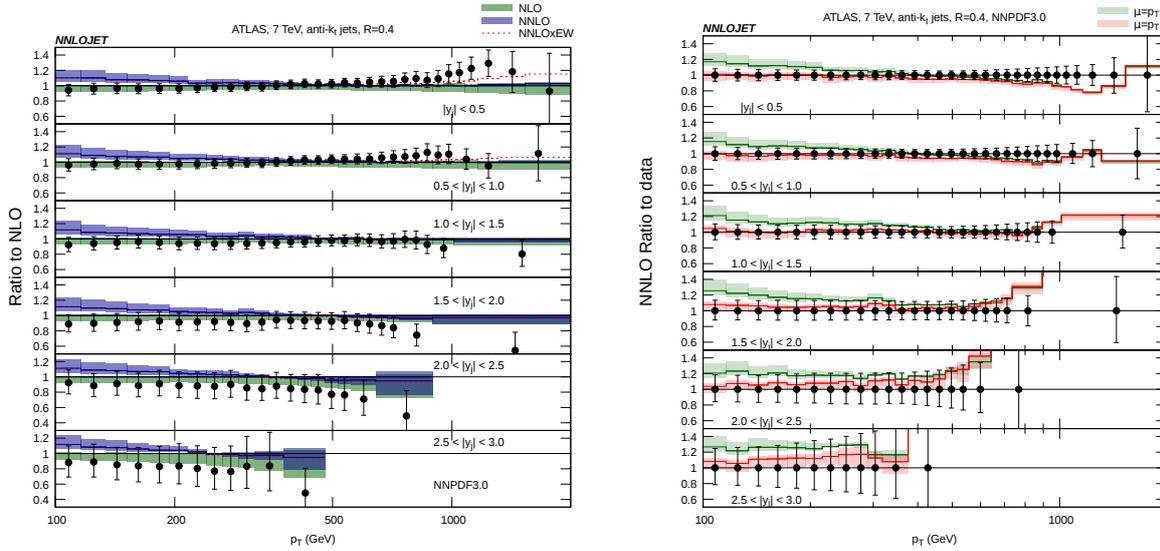


Figure 9: Left: Predictions for single-inclusive jet production at the 7 TeV LHC using ATLAS binning and the anti- k_T algorithm with $R = 0.4$. The central scale choice of the leading jet p_T is taken and the scale variations, shown by the shaded regions, are found by varying the renormalization and factorization scales up and down simultaneously by a factor of 2 [229]. Right: As before, but comparing the NNLO predictions using a central scale choice of the leading jet p_T (green) and the individual jet p_T (red) [229].

resummation [256–260]. It has been shown in [260] for the case of inclusive jet production at the LHC, that the approximate NNLO predictions from the expansion of threshold resummation agree well with the exact NNLO predictions for the all-gluon channel at large p_T . Over the full rapidity range the threshold expansion reproduces the fixed-order results down to a p_T of about 400 GeV with the same value shifted to lower p_T at forward rapidities. Such approximate NNLO predictions have been used in previous global analysis involving jet data [30, 32]. The jet cross sections are also sensitive to the jet algorithm used, in particular the value of the jet radius R . The use of a larger jet radius R at the LHC leads to a larger inclusive cross section and better convergence in the perturbative expansion. While this also reduces the non-perturbative corrections from QCD hadronization, it increases the effects from other non-perturbative phenomena such as the underlying event (UE) and multiple parton interactions (MPI). At the LHC, ATLAS uses $R = 0.4$ and $R = 0.6$, while CMS instead uses 0.5 and 0.7.

Impact on PDFs

Jet data from the Tevatron and the LHC Run I have already played an important role in global analyses [30–32, 261], although in the NNLO fits these only currently apply NLO or approximate NNLO theoretical predictions. Indeed, it was found that removing all jet data from the global analyses can lead to an increase of the gluon PDF uncertainties at large x by at least a factor of two [30], see Ref. [19] for an overview of such studies.

To illustrate the impact on PDFs of LHC inclusive jet data, in Fig. 10, taken from [244], we show the impact of the CMS 8 TeV jet data on the gluon PDF, by adding the data into a base fit with HERA DIS data only [34]⁴. The inclusion of both the single-inclusive jet data and the inclusive dijet data leads to a

⁴Such comparisons represent a non-trivial and useful indication of the impact of a given data set on the PDFs, and as such we

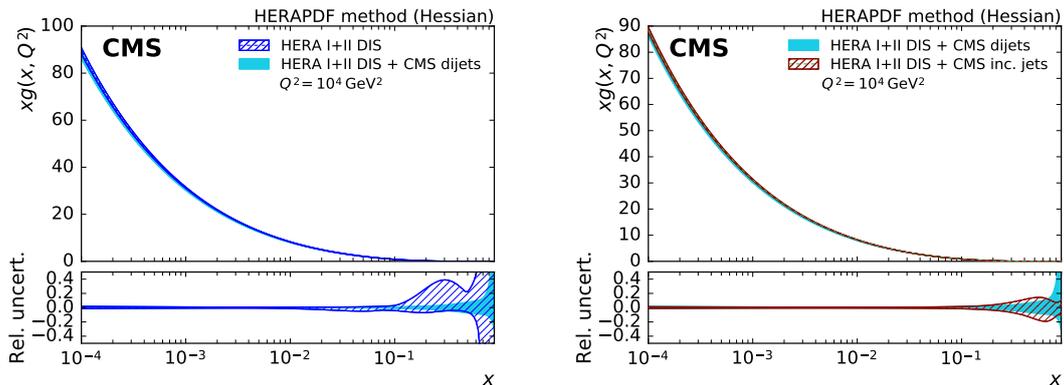


Figure 10: Effect of including the CMS 8 TeV jet data on the gluon PDF. Studies performed using `xFitter`, with PDF uncertainties obtained with the Hessian method. Left: Comparison of HERA DIS fit to that with the CMS dijet data included in addition. Right: Comparison of HERA DIS and CMS dijet fit to HERA DIS and CMS inclusive jet data fit [244].

sizeable reduction in the gluon PDF uncertainty at large x . Meanwhile in the same fit a reduction of the PDF uncertainty in the valence quark at high x is also observed, providing a complementary constraint to Drell–Yan and fixed–target DIS data. We note that such comparisons to HERA-only fits, presented here and in the following sections, provide an indication only of these data. A more complete analysis should compare instead to a full global fit, see Sect. 5.7.2 for further discussion.

With the full NNLO predictions on inclusive jet and dijet production now available, we can foresee significant advances in pinning down the gluon PDF at large x , in particular using the increasingly precise inclusive jet data from LHC Run I and Run II. Moreover, in the latter case, the increased luminosity will allow jet production to be probed deep into the TeV region, thus accessing PDFs at larger values of x where available constraints are rather scarce.

3.4. Inclusive gauge boson production

The inclusive production of electroweak gauge bosons, the so-called Drell–Yan process [262], has been of enormous historical importance, providing the first window on the quark flavour separation in the proton beyond that contained in DIS structure functions alone. These now provide the backbone of global PDF fits together with the fixed–target and HERA structure function data.

PDF sensitivity

The lowest–order contributions to W and Z/γ^* production proceed via the following partonic subprocesses:

$$\bar{u}\bar{d}, c\bar{s} \quad (u\bar{s}, c\bar{d}) \rightarrow W^+, \quad (38)$$

$$\bar{d}\bar{u}, s\bar{c} \quad (s\bar{u}, d\bar{c}) \rightarrow W^-, \quad (39)$$

$$q\bar{q} \rightarrow Z/\gamma^*, \quad (40)$$

will present examples of these throughout this review. Nonetheless, in general some caution is recommended when interpreting the quantitative conclusions of these studies, see Section 5.7 for further discussion.

where we show the Cabibbo suppressed contributions in brackets and where q corresponds to all active quark flavours. These processes can therefore tell us about the flavour decomposition of the proton, given that each flavour subprocess carries a different weight in the total cross section. To examine the dominant PDF sensitivity we can approximate the CKM matrix as diagonal, and thus ignore the bracketed contributions. In this case it is informative to consider the ratio of W^+ to W^- production, differential in the rapidity y_W of the produced boson [109],

$$R_{\pm} = \frac{d\sigma(W^+)/dy_W}{d\sigma(W^-)/dy_W} = \frac{u(x_1)\bar{d}(x_2) + c(x_1)\bar{s}(x_2) + 1 \leftrightarrow 2}{d(x_1)\bar{u}(x_2) + s(x_1)\bar{c}(x_2) + 1 \leftrightarrow 2}. \quad (41)$$

and the corresponding W asymmetry

$$A_W = \frac{d\sigma(W^+)/dy_W - d\sigma(W^-)/dy_W}{d\sigma(W^+)/dy_W + d\sigma(W^-)/dy_W} = \frac{u(x_1)\bar{d}(x_2) + c(x_1)\bar{s}(x_2) - d(x_1)\bar{u}(x_2) - s(x_1)\bar{c}(x_2) + 1 \leftrightarrow 2}{u(x_1)\bar{d}(x_2) + c(x_1)\bar{s}(x_2) + d(x_1)\bar{u}(x_2) + s(x_1)\bar{c}(x_2) + 1 \leftrightarrow 2}. \quad (42)$$

We will for simplicity consider the W rapidity, rather than the experimentally observable rapidity of the charged lepton from the W decay, in what follows. These variables are clearly correlated; we will comment further on this at the end.

Thus these ratios are in general sensitive to a fairly non-trivial combination of quark and anti-quark PDFs evaluated at the following values of x :

$$x_1 = \frac{M_W}{\sqrt{s}} e^{+y_W}, \quad x_2 = \frac{M_W}{\sqrt{s}} e^{-y_W}. \quad (43)$$

While these expressions completely define the PDF sensitivity of these observables at LO, it is informative to consider various kinematic limits, where these expressions simplify and more straightforward approximate dependences become apparent. Including only the (dominant) u and d contributions, we can in particular consider the cases of central and forward W production

$$\text{Central :} \quad y_W \sim 0 \quad x_1 \sim x_2 = x_0, \quad \bar{u}(x_{1,2}) \sim \bar{d}(x_{1,2}), \quad (44)$$

$$\text{Forward :} \quad y_W \gtrsim 2, \quad x_1 \gg x_2, \quad q(x_1) \sim q_V(x_1), \quad \bar{u}(x_2) \sim \bar{d}(x_2), \quad (45)$$

where $x_0 = M_W/\sqrt{s}$ and $q = u, d$, and we assume that $\bar{u} \sim \bar{d}$ at low x for demonstration here. At the LHC we have $x_0 = 0.005 - 0.01$, while in the forward region $x_2 \ll 1$, and therefore the $\bar{d} \sim \bar{u}$ approximation is a very good one. For the case of negative W rapidity we can of course simply interchange $x_1 \leftrightarrow x_2$.

In the central region, applying the simplification of Eq. (44) and dropping the c, s contributions we find

$$R_{\pm} \sim \frac{u(x_0)}{d(x_0)}, \quad (46)$$

$$A_W \sim \frac{u_V(x_0) - d_V(x_0)}{u(x_0) + d(x_0)}. \quad (47)$$

Thus A_W is sensitive to the valence difference, while R_{\pm} is sensitive to the ratio of u to d at $x_1 \sim x_2 \sim x_0$. For these reasonably low x values, the valence u and d quarks are fairly small, and so we roughly expect $R_{\pm} \sim 1$ and $A_W \sim 0$, with the departures from these values being due to the precise flavour content of the proton, in particular the fact that the valence distributions are not completely negligible in this region.

In the forward region, applying the simplification of Eq. (45) and again dropping the c, s contributions we find that the following relations hold

$$R_{\pm} \sim \frac{u_V(x_1)}{d_V(x_1)}, \quad (48)$$

$$A_W \sim \frac{u_V(x_1) - d_V(x_1)}{u_V(x_1) + d_V(x_1)}. \quad (49)$$

Thus these provide (equivalent) sensitive constraints on the u/d ratio at high x .

Considering now the case of Z production, for forward production we find, using the same approximations as before, that the differential ratio of W to Z production is given by

$$\frac{d\sigma(W^+)/dy_W + d\sigma(W^-)/dy_W}{d\sigma^Z/dy_Z} \approx \frac{u_V(x_1) + d_V(x_1)}{0.29u_V(x_1) + 0.37d_V(x_1)}, \quad (50)$$

where the factors in the denominator arise from the electroweak Z -quark couplings. For the central region a similar result, evaluated at x_0 , is found up to an overall factor of 2. Therefore, we can see that the W^{\pm} and Z cross sections provide very similar information about the u and d quarks.

Up to this point we have omitted the contribution from the strange quarks to W and Z production. Generally speaking this is washed out when considering ratio observables, justifying their omission above, although the W asymmetry displays some sensitivity to the strange difference $s - \bar{s}$. On the other hand, the contribution to the absolute cross sections is not negligible, in particular at lower x . Thus for example the Z cross section at central rapidity becomes, for five active quark flavours,

$$\frac{d\sigma^Z}{dy_{\parallel}} \sim 0.29(u(x_0)\bar{u}(x_0) + c(x_0)\bar{c}(x_0)) + 0.37(d(x_0)\bar{d}(x_0) + s(x_0)\bar{s}(x_0) + b(x_0)\bar{b}(x_0)). \quad (51)$$

Therefore, provided the absolute cross section data are sufficiently accurate and the other quark flavours are sufficiently well determined, this may for example be sensitive to the currently less well determined strange quark distribution. Moreover, this is not a case of a simple overall normalization; as the Z rapidity increases the valence u, d contributions will become increasingly dominant, and the contribution from the strange (and the heavy flavours) will decrease. Thus the shape of the Z rapidity distribution is sensitive to the proton strangeness, as well as to that the contribution from heavy flavour PDFs. Similar considerations also apply for the absolute cross sections, where in particular the ratios of total W^{\pm} and Z cross sections at the LHC are known to be quite sensitive to the strangeness PDF [263].

Below the Z peak region, the Drell–Yan process is dominated by an off-shell intermediate photon, with cross section given by

$$\frac{d\sigma^{\text{DY}}}{dy_{\parallel}} \sim \sum_i e_i^2 (q(x_1)\bar{q}(x_2) + q(x_2)\bar{q}(x_1)), \quad (52)$$

where the sum runs over all active quark flavours. Thus in comparison to Eq. (51), a different combination of the quark and anti-quark PDFs is probed, due to the differing electromagnetic couplings. In particular, the relative $u\bar{u}$ to $d\bar{d}$ contribution is now a factor of ~ 5 higher in comparison to the Z cross section. At the LHC, low mass Drell–Yan production therefore provides complementary flavour information in the low to intermediate x region. In addition, as the cuts on the final-state lepton transverse momenta tend to increase the relative importance of the higher order contributions, for which the Z p_{\perp} can be non-zero, this can be sensitive to the gluon PDF at lower x , which contributes through the NLO $g \rightarrow q\bar{q}$ splitting. On the other hand, high mass Drell–Yan production (with $M_{\parallel} \gg m_Z$) is sensitive to the q, \bar{q} PDFs at high x , in particular the anti-quarks, which are rather less well determined in this region.

A further constraint is provided by considering the Drell–Yan process on fixed proton and neutron (in practice, deuteron) targets. By using isospin symmetry the PDFs between the proton and the neutron can be related

$$u^p = d^n \quad d^p = u^n , \quad (53)$$

allowing an extra handle on the proton flavour decomposition. In particular, such fixed target experiments generally have larger acceptance in the $x_1 \gg x_2$ region (where x_1 is defined with respect to the proton beam) for which the first term in Eq. (52) is dominant, with $q(x_1) \sim q_V(x_1)$. It is then straightforward to show that

$$\frac{\sigma^{pn}}{\sigma^{pp}} \sim \frac{\bar{d}(x_2)}{\bar{u}(x_2)} . \quad (54)$$

That is, they are sensitive to quark sea decomposition in the intermediate to high $x_2 \sim 0.01 - 0.3$ region probed by these fixed target experiments [264]. This however comes with the added complication that the cross section must be corrected to account for the fact that the neutron is bound in a deuteron nucleus, and therefore the ‘free’ neutron PDF is not directly probed. Fixed–target pp scattering does not suffer from this issue, and is sensitive to the quark sea (dominantly, the \bar{u}) in the same x region, but is much less directly sensitive to the \bar{d}/\bar{u} decomposition.

Turning now to the case of W, Z production at the Tevatron, the fact that we have $p\bar{p}$ collisions affects the flavours probed. In particular, we can use charge–conjugation symmetry to write

$$q^p = \bar{q}^{\bar{p}} . \quad (55)$$

In fact, it is straightforward to show that in the region of valence quark dominance, the cross section ratio R^\pm and the asymmetry A_W are again sensitive to the u/d ratio and the valence difference $u_V - d_V$, while the Z cross section again provides similar information to W^\pm cross section sum. Nonetheless, these conclusions are only approximately true, and the presence of a \bar{p} beam provides complementary flavour information.

Finally, we emphasise that we have considered here the distributions with respect to the (unobservable) rapidity of the W boson to simplify the discussion. In general we should correctly account for the kinematics, as well as weight the corresponding $q\bar{q}$ contributions by the appropriate W decay distributions. This in fact provides a further handle on the flavour sensitivity of this observable, as by changing the p_\perp cut on the charged lepton, different weights of the different quark contributions are achieved, see *e.g.* [124, 265] for further details. Nonetheless, the forward and central W rapidity regions are certainly correlated with the equivalent lepton rapidity regions that are measured experimentally, and so the above discussion provides a qualitative guide for the PDF sensitivity of W boson production. However, as we will discuss below, the current simulation codes for W and Z production include the full kinematics of the leptonic decays, and therefore there is no need to explicitly correct back to the W rapidity.

Experimental data

A non–exhaustive list of the available Drell–Yan measurements is as follows:

- The most precise fixed target Drell–Yan data come from the E866/NuSea [266] experiment at Fermilab, while the E906/SeaQuest experiment [267] will extend out to higher x , and is currently taking data.
- The Tevatron collider has produced a range of data on W and Z production, including measurements of the Z rapidity distribution [268, 269] and W production for both the lepton [270, 271] and the W [272, 273] asymmetries.

- Early LHC measurements of the Z rapidity distribution presented by CMS [274] and ATLAS [275].
- CMS Drell–Yan data at 7 TeV [276], for $15 < M_{ll} < 1500$ GeV and at 8 TeV [277], which increased the upper mass limit to 2000 GeV. These are presented double differentially in the rapidity and invariant mass of the lepton pair.
- ATLAS 7 TeV Drell–Yan invariant mass distribution (integrated over rapidity) at high [278] ($116 < M_{ll} < 1500$ GeV) and low [279] ($26 < M_{ll} < 66$) invariant masses.
- CMS [280, 281] 7 TeV W asymmetry, and ATLAS [275] W^+ and W^- cross section data.
- ATLAS high precision W and Z, γ^* data [282], using the full 4.6 fb^{-1} data set at 7 TeV. The Drell–Yan rapidity distribution is presented double differentially in three intervals of lepton pair invariant mass, over the $45 < M_{ll} < 150$ GeV range. In the Z peak and higher mass regions the measurement was also extended out to $|y_{ll}| = 3.6$.
- LHCb Z rapidity distributions at 7 [283, 284], 8 [285] and 13 [286] TeV.
- LHCb lepton rapidity distributions for W^+ and W^- production at 7 [287], and 8 [288, 289] TeV.
- In Ref. [288], cross section ratios between the 7 and 8 TeV LHCb W and Z measurements are presented, with the cancellation in various systematic uncertainties providing a more precise PDF sensitivity.

Thus at the LHC numerous measurements have already been presented, with many others in preparation. The Z/γ^* data are available over a wide range of invariant masses, providing extensive coverage in x . The W production data are increasingly presented as individual cross sections, including the correlated error information, to provide the maximum possible constraints. While the majority of cases, such as with the ATLAS and CMS measurements, are limited to the central rapidity region, that is, a lepton pseudorapidity of $|\eta_l| < 2.4$, this reach is extended by exploiting the forward acceptance of the LHCb detector, for which $2 < \eta_l < 4.5$ is accessed. This allows the high and low x region to be probed.

The most recent ATLAS W and Z, γ^* data [282], which uses the full 4.6 fb^{-1} data set at 7 TeV demonstrates the level of precision that is now being achieved. The Z rapidity distribution and W asymmetry are shown in Fig. 11 (in the latter case the individual W measurements are available). The high experimental precision is clear, in particular in the Z distribution where excluding the luminosity uncertainty it is as low as $\sim 0.3\%$ at central rapidity, while the error on the PDF predictions, as well as the spread between sets, is significantly larger.

Theoretical calculations and tools

W and Z boson production is arguably the simplest process one can consider at a hadron collider, and indeed it was the first hadroproduction process for which the NNLO calculation became available, with the total cross sections being calculated in the early 90s [141]⁵. A decade later, in [291, 292] the NNLO corrections to the differential W and Z rapidity distributions was presented for the first time. A more direct comparison with experimental observables was provided in [293, 294] which presented the NNLO calculation fully differential in the final–state leptons, including in addition spin correlations, finite width effects and $\gamma - Z$ interference. This was accompanied by the public release of the FEWZ simulation

⁵A mistake in the one–loop real emission contribution was reported in [290].

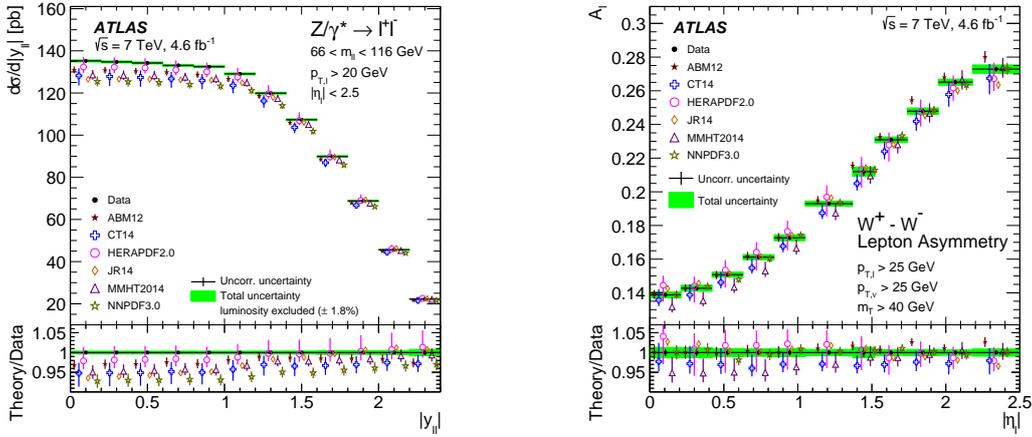


Figure 11: The ATLAS 7 TeV measurements of Z rapidity distribution (left) and W asymmetry (right) based on $\mathcal{L} = 4.6 \text{ fb}^{-1}$, compared with different NNLO PDF sets. Taken from [282].

code, with subsequent improvements reported in [295, 296] and [297], where NLO EW corrections (first calculated in [298–303]) were included. The DYNMLO [304] parton-level MC provides an alternative tool for generating W and Z production, again including spin correlations, finite width effects, and γ – Z interference, but currently without EW corrections. This code allows for arbitrary user-defined cuts on the final-state partons and leptons to be imposed and histograms to be made, in contrast to FEWZ, where a selection of pre-determined cuts and histograms may be applied.

These two codes differ in their theoretical treatment of the processes, in particular in the method that is applied to achieve the (non-trivial) cancellation of IR singularities at intermediates steps in the calculation. While FEWZ uses the local ‘sector decomposition’ method [293, 294] that provides an automated method for extracting and canceling the IR poles, DYNMLO applies an alternative non-local ‘ q_T -subtraction’ approach [305] which uses the transverse momentum q_T of the produced W or Z as a cut variable, treating the calculation in a different way above and below some q_T^{cut} .

It is important to emphasize that, at the time of writing this Report, these two codes can still lead to non-negligible differences in their predictions for the same cross sections even for identical input parameters. For example, in the recent ATLAS high precision W and Z/γ^* analysis [288] the difference in the fiducial cross section predictions can be as high as $\sim 1\%$, that is larger than the experimental uncertainties. This is due to the differing subtraction procedures, which affects the predicted boson p_\perp distributions. The differences between the predicted cross sections are generally more significant when more restrictive cuts on the final-state leptons are imposed; for the total W , Z cross sections the codes agree to within 0.2% [288]. A closer investigation of this issue and its impact on PDF determination will clearly be essential, given the that the experimental uncertainties of existing and future measurements is at the permille level.

More recently the MCFM event generator [306] has extended the NLO simulation of W and Z production to NNLO [307]. This takes a similar non-local approach to DYNMLO, but using the N -jettiness variable rather than the q_T . Here, it is shown that a careful and process-dependent choice of the cut on the 0-jettiness variable, τ_0^{cut} (the equivalent of q_T^{cut} above) is required in order to balance the requirements of sufficient statistical precision and control over systematic power corrections that increase in importance as this cut is increased.

Finally, public codes including transverse momentum resummation are also available. In this respect,

the DYRes [308] program combines NNLO fixed-order with NNLL resummation, while ResBos [309] combines NLO fixed-order with NNLL resummation. However, typical observables that are used in PDF fits are chosen to be largely insensitive to such resummation effects, which are most important as the W , Z transverse momentum becomes small, and so these codes are not used in most PDF fits. The exception are the CT10 and CT14 global analysis, for which the theoretical predictions for W^\pm and Z production are based on ResBos, with the motivation that the NLO+NNLL result is very close to the full NNLO calculation. Moreover, this choice could be advantageous since the lepton charge asymmetry may still be sensitive to q_T resummation effects [310] for some specific choices of the cut p_t^l in the charged lepton transverse momentum.

Impact on PDFs

As described above, inclusive gauge boson production has played a crucial role in determining the quark flavour decomposition of the proton. Indeed, these have been included in all major PDF analyses for some time, from earlier fixed target data through to measurements at the Tevatron and increasingly at the LHC. Two recent LHC results are shown in Fig. 12. In the left panel we show the CMS fit [311] to the down valence quark distribution. The baseline fit is to the HERA I+II data only, which is compared to the result including the CMS 8 TeV W boson production data. The change in shape and sizeable reduction in the PDF uncertainty over a wide range of x is clear.

In the right panel we show the impact on the strange quark fraction relative to the light quark sea

$$R_s = \frac{s + \bar{s}}{\bar{u} + \bar{d}}, \quad (56)$$

of the ATLAS high precision W and Z/γ^* data [282]. As described in Section 3.4, provided the light quark flavours are sufficiently well determined, and the data are sufficiently precise, the size and shape of the W, Z rapidity distributions can provide constraints on the strange quark PDFs. This is clear from the figure, where the fit to the ATLAS data predict a significantly higher value of R_s in comparison to previous PDF fits, which do not include the ATLAS data. A hint of this effect is seen in the earlier ‘ATLAS–epWZ12’ result [312], but it is only with the more recent high precision data that a clear effect becomes apparent. The size and shape of the strange PDF will be further discussed in Sect. 6.3.

As mentioned above, LHCb data on forward weak boson production provides important information about the quarks and anti-quarks flavour separation in the large- x region, beyond the kinematic coverage of the ATLAS and CMS measurements. To illustrate this, in Fig. 13 we show the down quark PDF at $Q = 100$ GeV, comparing the results of the baseline NNPDF3.1 fit with those of the same fit but with all LHCb data excluded [313]. We find that the LHCb data tends to enhance the value of $d(x, Q^2)$ by almost one-sigma in the $x \simeq 0.3$ region, with a marked decrease of PDF uncertainties. The impact is in fact found to be significant for all quark PDFs in this region, highlighting the important role that such data plays.

3.5. The p_T distribution of Z bosons

The LHC has provided precision measurements of the inclusive transverse momentum spectra of the Z boson, which may be exploited for the purposes of PDF fitting, as well as for other applications such as the tuning of Monte Carlo event generators. There are three distinct regions of the p_T spectrum. At small $p_T \ll m_Z$, the fixed-order predictions diverge due to the higher-order logarithms generated by soft gluon radiation. Here, QCD resummation techniques are needed in order to achieve reliable predictions, see [314–321]. Such predictions require additional non-perturbative input that cannot be calculated from first principles [318, 322, 323], and therefore the Z p_T distribution cannot be reliably used for the extraction

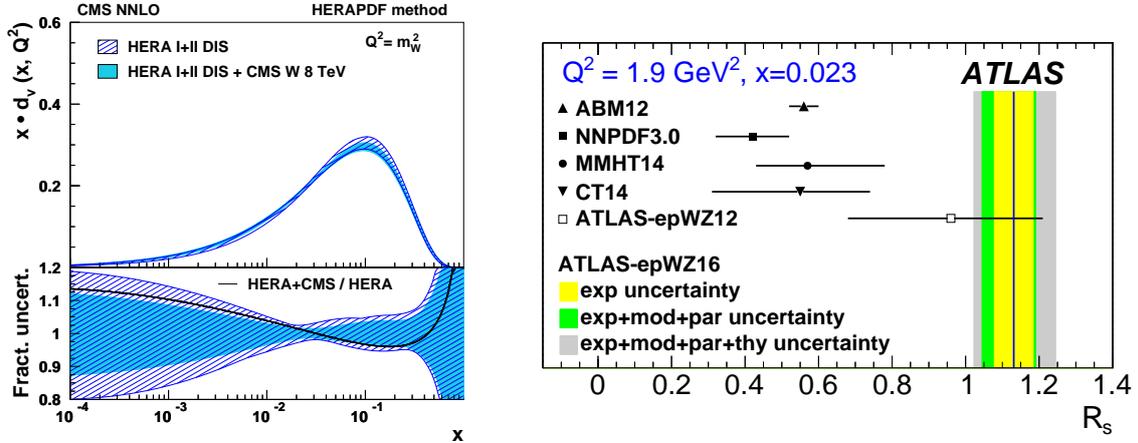


Figure 12: (Left) The $d_V(x, Q^2)$ PDF, with the CMS fit to HERA I+II data only, and including the CMS 8 TeV W production data. Taken from [311]. (Right) The ratio of the strange quark to the light quark sea, R_s , with different PDF predictions and the result of the ATLAS fit to HERA I+II data and the high precision W and Z/γ^* data shown. Taken from [282].

of the collinear PDFs in this region. Moreover, in this region, the distribution is highly correlated with and provides no additional sensitivity relative to the inclusive Z boson production cross section.

On the other hand, at large $p_T \gg m_Z$, the fixed-order predictions can also receive large logarithmic contributions due to soft gluon radiation at the partonic threshold of the Z boson and the recoiling jet [324, 325]. It has been shown that those contributions can increase the integrated cross sections with $p_T > 200$ GeV by $\sim 5\%$ compared to the NLO prediction at the LHC [324]. The third region is defined by intermediate $p_T \sim m_Z$ values, where the standard fixed-order predictions can be trusted. Note that resummed/matched calculations can also be reliably used in this intermediate p_T region, in particular since the effect of the non-perturbative effects will not be strong enough to affect determination of PDFs, because the resummed cross section is dominated by the perturbative logs even at $p_T = 0$. Therefore, is in this region that the p_T distribution of Z bosons can provide additional constraints on the PDFs, in particular on the gluon.

PDF sensitivity

At leading order, Z boson production with finite transverse momentum includes the following partonic subprocesses

$$q\bar{q} \rightarrow Zg, \quad gq \rightarrow Zq, \quad g\bar{q} \rightarrow Z\bar{q}. \quad (57)$$

In the leptonic channel where experimental measurements are the cleanest, the kinematics of the Z boson, namely the transverse momentum p_T and rapidity y_Z , can be reconstructed from the momenta of the lepton pair produced in the Z decay. The momentum fractions of the initial-state partons are given by

$$x_1 = \frac{m_T}{\sqrt{s}} e^{y_Z} + \frac{p_T}{\sqrt{s}} e^{y_j}, \quad x_2 = \frac{m_T}{\sqrt{s}} e^{-y_Z} + \frac{p_T}{\sqrt{s}} e^{-y_j}, \quad (58)$$

where \sqrt{s} is the centre-of-mass energy of the two incoming hadrons, $m_T = \sqrt{M_Z^2 + p_T^2}$ is the transverse mass of the Z boson and y_j is the rapidity of the recoiling parton. For inclusive production with respect to the hadronic recoil, that is integrating over y_j , these momentum fractions are therefore not uniquely determined, although for LO kinematics lower limits can be derived. Typically, the LHC experiments

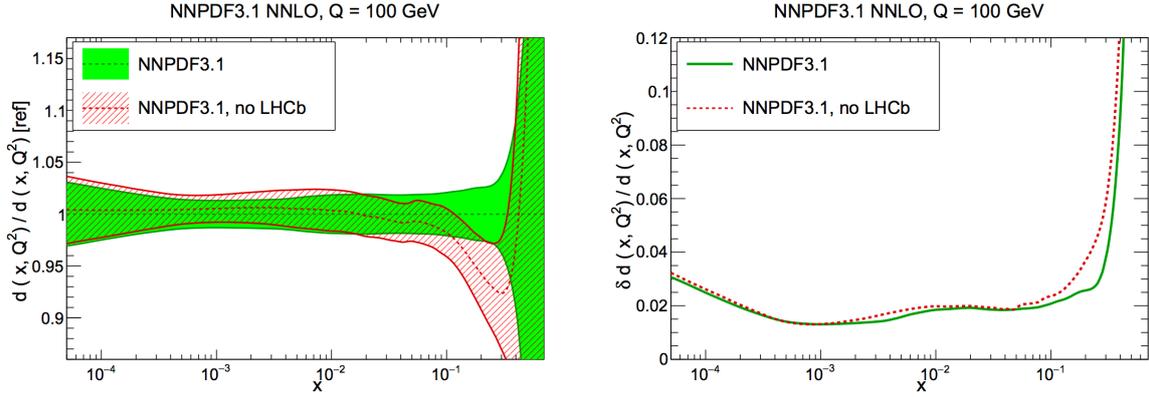


Figure 13: The down quark $d(x, Q^2)$ at $Q = 100$ GeV, comparing the results of an NNPDF3.1 baseline fit with those of the same fit without LHCb data [313]. We show the PDF ratio normalized to the central value of NNPDF3.1 (left) and the relative PDF uncertainties (right).

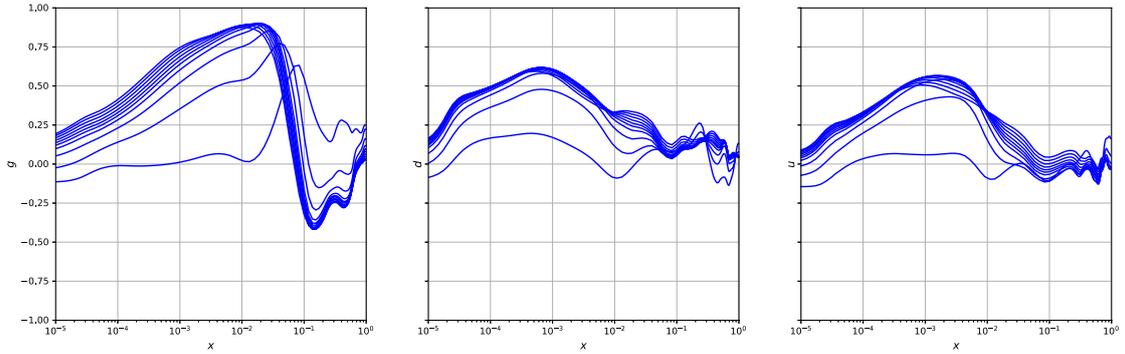


Figure 14: Correlations coefficients between the cross sections in various p_T bins and the gluon, down and up quark PDFs as a function of x , from Ref. [326]. The binning corresponds to the ATLAS 8 TeV measurement [327] within the rapidity interval $0 < |y_Z| < 0.4$.

measure the double-differential cross sections in p_T and y_Z at the Z peak, although the off-shell region, where the contributions from virtual photon may be important, can also be considered.

The cross sections at moderate and large transverse momentum are dominated by contributions from gluon and quark scattering and are strongly correlated with the gluon PDF in the region relevant to Higgs boson production at the LHC. This is illustrated in Fig. 14, which shows the PDF-induced correlations between the cross sections in different p_T bins, with $0 < |y_Z| < 0.4$, and the gluon, down and up quark PDFs at various x values [326]. We can see that indeed the correlations with the gluon at $x \sim 10^{-2}$ almost reach 0.9. Moderate correlations with the quark PDFs at $x \sim 10^{-3}$ are also observed. This highlights the fact that the Z p_T distributions provides in principle a handle on the gluon in a region of x that lies between that covered by HERA structure functions, at smaller values of x , and that covered by inclusive jets and $t\bar{t}$ production, at larger x .

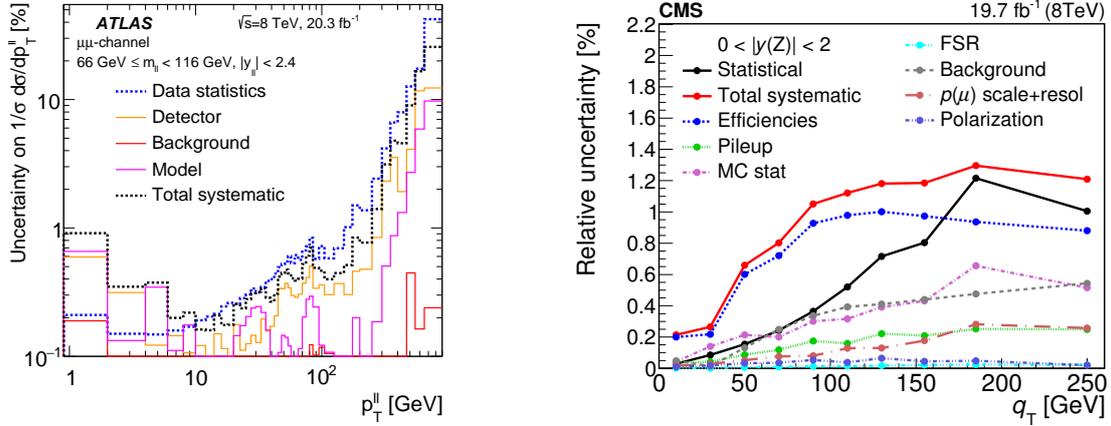


Figure 15: Left plot: the decomposition of the various experimental statistical and systematic uncertainties on the normalized p_T spectra of the Z boson in dimuon channel from ATLAS at 8 TeV [327]. Right plot: the corresponding experimental uncertainties for the normalized p_T spectra of the Z boson in dimuon channel from CMS at 8 TeV [328].

Experimental data

The experimental precision for the measured Z p_T distribution has reached the percent level for both ATLAS [327] and CMS [328] at the LHC Run I, due to both the clean dilepton final state as well as the high statistics of the signal. The ATLAS measurement extracts the cross sections at three different (‘Born’, ‘bare’ and ‘dressed’) particle levels in terms of the effect of final–state photon radiation. The Born and bare levels are defined from the lepton kinematics before and after final–state QED radiation, while the dressed level is defined by further combining the momentum of the lepton with photons radiated within a certain cone. These distributions can be presented with the Z boson rapidity integrated over, or separated into different rapidity intervals, and can be on or off the Z peak. In addition, measurements of the distributions with respect to the angular variable ϕ_η^* [327], which is proportional to $p_{T,Z}$ at small transverse momentum, are also available. As ϕ_η^* only depends on the direction of the lepton momenta, which are better measured than the momenta themselves, this allows the experimental systematics to be reduced.

We summarise the available measurements on p_T spectra of the Z boson relevant to constraining the PDFs below:

- The normalized Z p_T distribution in different rapidity intervals by the ATLAS collaboration [329, 330] at LHC Run I (7 TeV).
- The normalized and unnormalized distributions of lepton pairs with respect to p_T or ϕ_η^* in different rapidity intervals by the ATLAS collaboration [327] at LHC Run I (8 TeV).
- The normalized Z p_T distribution integrated over rapidity by the CMS collaboration [274] at LHC Run I (7 TeV).
- The normalized and unnormalized double–differential Z distribution in p_T and rapidity by the CMS collaboration [328] at LHC Run I (8 TeV).
- The normalized Z p_T distribution and the ratio to the W p_T distribution by the CMS collaboration [331] at LHC Run I (8 TeV).

- The unnormalized Z distribution in ϕ_η^* in the forward region by the LHCb collaboration [283, 285] at LHC Run I (7 and 8 TeV).
- The normalized Z distribution in p_T or ϕ_η^* in the forward region by the LHCb collaboration [284] at LHC Run I (7 TeV).
- The unnormalized distribution of the lepton pair with respect to ϕ_η^* in different rapidity intervals by the D0 collaboration [332] at Tevatron Run II (1.96 TeV).

A summary of the experimental uncertainties for the ATLAS and CMS 8 TeV measurements [327, 328] of the normalized Z p_T distribution is shown in Fig. 15. The luminosity uncertainty and some of the systematic errors largely cancel in the normalized distributions. Both ATLAS and CMS have measured the Z p_T up to about 1 TeV, while ATLAS has a finer binning at small p_T , motivated by the usefulness of low p_T measurements for MC event generator tuning.

For ATLAS, the total systematic uncertainty is well within 1% for p_T smaller than 200 GeV and is $O(10\%)$ at the higher p_T tail. The statistical uncertainties are tiny, starting at $\sim 0.2\%$ for $p_T \sim 10$ GeV, and are within 1% in most of the relevant p_T region for both ATLAS and CMS. The PDF uncertainties from individual PDF groups are about 2%, which is already larger than the experimental errors in general, even before considering the spread between different PDFs; such data can therefore provide valuable PDF constraints. We emphasise that the high precision of these measurements is an unprecedented challenge for QCD calculations, which need to achieve a permille level accuracy to match the experimental precision.

Theoretical calculations and tools

The NLO QCD corrections to the Z p_T distribution were calculated a long time ago [333–335], while more recently the EW corrections have been studied extensively [336–339]. The NLO QCD corrections are found to be sizeable at LHC energies, and large QCD scale variations are found in the predicted p_T spectra, rendering such predictions unsuitable for PDF determination. However, the NNLO QCD corrections have recently been calculated by two independent groups, in one case using the antenna subtraction method [46, 253, 340] and in the other the N -jettiness subtraction method [341, 342]; these are found to be in good agreement. While the original calculations are for Z +jet production, these can readily be translated to the case of inclusive production of Z boson for larger values of p_T .

At NNLO, the theoretical uncertainty due to the QCD scale variation is found to be greatly reduced, allowing the measurements of the Z boson p_T spectra to be included for the first time for precision PDF determination. Moreover, these calculations include the leptonic decays of the Z boson and thus parton-level selection cuts may be applied to the theoretical predictions, allowing a direct comparison with the measured fiducial cross sections, without relying on any experimental phase space extrapolation. To illustrate this, Fig. 16 (taken from [46]) shows the NLO and NNLO predictions for the unnormalized and normalized Z boson p_T spectra at the 8 TeV LHC. The central values of the renormalization and factorization scale are set to the transverse mass of the Z boson, with scale variations calculated by varying these simultaneously by a factor of 2 up and down. The NNLO corrections are moderate for the unnormalized distribution, being about 5% at low p_T and 9% at high p_T . The remaining scale variations range from 1% to 6% depending on value of p_T . The EW corrections are small at moderate transverse momentum but can be sizeable in the tail region, reaching around $\sim 10\%$ for p_T values greater than 600 GeV. However, the statistical errors in the tail region are currently quite large, preventing a direct probe of these EW effects. For the normalized distribution, the denominator corresponds to the inclusive Z production cross section at NNLO in same fiducial region and with independent scale variations, namely the uncertainties estimated for the numerator and denominator are considered as uncorrelated. That is due to the fact that inclusive cross section receives

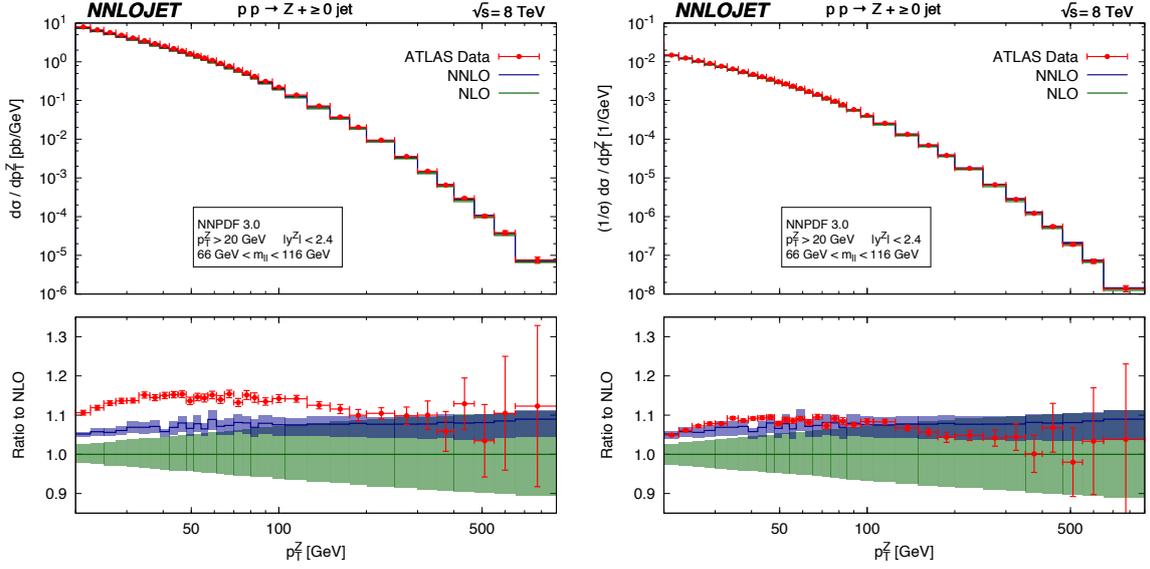


Figure 16: Left (right): the absolute (normalized) transverse momentum distribution of the inclusive Z boson production at LHC 8 TeV [46] computed with NNPDF3.0, compared with the corresponding ATLAS measurements. The green and blue bands indicate the NLO and NNLO predictions with the corresponding theoretical uncertainties from scale variations. The fiducial cuts on the final state charged leptons are $p_{T,l} > 20$ GeV and $|\eta_l| < 2.4$.

dominant contributions from small p_T region and is calculated upto one order lower in α_S . The size of the QCD corrections are found to be similar to the unnormalized case because the NNLO corrections to the inclusive Z production cross section are very small.

A detailed phenomenological study and comparison of the NNLO calculation to the ATLAS and CMS 8 TeV measurements has been presented in [46, 340]. Good agreement between the NNLO theory and data for the normalized distribution in ranges from $20 \text{ GeV} < p_T < 900 \text{ GeV}$, in all rapidity intervals, is observed. However, there is some discrepancy in the comparison to the ATLAS unnormalized distributions for using NNPDF3.0 PDFs, see Fig. 16, with the data tending to overshoot the theory over a wide p_T range. The agreement in normalization is better if using NNPDF3.1 PDFs, of which the ATLAS data was included in the fit, or using most recent PDFs from other major PDF groups. On the other hand, the NNLO prediction for the shape of the p_T distribution is in good agreement with the data down to a p_T value of 4 GeV, and is largely improved in comparison to the NLO predictions.

Impact on PDFs

The impact of the Z boson p_T data at LHC Run I has recently been studied within a global analysis framework [169, 326]. Here, an additional uncorrelated theoretical uncertainty of $\sim 1\%$ has been added to all p_T bins to account for the theoretical uncertainty due to the residual MC integration error in the NNLO calculations. Without including these errors, it was found that NNLO predictions can not describe the data well, especially in the case of the normalized distributions. Some tension is also found between the ATLAS 7 TeV normalized p_T distribution [330] and the 8 TeV p_T distribution from both ATLAS and CMS [327, 328]. The ATLAS 7 TeV data also pulls the PDFs in a very different direction with respect to the HERA inclusive DIS data [326]. In [169] it is concluded that the inclusion of the ATLAS 7 TeV normalized data in the global analyses does not appear to be justified.

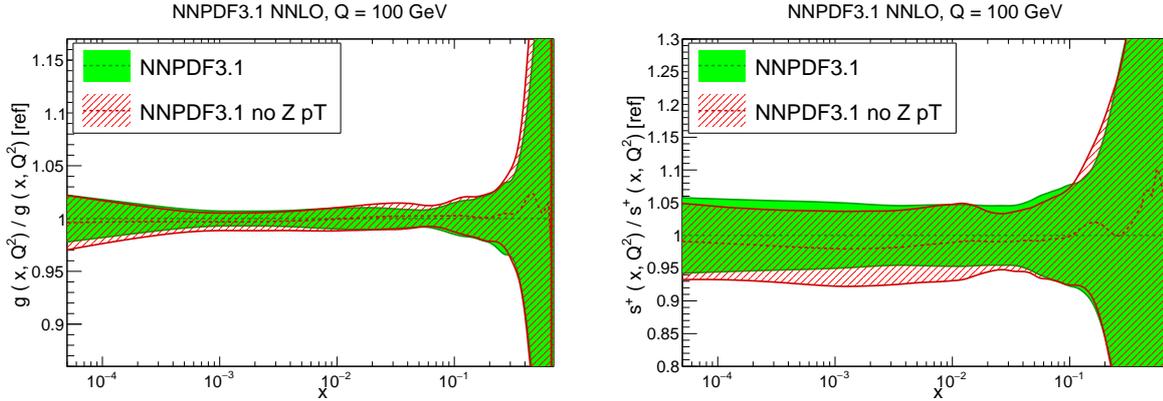


Figure 17: Impact of the Z boson transverse momentum measurements from ATLAS and CMS 8 TeV on the gluon PDF and the total strangeness in the NNPDF3.1 global analyses [169].

In Fig. 17 the impact of the ATLAS and CMS 8 TeV data on the NNPDF3.1 global analyses [169] is shown, by comparing the changes of the PDFs in the analysis with and without the Z p_T data sets included. The uncertainty in the gluon PDF is seen to be slightly reduced in the x region of $10^{-2} \sim 10^{-1}$. In the same region the gluon PDF receives constraints from precision measurements on top quark pair production and HERA inclusive DIS, both of which are present in the same analysis. It is also found that the 8 TeV data lead to a moderate reduction in the PDF uncertainty on the total strangeness.

It is important to highlight that in Fig. 17 the impact of the ATLAS and CMS Z p_T data is rather moderate due to the fact that many other datasets in the fit already constrain the PDFs, and in particular the gluon, in the relevant x region. As shown in [326], the effects of these measurements is rather more marked if the input dataset does not contain other data such as jets or top quark pair production data that also pin down the gluon, see Sect. 6.1 for further discussion.

3.6. Direct photon production

In this section we discuss the PDF constraints that can be derived from isolated prompt photon production at hadron colliders.

PDF sensitivity

The leading parton-level processes for direct (also known as ‘prompt’) photon production, that is, where the photon is produced by point-like emission from a quark, are given by

$$\text{QCD Compton : } \quad qg \rightarrow q\gamma, \quad (59)$$

$$\text{Annihilation : } \quad q\bar{q} \rightarrow g\gamma. \quad (60)$$

The QCD Compton process gives the dominant contribution, in particular at the LHC, and is directly sensitive to the gluon PDF. For LO kinematics, the momentum fraction carried by the incoming gluon is directly proportional to the transverse energy E_\perp^γ of the produced photon, and thus for higher E_\perp^γ this process provides a direct probe of the gluon PDF at high x . Moreover, this represents the highest rate electroweak process at the LHC, while the produced photon directly reflects the production kinematics, without for example requiring any additional hadronization corrections, as in the case of jet production. This can therefore provide a valuable tool with which to constrain the gluon.

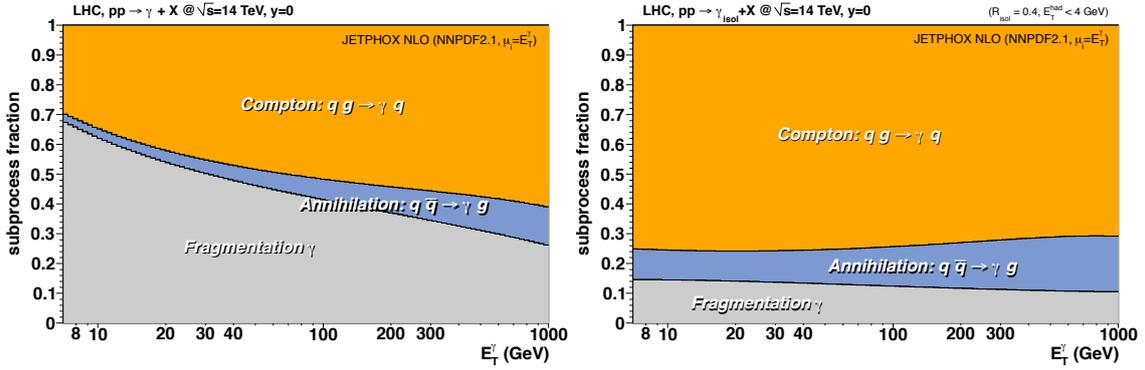


Figure 18: Relative contributions from Compton (qg), annihilation ($q\bar{q}$) and fragmentation to prompt photon production at central rapidities at the 14 TeV LHC, before (left) and after (right) the application of isolation cuts. Figures taken from [343].

However, the PDF interpretation of prompt photon production is not without its complications. In particular, the ‘direct’ process described above is not the only way in which high E_{\perp} photons can be produced in hadronic collisions. We must also include the ‘fragmentation’ contribution, whereby a standard $2 \rightarrow 2$ QCD scatter involving a final-state quark (or anti-quark) produces a photon through a collinear $q \rightarrow q\gamma$ emission. While the parton-level process for this fragmentation component carries an extra power of α_s compared to the direct production mechanism, the collinearly enhanced photon emission is effectively of order α/α_s , and thus fragmentation enters at the same order.

Technically speaking, this fragmentation emerges from the higher order corrections to the direct process. These correspond to multiple collinear splittings of a high p_{\perp} parton which end up with a photon, and that can be absorbed into universal ‘fragmentation functions’. These cannot be calculated perturbatively, but rather must be fit to data, for example in e^+e^- annihilation to hadrons. This introduces a potentially significant additional source of uncertainty, since our understanding of the non-perturbative quark-to-photon fragmentation mechanism is poor.

In fact, the situation is greatly improved by noting that physically this fragmentation process corresponds to the same multiple emission process that generates final-state jets, and indeed such fragmentation photons are typically accompanied by significant additional hadronic activity in the vicinity of the detected photon. This is to be contrasted with direct emission, where at LO the produced photon and outgoing quark are produced completely back-to-back. The direct mechanism may therefore be greatly enhanced by introducing ‘isolation’ criteria whereby the total sum of the transverse energy of the hadrons present in some cone R centred on the photon is less than a given value. These also reduces the additional ‘non-prompt’ background due to the electromagnetic decay of hadrons. The impact of such a cut is shown in Fig. 18, where it seen that the contribution from the less well known fragmentation contribution is small. This also demonstrates the dominance of the direct Compton production process for prompt photon production at the LHC.

Historically, isolated photon production represented one of the first non-DIS processes used to constrain PDFs in global analyses, and was used in such early fits as Refs. [103–105]. However, the difficulties in describing fixed-target photon production measurements from the E706 experiment [344, 345] raised questions about the reliability of this process for PDF fits within the perturbative QCD framework, and the potential sensitivities to non-perturbative effects. Combined with the increasing availability of high precision jet data from the Tevatron, which also constrain the high x gluon, this lead to the process falling

out of favour in the PDF fitting community. The last PDF set to include any such data is the MRST99 [346] fit.

However, the subsequent studies of [343, 347] (see also [348] for a study of the related $\gamma + \text{jet}$ process) have shown that by increasing the centre-of-mass energy \sqrt{s} from fixed target to collider energies and, as discussed above, imposing a suitable isolation condition on the produced photon, the process may be brought under reasonable theoretical control. Moreover, a comparison of the NLO perturbative QCD predictions to the ATLAS Run I measurements [349, 350] exhibits an adequate description of the data, albeit with fairly large $\sim 10 - 15\%$ scale variation uncertainties. These largish theoretical uncertainties have been recently reduced down to a few percent thanks to the availability of the NNLO calculation, discussed below. Therefore, isolated photon production may well provide a useful tool for PDF constraints for the next generation of global fits.

Experimental data

The available hadron collider data on isolated photon production is summarised below:

- The most recent data at 1.96 TeV from the CDF [351] and D0 [352] collaborations extends up to $E_{\perp}^{\gamma} < 0.5$ TeV and 0.3 TeV, respectively, while the photon pseudorapidity is restricted to have $|\eta^{\gamma}| \lesssim 1$. In the CDF case this corresponds to the full Run II 9.5 fb^{-1} data set, and so represents the final legacy measurement.
- The ATLAS 7 TeV measurement [353] covers up to $E_{\perp}^{\gamma} < 1$ TeV, while those at 8 [349] and 13 [350] TeV extend up to $E_{\perp}^{\gamma} < 1.5$ TeV. These datasets correspond to the full available integrated luminosities of 4.6 fb^{-1} and 20.2 fb^{-1} at 7 and 8 TeV, respectively, while the 13 TeV measurement uses a 3.2 fb^{-1} data set.
- The most precise CMS data at 7 TeV [354], corresponding to 36 pb^{-1} of integrated luminosity and extending to $E_{\perp}^{\gamma} < 0.4$ TeV.
- Data from a smaller sample at 2.76 TeV have also been taken by ATLAS [355] and CMS [356], to be used as baseline for heavy ion collision studies.

For the ATLAS and CMS measurements, the photon pseudorapidity is restricted to satisfy $|\eta^{\gamma}| \lesssim 2.4$. In addition to these collider datasets, there exist a number of older measurements taken from fixed-target experiments (see [343, 347] for a summary) but these are characterised by large experimental uncertainties, as well as by an inadequate treatment of photon isolation in some cases, and thus they should not be considered for PDF studies.

Theoretical calculations and tools

For the past 15 years, the theoretical state-of-the-art was provided by the JETPHOX [357] MC generator, which implements both the direct and the fragmentation contributions consistently at NLO. The NLO EW corrections have also been calculated in [358]. Recently, the first NNLO calculation of direct photon production has been reported [44], although the (numerically small) fragmentation component still remains at NLO. The NNLO prediction for the ATLAS 8 TeV data [349] is compared to the corresponding NLO calculation and found to lie consistently within the NLO scale uncertainty band, with the central value being $\sim 5\%$ higher. Moreover, the NNLO scale uncertainty is found to be greatly reduced, leading now to scale variations at the 3% level at most.

In order to achieve the best description of the ATLAS measurement, in [44] the leading-log EW Sudakov corrections of [339] are included, and the coupling α is evaluated at the scale m_Z , as recommended

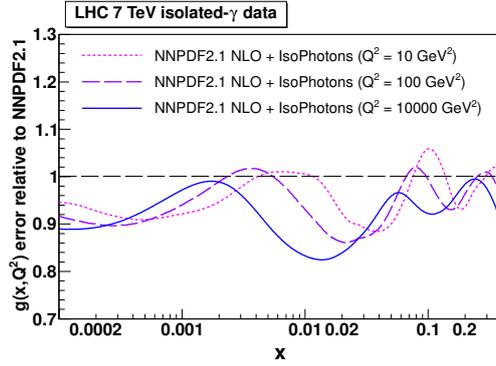


Figure 19: Relative reduction in the NNPDF2.1 NLO gluon PDF uncertainty when including a 36 pb^{-1} LHC 7 TeV isolated photon data set via Bayesian reweighting. Taken from [343].

in [339]. In particular, the EW corrections are found to reduce the cross section by as much as 10%, that is, significantly outside the QCD scale variation band, at the highest E_{\perp} , improving the shape description. The results of the PeTeR code [359], which combines the NLO calculation with N^3LL threshold resummation in addition to these EW corrections, is found to lie close to the combined NNLO+EW prediction, but with a larger uncertainty band, indicating that the data may not be too sensitive to such additional resummation effects. This is not completely unexpected, since direct photons are produced far from the kinematic threshold where the effects of resummation are expected to be more important.

Therefore, while the NNLO calculation is a very encouraging step towards including isolated photon data in high precision PDF fits, there remain still some further theoretical issues to be investigated, related to the impact of EW corrections and, as discussed in [44] the choice of photon isolation, which can also affect the comparison with the experimental data.

Impact on PDFs

Currently, no up to date studies of the impact of isolated photon data on the PDFs have been performed, in particular taking into account the new NNLO calculation and the high precision LHC data. However, in [343] (see also [347]) the impact of a range of data, including the earlier 36 pb^{-1} ATLAS and CMS measurements at 7 TeV, on the PDFs has been studied in detail through a reweighting of the NNPDF2.1 set. In Fig. 19 the impact of these LHC data on the gluon PDF is shown. A moderate reduction in the uncertainty, of up to 20%, is found in the intermediate x region. Interestingly, this overlaps with the kinematically region relevant for Higgs boson production via gluon fusion at the LHC, and indeed a $\sim 20\%$ reduction in the Higgs production cross section is found. Given these results correspond to a reasonably limited LHC data set, it would be interesting to see the impact of the latest LHC data, as well as that of the NNLO corrections.

3.7. Top quark production

In this section, we discuss the PDF information that can be obtained from top quark pair production measurements, and we also briefly review the constraints that can be potentially be derived from single top quark production.

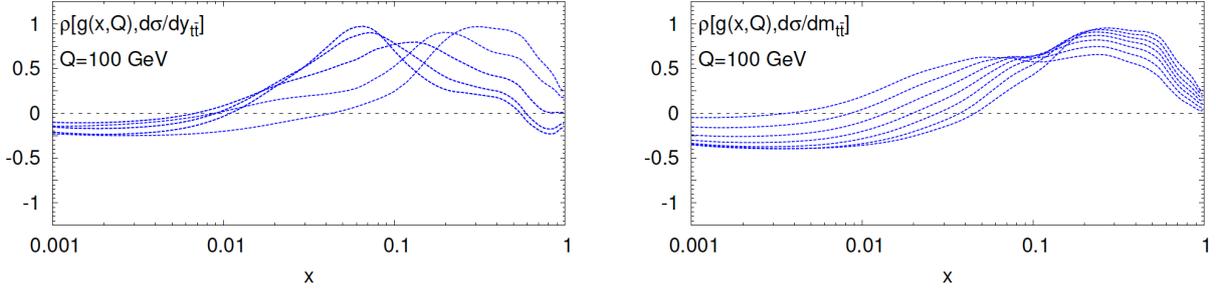


Figure 20: The correlation coefficient between the gluon PDF at $Q = 100$ GeV and the theory predictions for the absolute differential distributions in $y_{t\bar{t}}$ (left) and $m_{t\bar{t}}$ (right plot) at $\sqrt{s} = 8$ TeV, as a function of x . Each curve corresponds to a specific measurement bin. The higher the absolute value of the correlation coefficient, the bigger the sensitivity to the gluon for those specific values of x .

PDF sensitivity

The production of top quark pairs at the LHC is dominated by gluon–gluon fusion, which represents around 85% of the total cross section⁶. Therefore, provided that other sources of theoretical uncertainties such as missing higher orders and the values of the top mass m_t can be kept under control, including top quark production data into the global PDF fit has the potential to constrain the gluon in the large- x region, which is currently affected by large uncertainties.

To illustrate the kinematical sensitivity of top quark pair production to the gluon, in Fig. 20 we show the correlation coefficient $\rho[g(x, Q), d\sigma]$ between the gluon PDF at $Q = 100$ GeV and the theory predictions for the differential distributions in $y_{t\bar{t}}$ and $m_{t\bar{t}}$ at $\sqrt{s} = 8$ TeV, as a function of x . In this comparison, each curve corresponds to specific measurement bin. In the figure, the higher the absolute value of the correlation coefficient, the bigger the sensitivity to the gluon for those specific values of x . We observe that this sensitivity is high for values of x up to $x \simeq 0.6 - 0.7$, beyond the reach of other processes sensitive to the gluon, such as inclusive jet production. Moreover, it is important to emphasise that the availability of differential distributions significantly extends the kinematical coverage beyond that provided by the total inclusive cross sections.

Experimental data

The available data on top quark pair production are summarised below:

- Earlier measurements, presented at the total cross section level, have been performed first at the Tevatron [360] and then by ATLAS and CMS [361–366] at 7, 8, and 13 TeV.
- Single–inclusive differential distributions of top quark pair production have been presented by ATLAS [367] and CMS [368]. These include measurements both at the top quark level extrapolated to the full phase space $(p_T^t, y_{t\bar{t}}, m_{t\bar{t}})$, and in terms of directly observable quantities (charged lepton p_T and rapidity, b –tagged jet kinematics etc).
- Double–differential distributions for top quark pair production may also be performed, *i.e.*, as a function of $p_{t\bar{t}}$ and $m_{t\bar{t}}$, as illustrated by the recent CMS measurement [369] of normalized double–differential distributions.

⁶Note that this statement is not true at the lower values of \sqrt{s} of the Tevatron, where top quark pair production is dominated instead by quark anti–quark annihilation.

The differential measurements are often presented both as absolute distributions as well as normalized to the total fiducial cross section, in order to benefit from a number of cancellations between experimental systematic uncertainties.

Theoretical calculations and tools

The NNLO QCD calculation of the total $t\bar{t}$ production cross section has been available since 2013 [142, 370–372], including the resummation of logarithmically enhanced threshold corrections up to NNLL [373]. More recently, the full NNLO corrections to the single inclusive distributions in top quark pair production have been computed [41, 374, 375]. Differential NNLO results are available for the rapidity of the top quark and the top pair system, y_t and $y_{t\bar{t}}$, the transverse momentum of the top quark, p_T^t , and the invariant mass of the top pair $m_{t\bar{t}}$, though not for other variables such as $p_T^{t\bar{t}}$ since these vanish at leading order.

Moreover, when differential distributions probe the TeV region, electroweak corrections (including taking into account photon-initiated processes, see Sect. 7.1) also become relevant, and need to be included in the theoretical calculations. In [376] (see also [377]), the NNLO QCD calculations were combined with the state-of-the-art NLO EW corrections, in the latter case including not only the $\mathcal{O}(\alpha_s^2\alpha)$ but also the $\mathcal{O}(\alpha_s\alpha^2)$ and $\mathcal{O}(\alpha^3)$ contributions. This study showed that an accurate description of the tails of the kinematical distributions, such as the high- p_T^t and high- $m_{t\bar{t}}$ regions, must include NLO EW corrections.

An important limitation of the calculations discussed above is that they are restricted to stable top quarks. On the other hand, when experimental measurements are presented at the top quark level, they are extrapolated from the fiducial cross sections using some theoretical model, thus possibly biasing the result by an amount which is difficult to quantify. Ideally, one would like a fully differential calculation with NNLO corrections included both for production and decay, in order to directly compare with experimentally observable quantities. An important milestone in this respect was the recent calculation of top quark pair-production and decay [378] which provides predictions for observables constructed from top quark leptonic and b-tagged jet final states, based on an approximation to the exact NNLO corrections for production and the exact NNLO corrections to the decay.

Concerning the tools for the inclusion of top quark differential data into PDF fits, there exist two basic approaches. The first one is based on computing APPLgrids for the NLO calculation using either MCFM or Sherpa (see also Sect. 3.11), and then supplementing these with the NNLO/NLO bin-by-bin K -factors from [41, 374]. An improved strategy has been made feasible by the recent availability of FastNNLO tables [65] that allow the efficient calculation of NNLO top quark pair distributions for arbitrary PDF sets and input $\alpha_s(m_Z)$ values. The latter option provides a more precise evaluation of the PDF-dependent NNLO corrections, although as shown explicitly in [379] the dependence of these K -factors on the PDF set is very small.

Impact on PDFs

The availability of the NNLO calculation of the total cross sections for top quark pair production has made it possible to include top quark data from the Tevatron and the LHC consistently into a NNLO PDF fit. By applying Bayesian reweighting to NNPDF2.3, it was shown in Ref. [380] that top quark data could reduce the PDF uncertainties in the large- x gluon by up to 20% for $x \simeq 0.2$ (see also previous related work in [381]). Several other global fits, such as ABMP16 and MMHT14, also include total $t\bar{t}$ cross sections in their baseline fits. While these results provided an encouraging indication of the PDF constraining potential of $t\bar{t}$ production, the full exploitation of this potential clearly required the use of differential distributions.

The impact of the $\sqrt{s} = 8$ TeV top quark pair differential data from ATLAS and CMS at 8 TeV on the NNPDF3.0 analysis has been quantified in [379]. Here, it was shown that the constraints on the large- x gluon were at this point competitive with those provided by inclusive jet production, despite the smaller

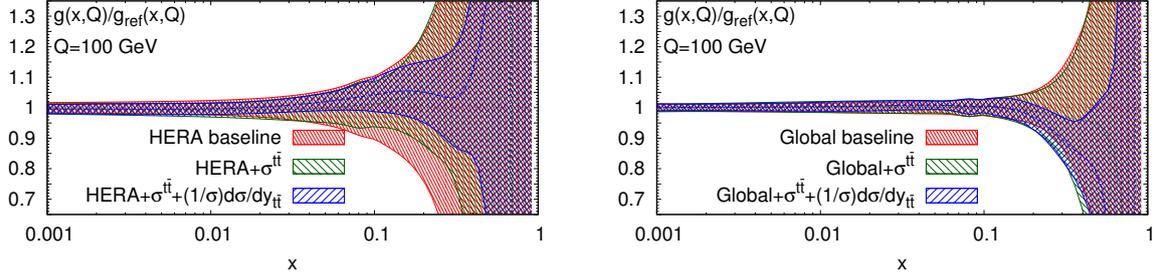


Figure 21: The impact of the LHC 8 TeV inclusive top quark pair data on the gluon PDF [379]. The results with the total cross section data only, and with the normalized $y_{\bar{t}t}$ distribution included in addition are shown, relative to HERA-only (Left) and global fit (Right) baselines is shown.

number of experimental data points. See also [382] for related work based on approximate NNLO calculations. An important result of the investigations of [379] was that the constraints from the normalized distributions were in general superior to those from their absolute counterparts, most likely because of the cancellation of systematic uncertainties that takes place in this case. In addition, top quark differential distributions at 8 TeV from the LHC have been included in the recent NNPDF3.1 global analysis, and other groups have also studied the impact of this data into their PDF fits in a preliminary form.

A challenge in the study of [379] was the observed tension between some of the ATLAS and CMS distributions, such as $m_{\bar{t}t}$, which prevented their simultaneous inclusion in the global fit. While the underlying cause of these discrepancies is still under investigation, this limitation was bypassed by identifying pairs of distributions which could be fitted with good quality at the same time and that exhibited comparable constraining power. Further investigations of this issue, including 13 TeV data and comparisons between data and theory in terms of lepton and jets observables, should be able to shed more light on the origin of such tension.

In order to illustrate the impact of the top quark pair production data on the large- x gluon, in Fig. 21 we should how the PDF uncertainties and central value of the gluon are affected by the inclusion of LHC 8 TeV data, taken from the above study. In particular, we show the results with only the total cross section data, and with the normalized $y_{\bar{t}t}$ distribution data included in addition, relative to HERA-only (left panel) and global fit (right panel) baseline fits. We can see that the impact of the total cross section data is moderate, in particular for the global fit. On the other hand, when the differential distribution data is included one clearly sees that the impact is much more significant, highlighting the increase in constraining power of the differential distributions in comparison to the total cross section data, in particular in the large- x region, where PDF uncertainties can be reduced by more than a factor of two.

Single top production

In addition to top quark pair production, single top production can also provide in principle useful PDF-sensitive information. Such a process can proceed via the scattering of a bottom quark with a light quark, see Fig. 22 (left) for a typical diagram, and will therefore provide information about the b quark PDF. In addition, due to the presence of the b quark in the initial state, it provides an important testing ground for the different heavy quark flavour schemes used in the calculation, analogously to those described in Sect. 2.5 for the case of DIS structure functions. That is, it is possible to use a $n_f = 4$ massive scheme, a $n_f = 5$ massless scheme, or a matched scheme interpolating between the two, see the discussion in Refs. [383, 384].

State of the art calculations of this process are based on NNLO QCD theory both for the total cross

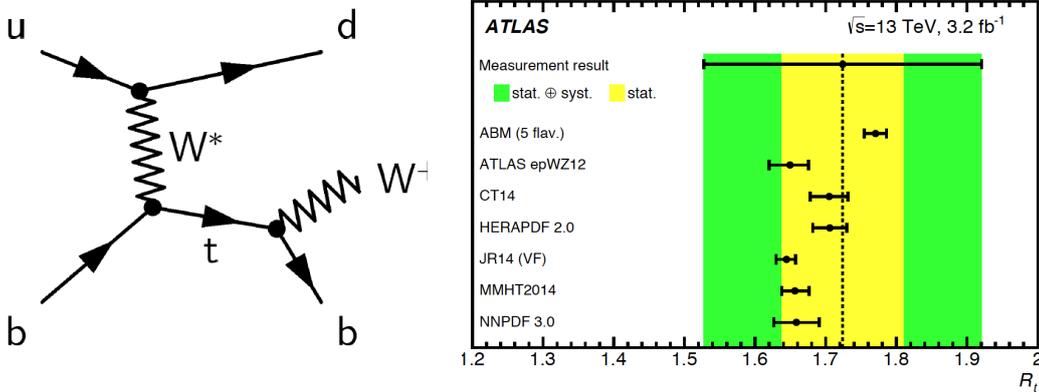


Figure 22: Left plot: one of the Feynman diagrams for single top production at leading order, illustrating its sensitivity to the b quark PDF. Right plot: comparison of the theoretical predictions for the ratio $R_t = \sigma_t/\sigma_{\bar{t}}$ from different PDF sets and the corresponding ATLAS measurements at $\sqrt{s} = 13$ TeV from [388].

sections and for differential distributions [385, 386], and LHC measurements at 8 TeV and 13 TeV of total cross sections (including ratios of top to anti-top production) as well as single inclusive distributions are already available [387, 388], although some of them only in preliminary form.

Moreover, since the production of top and anti-top quarks is generated by different initial-state partons, cross section ratios such as $R_t \equiv \sigma_t/\sigma_{\bar{t}}$ can provide important information on the quark flavour separation of the proton, specifically in the ratio u/d between valence up and down quarks at large- x . To illustrate this point, we show in Fig. 22 (right) a comparison of the theoretical predictions for the R_t ratio from different PDF sets and the corresponding ATLAS measurements at $\sqrt{s} = 13$ TeV from [388]. While experimental uncertainties are still large, due to the limited statistics, we can see that the measurement may eventually become sensitive to differences between PDF sets.

In addition, similar comparisons could also be performed for differential distributions, either at the level of top kinematic variables or observable quantities constructed from leptons and b -jets. In the case of the ATLAS 8 TeV measurements [388], these distributions are provided including the full experimental covariance matrix, and therefore all the ingredients are available in order to quantify for the first time the impact of the LHC single top production data on the PDFs.

3.8. Charm production in pp collisions

In this section we discuss the impact of open D meson production at hadron colliders on the gluon at small x .

PDF sensitivity

The production of charmed mesons at hadron colliders is dominated by the $gg \rightarrow c\bar{c}$ subprocess, and therefore it provides a sensitive probe to the gluon PDF at small- x . In particular, the forward measurements from the LHCb experiment at 5, 7, and 13 TeV [389–391] provide information on the gluon at values of x as small as $x \simeq 10^{-6}$, well below the kinematic reach of the HERA structure function data, and thus in a region where PDF uncertainties are large due to the limited amount of experimental information available.

This very small- x region is not only important for the description of soft and semi-hard QCD processes [392], it is also crucial for the calculations of signal and background processes [23–25] for ultra-high

energy neutrino astrophysics. In the latter case, the small- x gluon is relevant both for the calculation of signal event rates, via the interaction cross-section between UHE neutrinos and target nucleons (ice or water), as well as for the calculation of the rates for the dominant background process, the production of charm quarks in cosmic ray collisions in the atmosphere which then decay into so-called ‘prompt’ neutrinos and which dominate the atmospheric neutrino flux at high energies. With this motivation, various groups have studied the impact of the LHCb charm measurements on the gluon PDF using different theoretical frameworks [26, 27, 393, 394], as we discuss in more detail below.

Experimental data

The unique forward coverage of the LHCb experiment allows the very small- x region to be accessed, by means of the production of low-mass final states such as D mesons and J/Ψ mesons. LHCb has presented measurements of D meson production at three center of mass energies, $\sqrt{s} = 5, 7$ and 13 TeV [389–391]. In particular, these datasets are available double differentially in transverse momentum (p_T^D) and rapidity (y^D) for a number of final states such as D^0, D^+, D_s^+ and D^{*+} , which also contain the contribution from charge-conjugate states. The availability of the cross-sections for the same processes at different values of \sqrt{s} is particularly appealing for PDF studies, since many experimental and theoretical systematic uncertainties cancel in their ratios [245].

The original LHCb measurements of charm production were affected by an experimental issue, whose effects were particularly marked in the central rapidity region. This also affected the B meson production cross-section, where QCD calculations were far off the LHCb data for $\eta_B \simeq 2.0$ [395]. The issue has since been corrected, and the LHCb collaboration has presented revised versions of their measurements, which exhibit now an improved agreement with the theoretical predictions. Note that this issue did not affect the PDF interpretation of the LHCb D meson data, for which the dominant sensitivity comes from the region with forward rapidities, which corresponds to the small- x PDF region.

Theoretical calculations and tools

The calculation of charm production at NLO was first performed almost three decades ago [396, 397]. This fixed-order calculation can be matched to NNLL resummation [160] using a GM-VFN scheme, although these corrections are numerically small in the kinematic region accessible to the LHCb measurements. The parton-level charm production cross sections then need to be corrected for hadronization into the experimentally observed D mesons, and this can be done either analytically, using the heavy quark fragmentation functions extracted from LEP data [398], or using hadronization models as implemented in Monte Carlo event generators such as Pythia8 [399] or POWHEG [400]. The recent calculation of NNLO corrections to differential distributions for $t\bar{t}$ production in hadron collisions, see Sect. 3.7, opens up the possibility in the near future of a NNLO calculation for charm production as well.

One important challenge in order to include the LHCb D meson measurements into a PDF fit is the large scale uncertainties from the NLO calculation, which can be up to $\mathcal{O}(100\%)$. In order to tackle this limitation, it is convenient to define different cross sections ratios in order to cancel most of these theoretical errors. In particular, the following two families of observables have been advocated:

$$\begin{aligned}
 N_X^{ij} &= \frac{d^2\sigma(X \text{ TeV})}{dy_i^D d(p_T^D)_j} \bigg/ \frac{d^2\sigma(X \text{ TeV})}{dy_{\text{ref}}^D d(p_T^D)_j}, \\
 R_{13/X}^{ij} &= \frac{d^2\sigma(13 \text{ TeV})}{dy_i^D d(p_T^D)_j} \bigg/ \frac{d^2\sigma(X \text{ TeV})}{dy_i^D d(p_T^D)_j},
 \end{aligned} \tag{61}$$

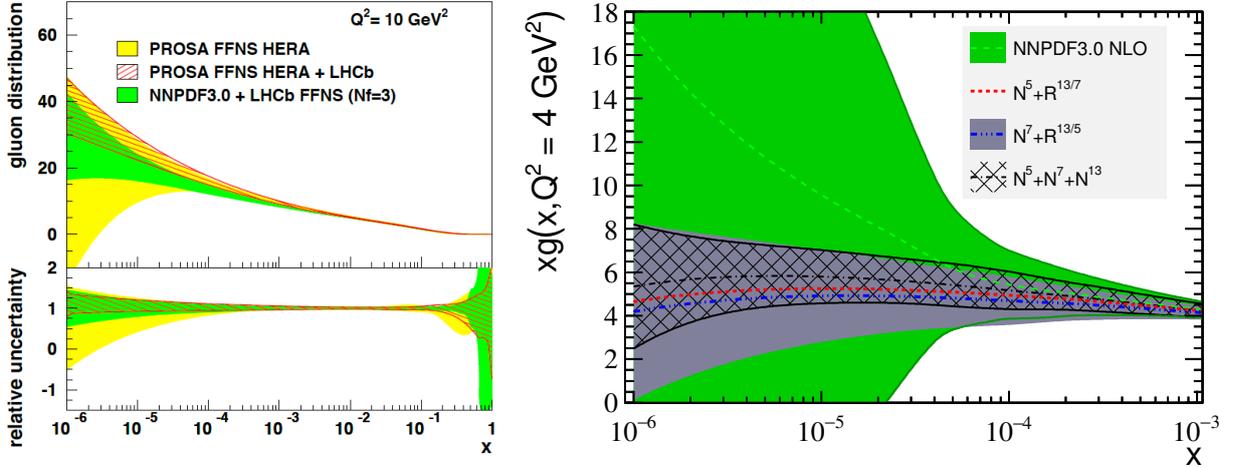


Figure 23: Left: the results of the gluon PDF from the PROSA fit of HERA structure functions and LHCb D meson 7 TeV data in the FFNS [26], compared with the corresponding fit based on the NNPDF3.0 framework [401]. Right: comparison between the small- x gluon at $Q^2 = 4 \text{ GeV}^2$ in NNPDF3.0 with the corresponding result after different combinations of the charm production data at LHCb have been included in the fit, from [27]. Specifically, we show here the central value and one-sigma PDF uncertainty bands for the $N^7 + N^{13/5}$ and the $N^5 + N^7 + N^{13}$ combinations, as well as the central value for the $N^5 + N^{13/7}$ case, see text for more details.

which benefit from the partial cancellation of the residual scale dependence from missing higher-orders, while retaining sensitivity to the gluon since different regions of x are probed in the numerator and denominator of these observables. In Eq. (61), y_{ref}^D is a specific rapidity bin chosen to normalize the absolute cross section measurements. As explained in [26, 27, 394, 401], its value should be such that the results of the PDF fit are stable with respect to small variations of this choice. Thanks to these normalizations, scale variations are greatly reduced, and in particular it can be shown [27] that the entire set of LHCb D meson measurements at 5, 7 and 13 TeV can be satisfactorily described using NLO QCD theory, with a χ^2 per point of order 1.

Impact on PDFs

In order to illustrate the impact of the LHCb open charm production data on the small- x gluon, in Fig. 23 we show the results of the gluon PDF from the PROSA fit, based on the HERA structure functions and LHCb D meson 7 TeV data in the FFNS [26], and the corresponding fit based on the NNPDF3.0 framework [401]. Comparing to the result of the HERA-only fit, we see that there is a marked reduction of the gluon PDF uncertainties down to $x \simeq 10^{-6}$. We also find that the PROSA and NNPDF results agree reasonably well with each other, both in terms of the central value and the size of the gluon PDF uncertainties at small x .

Next, in Fig. 23 we show a comparison between the small- x gluon at $Q^2 = 4 \text{ GeV}^2$ for NNPDF3.0 with the corresponding result after different combinations of the LHCb charm production data at 5, 7 and 13 TeV have been included in the fit [27]. We show the central value and one-sigma PDF uncertainty bands for the $N^7 + R^{13/5}$ and the $N^5 + N^7 + N^{13}$ combinations, as well as the central value for the $N^5 + R^{13/5}$ case, see Eq. (61) for the observable definitions. We can see that the charm data reduces the uncertainty on the small- x gluon by up to almost an order of magnitude at $x \simeq 10^{-6}$. One can also observe that the central value of the resultant small- x gluon is robust with respect to the specific choice of LHCb D meson observables used in the fit.

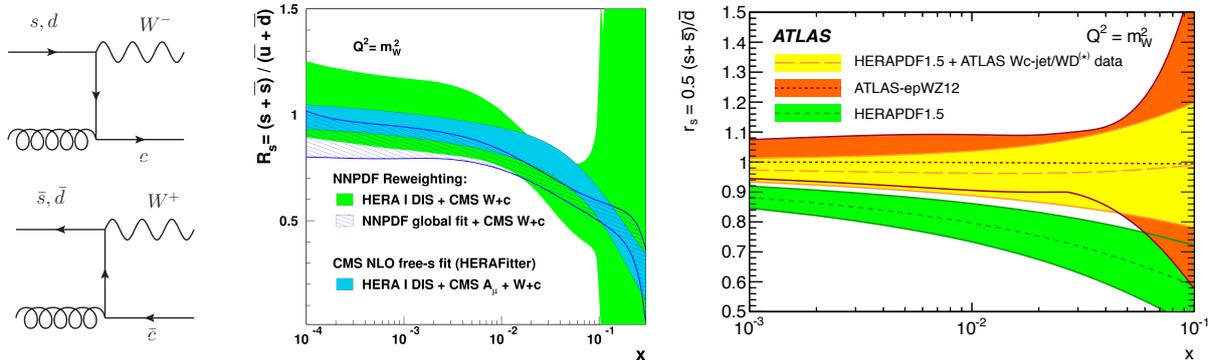


Figure 24: Left: the LO diagrams for W +charm production in pp collisions. Center: the strangeness ratio $R_S = (s + \bar{s})/(\bar{u} + \bar{d})$ at $Q = m_W$ comparing the results of the HERAFitter analysis from [404], based on the CMS $W + c$ measurements, with the corresponding results based on the NNPDF reweighting method. Right: the strangeness ratio $r_s = 0.5(s + \bar{s})/\bar{d}$, comparing the ATLAS-epWZ12 fit (which included the ATLAS 2010 W, Z inclusive rapidity distributions) with a new fit including the ATLAS $W + D$ measurements [405]. See text for more details.

The results of Fig. 23 highlight that forward D meson production allows a precision determination of the small- x gluon PDF, with implications from UHE neutrino astrophysics to future high-energy colliders (see also the discussion of Sect. 9.3). Furthermore, from the technical point of view, the inclusion of these data into the next generation of global PDF fits is facilitated by the availability of the aMCfast interface, which allows fast grids for NLO calculations matched to parton showers to be generated, as required for this process.

3.9. W production in association with charm quarks

The production of W bosons in association with D mesons is a direct probe of the strange PDF [402]. As shown in the left panel of Fig. 24, at leading order the $W + c$ process is proportional to the proton strangeness, and therefore for a long time measurements of this process have been advocated [403] as a way to provide direct information on the strange content of the proton. Moreover, by taking ratios or asymmetries of the $W^+ + \bar{c}$ and $W^- + c$ differential distributions, it is also possible to extract information on the strangeness asymmetry $s - \bar{s}$.

During Run I, the ATLAS, CMS, and LHCb collaborations have measured the production of W bosons associated with D mesons using different final states [404–406]. In the CMS case [404], the experimentally accessible $W + D$ cross sections were unfolded to the corresponding $W + c$ parton-level cross sections, facilitating the comparisons with theoretical calculations. On the other hand, the ATLAS experiment has presented their results [405] in terms of final-state quantities such as D mesons and c -tagged jets. While the two types of measurements provide in principle comparable information, the interpretation of the latter is somewhat more delicate since PDF fits are typically based on parton-level (as opposed to hadron-level) calculations. However, this is no longer a fundamental limitation as using the aMCfast [63] interface it is possible to use NLO+PS calculations directly into a PDF analysis, see Sect. 3.11.

In order to illustrate the impact of the LHC $W + D$ measurements in the strange PDF, in the central panel of Fig. 24 we show the strangeness ratio $R_S = (s + \bar{s})/(\bar{u} + \bar{d})$ at $Q = m_W$, comparing the results of the HERAFitter analysis from [404], based on the CMS $W + c$ measurements added to the HERA-I inclusive DIS data, with the corresponding results based on the NNPDF reweighting method. We can see that the CMS $W + c$ data indeed allows $s(x, Q^2)$ to be pinned down with good precision, and that the central value

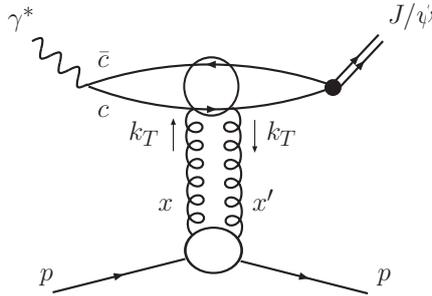


Figure 25: Photoproduction of heavy vector meson J/ψ [407].

is consistent with the global fit results, which favours a suppressed strangeness as compared to the \bar{u} and \bar{s} light quark sea. For instance, one sees that the CMS data prefers $R_S \simeq 0.6$ at $x \simeq 0.1$, though R_S tends to 1 as we move towards smaller values of x .

In Fig. 24 we also show the strangeness ratio $r_s = 0.5(s + \bar{s})/\bar{d}$, comparing the ATLAS–epWZ12 fit, which included the ATLAS 2010 W, Z inclusive rapidity distributions [275], with a new fit based on the ATLAS $W + D$ measurements [405]. In contrast to the CMS measurements, the ATLAS $W + D$ data favour a symmetric strange sea, namely $r_s \simeq 1$, which is also consistent with the constraints from the ATLAS 2010 W, Z inclusive rapidity distributions [275, 282] (as well as with the updated analysis from the 2011 dataset). While it is clear from Fig. 24 that the ATLAS and CMS W +charm measurements have opposite pulls on the strange PDF, the tension is only at the 1 or 2- σ level.

Further measurements of this process at 13 TeV should be able to shed more light on whether or not this tension between the ATLAS and CMS measurements persists. Here, a complete fit including hadron–level $W+D$ measurements (and thus avoiding unfolding to the parton level) would provide the cleanest possible interpretation of the PDF information from this process. A more detailed discussion of the strange PDF, and of the constraints from the different strange–sensitive processes, is presented in Sect. 6.3.

3.10. Central exclusive production

The Central Exclusive Production (CEP) process occurs when an object X and nothing else is produced in a hadronic collision, while the hadrons themselves remain intact after the collision. The photoproduction of heavy vector mesons, see Fig. 25, is one example of such a process which has possible implications for PDF determination. Here, one proton elastically emits a photon, while the other interacts via t –channel two gluon exchange. This may therefore access the gluon PDF at a comparatively low scale $Q^2 \sim M_V^2$ and $x \sim M_V/\sqrt{s}$, where it is so far quite poorly determined. Indeed, the production kinematics are rather similar to those of D meson production described in the previous section.

LHCb has measured the exclusive production of J/ψ and $\psi(2S)$ mesons at 7 TeV [408] and Υ production at 7 and 8 TeV [408]. Preliminary LHCb data on J/ψ and $\psi(2S)$ production at 13 TeV has also been reported in [409]. This uses the newly installed HeRSChel shower counters, which greatly extend the rapidity coverage for vetoing additional particle production, reducing significantly the non–exclusive background. In addition, the ALICE experiment has measured exclusive J/ψ production in p–Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV [410]. Due to the large Z^2 enhancement for photon emission from the Pb ion, this is dominated by the process where the lead ion emits a photon.

While clearly the photoproduction process shown in Fig. 25 proceeds through an initial–state gluon interaction, this does not correspond to a standard inclusive process where collinear PDFs are generally

introduced. However, as discussed in [407] under certain assumptions this process can be related to the standard gluon PDF and may therefore serve as a probe of it at low x and Q^2 . This analysis has subsequently been performed at NLO [411] (see also [412]) for the cases of J/ψ and Υ production. However, here it is found that the NLO correction is significantly larger than, and of the opposite sign to, the LO contribution, indicating a lack of perturbative stability and casting some doubt on the viability of this process as a PDF probe.

Nonetheless, work in the direction of at least partially solving this issue has been reported most recently in [413], where the stability is shown to be improved through a judicious choice of factorization and renormalization scale, and by imposing a cut on the NLO contribution to avoid double counting. An alternative method, not investigated so far, might be to consider ratio observables similar to those described in the previous section for the case of inclusive D meson production. It is also worth noting that the perturbative stability is naturally improved somewhat by considering the production of the higher mass Υ meson. Future studies will determine whether or not central exclusive production can then be reliably added to the toolbox of global PDF analyses.

3.11. Fast interfaces to (N)NLO calculations

Given the highly CPU time consuming nature of global PDF fits, the direct evaluation of the lengthy (N)NLO hadronic cross sections during the PDF fit is not feasible. For this reason, until around 2008, fits included hadronic data using LO hadronic cross sections supplemented by bin-by-bin K -factors, defined as

$$K_{\text{NLO}}^i \equiv \frac{\sigma_i^{\text{NLO}}}{\sigma_i^{\text{LO}}}, \quad (62)$$

using the same PDF set in the numerator and in the denominator. To ensure consistency of the procedure, these K factors can be computed iteratively until convergence was achieved. However, this approximation is known to have several deficiencies, the most important one being the reduced sensitivity to those partonic initial states that only enter the cross-section at NLO and beyond.

In order to improve upon this unsatisfactory situation, the method of fast interfaces was proposed. In these methods, the most CPU time consuming part of a (N)NLO calculation, namely the evaluation of the partonic matrix elements by integrating over a very large number of events, is precomputed *a priori* using a complete interpolation basis for the input PDFs. This way, the hadronic cross-sections can be reconstructed *a posteriori* by means of a very efficient matrix multiplication of the PDFs evaluation in a grid of (x, Q) points and the precomputed partonic matrix elements at the same grid points. These tools have become very popular within the community, and are used in the majority of modern PDF fits.

Following a common approach, two main tools have been developed, APPLgrid [61] and FastNLO [64]. The former is interfaced to the MCFM [307] and NLOjet++ [248] programs. The latter is interfaced to NLOjet++ as well as to the NNLO $t\bar{t}$ calculation of Ref. [41]. Therefore, using these tools, fast interfaces to the most commonly used codes in PDF fits can be constructed. More recently, the aMCfast interface [63] to MadGraph5_aMC@NLO [414] has also been developed. Given the automated character of this code, aMCfast allows the fast interpolation of arbitrary NLO processes, defined by the user at run time. It is also possible to produce fast grids for NLO calculations matched to parton shower MCs, which opens the way to include in the PDF fit hadron-level cross sections such as W boson in association with charmed mesons or forward D meson production at LHCb.

Currently, a significant amount of effort is being devoted to extending these fast interfaces to NNLO calculations. For instance, FastNLO has been used to produce interpolating grids for the NNLO calculation

of top-quark pair differential distributions [65], while both FastNLO and APPLgrid are being interfaced⁷ to the NNLOJET code [42, 43] making possible to generate NNLO grids for inclusive jets, dijets, and vector boson in association with jets.

To describe the basic strategy of these fast interpolation methods, we use the of APPLgrid notation for concreteness, but the general method is similar in the FastNLO and aMCfast cases. All these methods are based on representing the PDFs in (x, Q^2) by means of a suitable interpolation basis, computing a physical cross-section for a basis PDF set, and then reconstructing the same observable *a posteriori* using an arbitrary PDF set. One expands an arbitrary PDF $f(x, Q^2)$ in terms of a suitable basis of interpolating polynomials as follows

$$f(x, Q^2) = \sum_{i=0}^n \sum_{j=0}^m f_{k+i, \kappa+j} I_i^{(n)} \left(\frac{y(x)}{\delta y} - k \right) I_j^{(m)} \left(\frac{\tau(Q^2)}{\delta \tau} - \kappa \right), \quad (63)$$

where n and m are the interpolation orders in x and Q^2 respectively, $y(x) = \ln 1/x + a(1-x)$ and $\tau(Q^2) = \ln(\ln Q^2/\Lambda^2)$, and $I_i^{(n)}, I_j^{(m)}$ are interpolating functions, for instance Lagrange interpolating polynomials, though this expression holds generically for other choices. Here, a is a free parameter that controls the density of points in the large x region. In Eq. (63), k and κ are the grid nodes associated with a given values of x and Q^2 , and are defined as

$$k(x) = \text{int} \left(\frac{y(x)}{\delta y} - \frac{n-1}{2} \right), \quad \kappa(Q^2) = \text{int} \left(\frac{\tau(Q^2)}{\delta \tau} - \frac{m-1}{2} \right), \quad (64)$$

with $\text{int}(u)$ being the largest integer that is smaller than u .

After the representation of the PDFs as in Eq. (63) has been constructed, we need to evaluate cross sections using the same interpolation basis. Let us consider here for simplicity a hypothetical DIS structure function F that receives contributions from a single flavour. This NLO cross section may be computed by means of a MC program that generates a large number N of events, each one with weight ω_m and with associated values x_m and Q_m^2 . If p_m is the order of α_s for this specific event, the total cross-section can be written as

$$F = \sum_{i=1}^N \omega_i \left(\frac{\alpha_s(Q_i^2)}{2\pi} \right)^{p_i} f(x_i, Q_i^2). \quad (65)$$

The fast interpolation grid can then be constructed by, instead of computing F as in Eq. (65), introducing a weight grid $W_{i_y, i_\tau}^{(p)}$, and for each event only a fraction of the grid nodes is updated according to the expression

$$W_{k+i, \kappa+j}^{(p_i)} \rightarrow W_{k+i, \kappa+j}^{(p_i)} + \omega_i I_i^{(n)} \left(\frac{y(x_i)}{\delta y} - k \right) I_j^{(m)} \left(\frac{\tau(Q_i^2)}{\delta \tau} - \kappa \right); \quad (66)$$

Conceptually, the weight grid $W_{i_y, i_\tau}^{(p)}$ is the equivalent of computing the structure function F for a given combination of interpolating polynomials rather than for the original parton distributions.

An important factor here is that the most CPU time intensive computation, namely the calculation of the MC weights ω_m , needs only to be done once to fill the grid $W_{i_y, i_\tau}^{(p)}$, and the PDF can be decided *a posteriori* at virtually no extra computational cost. Indeed, it can be shown that the structure function can be reconstructed *a posteriori* from the weight grid using the following expression:

$$F = \sum_p \sum_{i_y} \sum_{i_\tau} W_{i_y, i_\tau}^{(p)} \left(\frac{\alpha_s(Q_{i_\tau}^2)}{2\pi} \right)^p f(x_{i_y}, Q_{i_\tau}^2). \quad (67)$$

⁷ See for instance the talk by C. Gwenlan at the Deep Inelastic Scattering Workshop 2017.

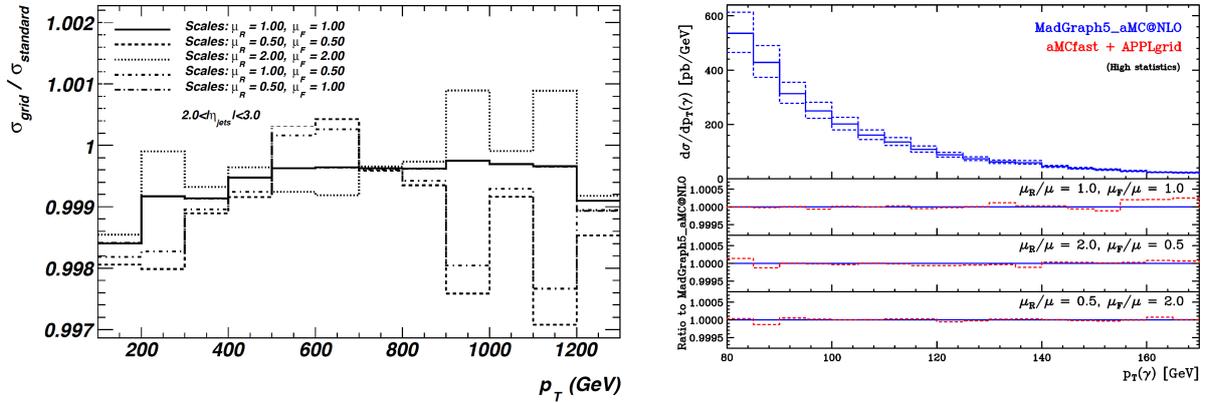


Figure 26: Left: the ratio between the NLOjet++ calculation of inclusive jet production at NLO with $\sqrt{s} = 7$ TeV in the rapidity interval $2 \leq y \leq 3$ and the corresponding *a posteriori* calculation based on APPLgrid, for different values of the factorization and renormalization scales. Right: the photon transverse momentum distribution in the $pp \rightarrow \gamma + \text{jet}$ process at 7 TeV, comparing the original MadGraph5_aMC@NLO calculation with the *a posteriori* result based on aMCfast and APPLgrid. The lower insets show the ratio between the two calculations for different choices of μ_R and μ_F , highlighting the good agreement in all cases.

In other words, the only information which is needed is the value of the PDFs and the strong coupling at the grid nodes i_y, i_τ . The method can be straightforwardly generalised to hadron-hadron collisions and to a generic composition of the initial parton state. In proton-proton collisions, the analog of Eq. (67) is given by

$$\sigma = \sum_p \sum_{l=0}^{n_{\text{sub}}} \sum_{i_{y_1}} \sum_{i_{y_2}} \sum_{i_\tau} W_{i_{y_1}, i_{y_2}, i_\tau}^{(p)(l)} \left(\frac{\alpha_s(Q_{i_\tau}^2)}{2\pi} \right)^p \mathcal{L}^{(l)}(x_{1, i_{y_1}}, x_{2, i_{y_2}}, Q_{i_\tau}^2), \quad (68)$$

where we have indicated that there are n_l contributing partonic subprocesses, each with the corresponding luminosity $\mathcal{L}^{(l)}$, which depend on the cross-section upon consideration. The more complex the process is, the larger the number of independent PDF luminosities that will be relevant for the calculation. Moreover, the extension of Eq. (68) to NNLO calculations is conceptually simple, although challenging in practice since NNLO codes are rather more complex than NLO ones.

In order to illustrate the high precision that these fast interfaces can achieve, we show two representative examples in Fig. 26. First of all, we show the ratio between the NLOjet++ calculation of inclusive jet production at NLO for 7 TeV in the rapidity interval $2 \leq y \leq 3$ and the corresponding *a posteriori* calculation based on APPLgrid, for different values of the factorization and renormalization scales. One sees that the differences between the original and the interpolated calculation are at the few permille level. Then we show the transverse momentum distribution of photons in the $pp \rightarrow \gamma + \text{jet}$ process at 7 TeV, comparing the original MadGraph5_aMC@NLO calculation with the *a posteriori* result based on aMCfast and APPLgrid. The lower insets show the ratio between the two calculations for different choices of μ_R and μ_F . Here, we also find excellent agreement between the original and interpolated calculations, now at the sub-permille level. Note that in all these methods, the interpolation accuracy can be arbitrarily increased by using denser grids in x and Q^2 .

While the fast interfaces that we have just described represent a very significant improvement in terms of CPU efficiency in comparison to the original (N)NLO calculations, one limitation of this approach can

be seen from the master formula for proton–proton collisions, Eq. (65): each time that the PDF set is varied, one needs to recompute its values in the (x, Q^2) nodes. In a PDF fit, this means that each time the input parametrization is modified during the iterative minimization, the DGLAP evolution equations need to be solved again, before the PDFs can be convoluted with the interpolated coefficient functions to obtain the hadronic cross-section. Therefore, even FastNLO or APPLgrid tables might not be fast enough in the case of PDF fits with extremely intensive minimization algorithms, such as the NNPDF approach, which involves training a neural network within a $O(300)$ parameter space.

To improve upon this shortcoming, the APFELgrid tool has recently been developed [62]. The goal of APFELgrid is to combine the interpolated partonic cross-sections provided by APPLgrid with the DGLAP evolution factors provided by APFEL, in a way that hadronic cross-sections can be reconstructed from a matrix multiplication requiring only as input the values of the PDFs at the x grid nodes at the input evolution scale Q_0 . In this way, the need to solve the DGLAP PDF evolution equations during the fit is bypassed. The combination of these two ingredients then leads to a very significant improvement in computation speed in comparison to Eq. (65) without any loss of numerical accuracy, and therefore allows much faster PDF fits. Mathematically, in the APFELgrid method an arbitrary hadronic cross-section is expressed as

$$\sigma_{pp \rightarrow X} = \sum_{k,l} \sum_{\delta,\gamma} \tilde{W}_{kl,\delta\gamma} f_k(x_\delta, Q_0^2) f_l(x_\gamma, Q_0) , \quad (69)$$

in terms of the PDFs at the parametrization scale Q_0 , where k, l run over all active parton flavours and δ, γ run over the nodes of the x interpolating grid.

To gauge the improvements in computational efficiency that can be archived by this method, in Fig. 27 we show a comparison of the timings per data point between the original APPLgrid computation of hadronic cross-sections, Eq. (67), with the same calculation based on the APFELgrid combination, Eq. (69), for a variety of LHC datasets. The improvement in computational speed is between a factor 100 and a factor 1000 depending on the specific dataset.

4. Fitting methodology

In this section we present the main features of the global PDF analysis framework. First, we discuss how the PDFs parameterization, and review the theoretical constraints that are imposed on this, such as the momentum and valence sum rules, as well as positivity. Then we discuss how to quantify the agreement of data and theory, and review various methods used in global analyses for minimization of the figure of merit, χ^2 . We also review the different strategies available to estimate and propagate PDF uncertainties, with emphasis on the three most important ones, namely the Hessian, the Monte Carlo, and the Lagrange multiplier methods.

Other related topics that are discussed in this section include how to combine individual PDF sets within a single set, the treatment of theoretical parametric uncertainties, the use of approximate methods to estimate the impact of new experiments without redoing the PDF fit, and the public delivery of PDFs by means of the LHAPDF6 interface.

4.1. PDF parametrization

We start by discussing different aspects related to the parameterization of the PDFs at the input scale Q_0 , namely the choice of functional form, the theoretical constraints from the momentum and valence sum rules and the PDF positivity, and the various quark flavour assumptions used in PDF fits.

APFELgrid/FK timings gcc-5.2.1 on i7-6500U CPU @ 2.50GHz

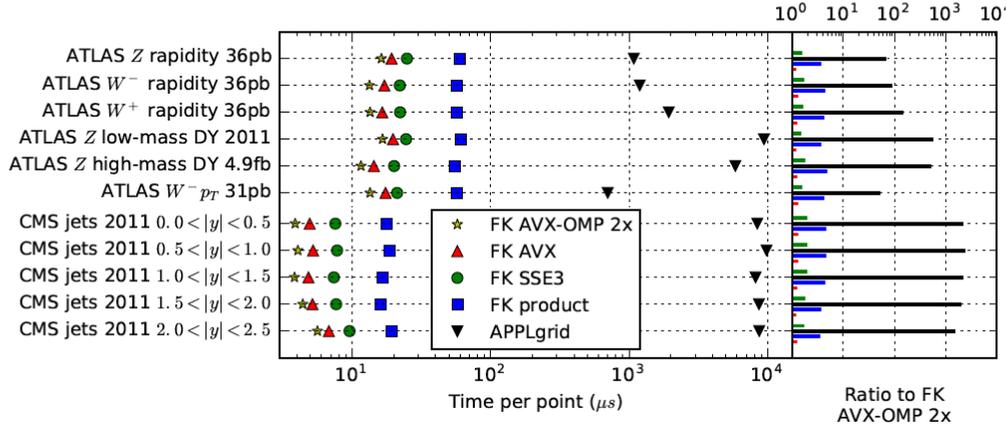


Figure 27: Comparison of the timings per data point between the original APPLgrid computation of hadronic cross-sections, Eq. (67), with the same calculation based on the APFELgrid combination, Eq. (69), for a variety of LHC datasets [62]. We find that the improvement in computational speed is between a factor 10^2 and a factor 10^3 depending on the specific process.

4.1.1. Choice of functional form

In order to extract the PDFs, a particular choice for its parameterisation in x at some input scale Q_0 must be assumed, which can then be fit to the available data. As described in Sect. 2.4, given the PDFs at some reference scale Q_0 , the DGLAP evolution equations can be used to determine the PDFs at any other scale Q . Thus the PDFs are typically parameterised at a low scale $Q_0^2 \sim 1 - 2 \text{ GeV}^2$, which can then be evolved up to the scale relevant to e.g. LHC phenomenology. These parametrizations usually adopt the generic form

$$xf(x, Q_0^2) = A_f x^{a_f} (1-x)^{b_f} I_f(x). \quad (70)$$

The $(1-x)^{b_f}$ term, with $b_f > 0$, ensures that the PDFs vanish in the elastic $x \rightarrow 1$ limit, as we would expect on basic physical grounds. Such a form is also expected from the quark counting rules [415] (see also the discussion on [416]). There, in this elastic limit all the momentum is carried by the struck parton and the remaining n_s quark become spectators. An analysis of the scaling behaviour for elastic scattering then predicts $b_f = 2n_s - 1$, that is $b_f = 3, 5$ and 7 for the valence, sea and gluon PDFs, respectively.

The x^{a_f} form dominates at low x ; in this region, the PDFs may be related to the high energy parton–proton scattering amplitudes, which can be calculated using the tools of Regge theory. This scenario predicts such a simple power–like form, with the precise value of the power a_f being related to the leading Regge trajectory that is exchanged; for non–singlet distributions (e.g. the valence quarks) this predicts $a_f \sim 0.5$ and for singlet distributions (e.g. the gluon and the sea) this predicts $a_f \sim 0$.

We emphasise that the above discussion only corresponds to quite general expectations (as opposed to direct QCD predictions), which do not for example account for the scale dependence of the PDFs. Thus while the high and low x form of Eq. (70) is usually adopted, in modern fits the values of the powers themselves are more generally left free where there is sufficient data to constrain them.

The $I_f(x)$ in Eq. (70) is the interpolating function, which determines the behaviour of the PDFs away from the $x \rightarrow 0$ and 1 limits, where it tends to a constant value. This is assumed to be a smoothly varying

function of x , for which a variety of choices have been made in the literature. The simplest ansatz, which has been very widely used, is to take a basic polynomial form in x (or \sqrt{x}), such as

$$I_f(x) = 1 + c_f \sqrt{x} + d_f x + \dots . \quad (71)$$

Forms of this type are for example taken by the CJ and HERAPDF groups as well as in the MSTW08 analysis. A similar approach, but where the polynomial enters as the exponent of a power of x or a simple exponential function, are taken by ABM and earlier CT sets, respectively.

Such a choice is appropriate for a relatively small number of parameters, say only two or three in addition to a_f and b_f . However, as the precision and amount of the data included in the fit increases, it becomes essential to allow for an increasingly flexible parameterisation. As discussed in [417], simply adding more parameters to (71) can quickly run into the technical issues that large coefficients appear, with large cancellations between the terms. This leads to an unstable χ^2 minimisation and implausibly large variations in x in certain regions. This issue may be solved by instead expanding the interpolating function in terms of a basis of suitably chosen functions with the generic form

$$I_f(x) = \sum_{i=1}^n \alpha_{f,i} P_i(y(x)) , \quad (72)$$

where $y(x)$ is some simple function of x . Two possible choices for the functions P_i are Chebyshev and Bernstein polynomials, which are used in the MMHT14 and CT14 sets, respectively. These are taken because each order of the polynomials is strongly peaked at different values of y , and hence x , significantly reducing the degree of correlation between the terms. In addition, as the order is increased these tend to probe smaller scale variations in x , so that the smoothness requirement for $I(x)$ naturally leads to smaller coefficients α at higher i . Thus, while formally equivalent to the simple polynomial expansion in Eq. (71), these are much more convenient for fitting as the number of free parameters n is increased.

An alternative approach is taken by the NNPDF group. Here, the interpolating function is modelled with a multi-layer feed forward neural network (also known as a *perceptron*), see Sect. 5.3 for more details. In practice, this allows for a greatly increased number of free parameters, with the latest default fit having 37 per PDF, that is around an order of magnitude higher than other sets. The form of Eq. (70) is still assumed, but these are pre-processing factors that speed up the minimisation procedure but which do not in principle have to be explicitly included. Nonetheless, the study of [416] has shown that the NNPDF fit does exhibit high and low x behaviour that is consistent with Eq. (70), providing further support for such an assumption in the choice of input PDF parametrization.

4.1.2. Sum rules

The flavour quantum numbers of the proton, uud with zero strangeness, are expressed in the three number or valence sum rules,

$$\int_0^1 dx [u(x, Q^2) - \bar{u}(x, Q^2)] = 2 , \quad (73)$$

$$\int_0^1 dx [d(x, Q^2) - \bar{d}(x, Q^2)] = 1 , \quad (74)$$

$$\int_0^1 dx [s(x, Q^2) - \bar{s}(x, Q^2)] = 0 . \quad (75)$$

Thus for the valence distributions we must have $a_f > 0$ for the exponents in Eq. (70) or these integrals will diverge. In others words, this constraint is consistent with the well known result that the $x f$ valence

distributions vanish as $x \rightarrow 0$. Although not shown explicitly, a similar constraint applies to the heavy quark PDFs as to the strange PDF. In the absence of any intrinsic heavy flavour, these are automatically satisfied.

The second moment of the sum of PDFs must also obey the momentum sum rule

$$\int_0^1 dx x \left(\sum_{f=1}^{n_f} [q_f(x, Q^2) + \bar{q}_f(x, Q^2)] + g(x, Q^2) + \dots \right) = 1, \quad (76)$$

which expresses the simple physical requirement, arising from energy conservation, that the total proton momentum must be equal to the sum of its constituents. In Eq. (76), the dots indicate possible additional contributions from other partons in the proton, for instance from the photon PDF $\gamma(x, Q^2)$ that should be accounted for in PDF sets with QED corrections (see Sect. 7.1). Eq. (76) also implies that for non-valence distributions, the exponent a_f may be negative, but must be greater than -1 to avoid giving a divergent contribution to the momentum sum rule.

The above sum rules provide additional constraints on the input PDFs, and are typically applied to fix certain parameters, for example the overall normalization A_f of the gluon or of specific quark flavour combinations. Provided these sum rules are satisfied at the input scale, it follows straightforwardly from the structure of the DGLAP evolution that they will be satisfied at any other scale μ . For instance, the $g \rightarrow q\bar{q}$ splitting can generate no net $q - \bar{q}$ component, and while the DGLAP evolution reshuffles the momentum carried between the different partons, it does not generate any violation of momentum conservation.

4.1.3. Quark flavour assumptions

Assuming that there are n_f active quark flavours at the input parametrization scale Q_0 , there will in general be $2n_f + 1$ PDFs to be parametrized and fitted from data. Assuming that the heavy quark PDFs are generated perturbatively, in addition to the gluon, in many cases the remaining 6 light quarks PDFs parametrized are not those in the flavour basis, namely

$$u, \quad \bar{u}, \quad d, \quad \bar{d}, \quad s, \quad \bar{s}, \quad (77)$$

but rather other convenient linear combinations, e.g. the valence $u_V = u - \bar{u}$ and $d_V = d - \bar{d}$ distributions, are often adopted. To give one example, in addition to the gluon the MMHT14 analysis takes as a fitting basis the following quark combinations:

$$u_V, \quad d_V, \quad \bar{d} - \bar{u}, \quad s + \bar{s}, \quad s - \bar{s}, \quad s + \bar{s} + 2(u + \bar{u} + d + \bar{d}). \quad (78)$$

As another example, the NNPDF3.0 fit parametrizes the PDFs at the input evolution scale in the so-called *evolution basis*, defined as the eigenvectors of the DGLAP evolution equations (see also Sect. 2.4),

$$\begin{aligned} \Sigma &= u + \bar{u} + d + \bar{d} + s + \bar{s} \\ T_3 &= u + \bar{u} - d - \bar{d}, \\ T_8 &= u + \bar{u} + d + \bar{d} - 2s - 2\bar{s}, \\ V &= u - \bar{u} + d - \bar{d} + s - \bar{s} \\ V_3 &= u - \bar{u} - d + \bar{d}, \\ V_8 &= u - \bar{u} + d - \bar{d} - 2s + 2\bar{s}, \end{aligned} \quad (79)$$

in addition to the gluon PDF. However, as any particular basis can be trivially related to another by a linear transformation, the resulting physics should not depend on these choices. On the other hand, different

flavour assumptions do often lead to different results in regions with scarce experimental constraints, such as the large- x region.

Historically, the strange quark has been rather less well determined than the u and d quark PDFs, and indeed in many earlier fits $s(x, Q^2)$ was not fitted from data but rather was fixed according to

$$s = \bar{s} \propto (\bar{u} + \bar{d}) . \quad (80)$$

Such a choice is still taken in the CJ15 and HERAPDF fits, due to the more restricted data set. With the increase in available data, the total strangeness $s + \bar{s}$ is now freely parameterised in all global fits. While the sum rule Eq. (75) requires there to be no overall strangeness in the proton, at a given x value there is no requirement for the $s_V = s - \bar{s}$ distribution to vanish, and indeed non-perturbative approaches such as the ‘meson cloud model’ [418] predict a non-zero strange asymmetry. However, the strange difference $s - \bar{s}$ is generally quite poorly determined and still broadly consistent with zero within current uncertainties. From the latest global fits, only MMHT14 and the NNPDF3 family fit the strange difference s_V , while for all other sets it is still assumed that $s = \bar{s}$. Note also that at NNLO, even if s_V is set to zero at the initial evolution scale, a non-zero strangeness (as well as charm and bottom) asymmetry will be generated dynamically by the DGLAP evolution equations [419].

The above discussion assumes that the charm PDF $c(x, Q^2)$ is generated by perturbative $g \rightarrow c\bar{c}$ splittings, in which case it is completely determined by the DGLAP evolution equations in terms of the light quark and gluon PDFs above the charm mass m_c , while it vanishes for $Q < m_c$. If the charm PDF is instead fitted, the input flavour assumptions need thus to be modified. In the case of NNPDF3.1 [56], the evolution basis of Eq. (79) is supplemented with $c^+ = c + \bar{c}$, which is also freely parametrized with a 37-parameter artificial neural network, while it is assumed that $c^- = c - \bar{c} = 0$. This option is also adopted in other recent studies where the charm PDF is fitted, such as in the CT14 IC analysis [55] (see also Sect. 6.4). Note that in general different flavour assumptions concerning the parametrized charm PDF are conceivable, for instance Eq. (79) could be also generalised by adding

$$T_{15} = u + \bar{u} + d + \bar{d} + s + \bar{s} - 3(c + 3\bar{c}) , \quad (81)$$

though this option has the drawback that the connection with charm-sensitive observables is less direct.

An important issue related to the PDF parametrization is that of the positivity. While, beyond LO, PDFs are scheme-dependent quantities and thus in principle can become negative, physical observables such as cross sections and structure functions should always be positive-definite. This constraint is incorporated in the (N)NLO global fits in different ways. For instance, in CT10 all PDFs are made by construction positive-definite, while MMHT14 allows the small- x gluon PDF to become negative. In the case of the NNPDF family of fits, no positivity constraints are imposed at the PDF level (except of course at the LO case), but during the fit the strict positivity of a range of physical cross sections is imposed by means of a Lagrange multiplier. Specifically, in the NNPDF3 sets the positivity of the following cross sections is imposed at $Q^2 = 5 \text{ GeV}^2$: $F_2^u, F_2^d, F_2^s, F_L, \sigma_{\text{DY}}^{u\bar{u}}, \sigma_{\text{DY}}^{d\bar{d}}$, and $\sigma_{\text{DY}}^{s\bar{s}}$. Note that in general this positivity constraint applies to all conceivable cross sections, including for instance those that involve hypothetical new particles, and is not restricted to the actual cross sections that are accessible experimentally.

4.2. Fit quality and minimization strategies

In this section we discuss how the quality of the agreement between experimental data and theoretical predictions can be quantified within a PDF fit, and the associated issue of the minimization strategy adopted to find the optimal set of PDF parameters starting from a figure of merit, χ^2 .

4.2.1. Fit quality and χ^2 definition

The quality of the agreement between experimental measurements and the corresponding theoretical predictions within a global fit is usually expressed in terms of the log-likelihood function, or χ^2 . When the correlations between the experimental systematic errors are not available, the χ^2 as a function of the PDF parameters is given by

$$\chi^2(\{a\}) = \sum_{k=1}^{N_{pt}} \frac{1}{\sigma_k^2} (D_k - T_k(\{a\}))^2, \quad (82)$$

where N_{pt} is number of data points, and σ_k are the total experimental errors, given by adding the statistical and systematic errors in quadrature. In this expression, $T_k(\{a\})$ are theoretical predictions, expressed in terms of the PDF parameters $\{a\}$, and D_k are the central values of the experimental measurement.

Modern experiments provide correlated sources of the various systematic uncertainties, in addition to the statistical and uncorrelated systematics. The simplest example is the luminosity error in collider experiments, which is fully correlated among all the bins from the same dataset. Typically, there are many other sources that are introduced in the process of any given analysis. In such cases, the χ^2 has the following form [117]

$$\chi^2(\{a\}, \{\lambda\}) = \sum_{k=1}^{N_{pt}} \frac{1}{s_k^2} \left(D_k - T_k - \sum_{\alpha=1}^{N_\lambda} \beta_{k,\alpha} \lambda_\alpha \right)^2 + \sum_{\alpha=1}^{N_\lambda} \lambda_\alpha^2, \quad (83)$$

for N_λ sources of correlated error. Here, s_k represents the total uncorrelated error, which is constructed by adding the statistical and uncorrelated systematic errors in quadrature. Each source of correlated systematic error is described by a nuisance parameter λ_α , with the error $\beta_{i,\alpha}$ correlated among all data points. Thus the induced systematic shift to the experimental measurement is $\sum_{\alpha} \beta_{k,\alpha} \lambda_\alpha$. The second sum on the right hand side of Eq. (83) includes the penalty terms to the χ^2 , assuming standard Gaussian distributions for the nuisance parameters.

In global PDF analyses we are more interested in the PDF parameters than the specific values that these nuisance parameters take. Therefore, for any given set $\{a\}$ we can first minimise the χ^2 with respect to the nuisance parameters λ_α to give the profiled log-likelihood function $\chi^2(\{a\}) \equiv \chi^2(\{a\}, \{\hat{\lambda}\})$. While naively we might worry that this would be a computationally intensive exercise, the simple quadratic dependence of the χ^2 on the λ_α allows the profiled nuisance parameter $\hat{\lambda}_\alpha$ to be solved for analytically, assuming purely Gaussian errors. Explicitly, we have

$$\hat{\lambda}_\alpha = \sum_{i=1}^{N_{pt}} \frac{(D_i - T_i)}{s_i} \sum_{\delta=1}^{N_\lambda} A_{\alpha\delta}^{-1} \frac{\beta_{i,\delta}}{s_i}, \quad (84)$$

with

$$A_{\alpha\beta} = \delta_{\alpha\beta} + \sum_{k=1}^{N_{pt}} \frac{\beta_{k,\alpha} \beta_{k,\beta}}{s_k^2}. \quad (85)$$

By substituting $\hat{\lambda}_\alpha$ into Eq. (83) we obtain the profiled χ^2 as a function of the PDF parameters,

$$\chi^2(\{a\}) = \sum_{i,j=1}^{N_{pt}} (T_i - D_i) (\text{cov}^{-1})_{ij} (T_j - D_j), \quad (86)$$

with the experimental covariance matrix and its inverse given by

$$(\text{cov})_{ij} \equiv s_i^2 \delta_{ij} + \sum_{\alpha=1}^{N_\lambda} \beta_{i,\alpha} \beta_{j,\alpha}, \quad (\text{cov}^{-1})_{ij} = \frac{\delta_{ij}}{s_i^2} - \sum_{\alpha,\beta=1}^{N_\lambda} \frac{\beta_{i,\alpha}}{s_i^2} A_{\alpha\beta}^{-1} \frac{\beta_{j,\beta}}{s_j^2}. \quad (87)$$

Thus, the profiled χ^2 is fully determined by the covariance matrix, which is itself constructed analytically in terms of the experimental statistical and systematic errors. In certain circumstances, for example in the case of some of the CMS and LHCb measurements, the experiments publish the covariance matrix directly, instead of the full breakdown of the experimental systematics. This is however less advantageous from the PDF fitting point of view.

One final subtlety concerning the construction of the covariance matrix arises due to the fact that experimental systematic errors are usually presented as relative errors $\sigma_{i,\alpha}$ with respect to the data, that is

$$(\text{cov})_{ij} = s_i^2 \delta_{ij} + \left(\sum_{\alpha=1}^{N_c} \sigma_{i,\alpha}^{(a)} \sigma_{j,\alpha}^{(a)} + \sum_{\beta=1}^{N_L} \sigma_{i,\beta}^{(m)} \sigma_{j,\beta}^{(m)} \right) D_i D_j . \quad (88)$$

Here, we have further separated these sources into N_c additive and N_L multiplicative errors; in the former case this counts those errors that are absolute in size, while in the latter those (such as the luminosity) which genuinely correspond to a relative uncertainty on the data. These have quite different statistical interpretations, and indeed it is known that the above experimental definition of the covariance matrix will result in a D'Agostini bias of the multiplicative errors [420] when used in a PDF fit.

To obtain unbiased results, one should use the ‘ t_0 ’ definition of the covariance matrix [420], given by

$$(\text{cov})_{ij} = s_i^2 \delta_{ij} + \left(\sum_{\alpha=1}^{N_c} \sigma_{i,\alpha}^{(a)} \sigma_{j,\alpha}^{(a)} D_i D_j + \sum_{\alpha=1}^{N_L} \sigma_{i,\beta}^{(n)} \sigma_{j,\beta}^{(m)} T_i^{(0)} T_j^{(0)} \right) . \quad (89)$$

That is, one rescales the multiplicative errors not by the data but by the theory prediction T_i^0 , from a previous iteration of the χ^2 minimization. An alternative prescription is the ‘ t ’ definition [3], where the multiplicative errors are rescaled by the same theoretical prediction as in the comparison to the data,

$$(\text{cov})_{ij} = s_i^2 \delta_{ij} + \left(\sum_{\alpha=1}^{N_c} \sigma_{i,\alpha}^{(a)} \sigma_{j,\alpha}^{(a)} D_i D_j + \sum_{\alpha=1}^{N_L} \sigma_{i,\beta}^{(m)} \sigma_{j,\beta}^{(m)} T_i T_j \right) , \quad (90)$$

or the *extended- t_0* and *extended- t* definitions, where both the additive and multiplicative errors are rescaled by the corresponding theory. A detailed discussion of the various possible prescriptions can be found in [3, 153].

4.2.2. Minimization strategies

The central PDF fits are determined by finding the global minimum of the log-likelihood function $\chi^2(\{a\})$. For PDF sets with a moderate number of free parameters, say less than $N_{\text{par}} \simeq 40$, numerical gradient-based algorithms are typically used. As a simple example, in Newton’s method, the trial solution for the global minimum is given by

$$a_i^{\text{trial}} = a_i^0 - \sum_{j=1}^{n_{\text{par}}} H_{ij}^{-1} d_j , \quad (91)$$

for the i th PDF parameter. Here a^0 is an arbitrary starting point in the PDF parameter space, d is the gradient and H^{-1} is the inverse of the Hessian matrix (defined in Sect. 4.3.1) at the same point. This solution is exact assuming a purely quadratic shape for the χ^2 , although in practice it can deviate significantly from this when it is far away from the global minimum. The above solution is therefore typically applied iteratively until the desired degree of convergence is achieved. However, the method will fail if the Hessian matrix H is not positive-definite, and can suffer from numerical instabilities. Various quasi-Newton methods

have been proposed to overcome these complexities in real applications, such as the Levenberg–Marquardt method used in MSTW/MMHT analyses [124], which is based on a dynamically determined combination of Newton’s method and the steepest decent method.

In gradient–based methods, the gradient and Hessian matrix must be calculated numerically by means of finite differences. Another class of widely used gradient based algorithms are the variable metric methods (VMM), where it is not necessary to calculate the Hessian matrix numerically. Instead, the matrix H is updated iteratively based only on information of the gradients. VMM is the default algorithm in the MINUIT package [421] and is used in CTEQ–TEA analyses [422].

As the number of free parameters is increased, the above methods will begin to suffer from numerical instabilities and issues with local minima. For the NNPDF analysis, where the typical number of parameters is an order of magnitude higher than in other sets, a genetic algorithm [423–425] is the appropriate choice, as demonstrated in [132]. The basic idea is to start from an ensemble of arbitrarily chosen samples of the PDF parameters. Random mutations with possible crossing–overs are then applied to generate a larger group of new samples. Those candidates predicting a lower χ^2 are then selected to form a new ensemble with the same size. This procedure is then iterated until a suitable convergence criterion is met, while care is taken to prevent over–fitting.

4.3. PDF uncertainties

It is essential to determine not only the best–fit PDFs, but also the uncertainties associated to such best–fit set in a systematic way, arising from instance from the data errors or methodological fit choices. Subsequently, it should be possible to propagate these uncertainties to the cross section predictions. The most widely–used methods to estimate PDF uncertainties fall into three main categories, known as the Hessian, Monte Carlo, and Lagrange multiplier methods, respectively. Each of these methods will be explained in turn in the following sections.

4.3.1. The Hessian method

The Hessian method to quantify PDF uncertainties was first developed in [116]. Here, we describe the basic ingredients of this method and of their subsequent refinements, mostly following the discussion of [124]. Given the χ^2 estimator, the best–fit values correspond to those for which this estimator has a global minimum, χ_{\min}^2 . In the vicinity of this minimum, the χ^2 can be approximated in terms of a quadratic expansion of the form

$$\Delta\chi^2 \equiv \chi^2 - \chi_{\min}^2 = \sum_{i,j=1}^{n_{\text{par}}} H_{ij} (a_i - a_i^0) (a_j - a_j^0), \quad (92)$$

where the n_{par} fit parameters are denoted by $\{a_1, \dots, a_{n_{\text{par}}}\}$, and where the best–fit values that minimize the χ^2 are indicated by $\{a_1^0, \dots, a_{n_{\text{par}}}^0\}$. In the quadratic expansion Eq. (92) we have introduced the Hessian matrix, defined as the matrix of second derivatives of the χ^2 with respect to the fit parameters, namely

$$H_{ij}(\{\vec{a}^0\}) \equiv \frac{1}{2} \frac{\partial^2 \chi^2}{\partial a_i \partial a_j} \Big|_{\{\vec{a}\}=\{\vec{a}^0\}}. \quad (93)$$

This Hessian matrix contains all the information necessary to quantify the PDF uncertainties. Indeed, for a generic function $\mathcal{F}[\{a_i\}]$ that depends on the PDFs and thus indirectly on the fit parameters, the associated uncertainty can be computed by means of linear error propagation

$$\sigma_{\mathcal{F}} = T \left(\sum_{i,j}^{n_{\text{par}}} \frac{\partial \mathcal{F}}{\partial a_i} (H)^{-1} \frac{\partial \mathcal{F}}{\partial a_j} \right)^{1/2}, \quad (94)$$

where $T = \sqrt{\Delta\chi^2}$ is the ‘tolerance factor’ that determines the matching between the allowed range of parameter variations around the best-fit values and the associated confidence interval of the PDF uncertainties. While textbook statistics suggest that $T = 1$ corresponds to a 68% confidence interval, in the context of a global fit there is ample evidence that somewhat larger values for the tolerance are required in the Hessian method, in particular to account for inconsistent experiments, theoretical uncertainties, and for methodological uncertainties such as the specific choice of functional form [417].

The main limitation of Eq. (94) is that in general the derivatives $\partial\mathcal{F}/\partial a_i$ are unknown. This problem can be bypassed by diagonalizing the Hessian matrix and then representing PDF uncertainties in terms of orthogonal eigenvectors. After this diagonalization procedure, Eq. (94) has the much simpler form

$$\sigma_{\mathcal{F}} = \frac{1}{2} \left(\sum_{i,j}^{n_{\text{par}}} [\mathcal{F}(S_i^+) - \mathcal{F}(S_i^-)]^2 \right)^{1/2}, \quad (95)$$

where S_i^{\pm} corresponds to the i -th eigenvector associated to positive and negative variations with respect to the best-fit value. Using the eigenvectors $\{S_i^{\pm}\}$ it is also possible to compute asymmetric PDF uncertainties using the prescription of Ref. [426].

For the tolerance factor $T = \sqrt{\Delta\chi^2}$, the original studies by the CTEQ and MRST group used values of $T = 10$ and $T = \sqrt{50}$ respectively. In more recent releases, the determination of this tolerance has been refined. In the case of the MSTW08 analyses for example (as well as the subsequent MMHT14 set), the tolerance is determined dynamically for each eigenvector by demanding that all data sets are included within a given confidence level variation. To illustrate this, in Fig. 28 we show the individual tolerance for each eigenvector of the MSTW08 global analysis, determined by the criterion that each separate experiment should be described within the 68% C.L. (inner error band) and 90% C.L. (outer error band) limit. The name of the experiment that determines the 90% C.L. tolerance for the various eigenvector directions is also shown. For the various eigenvector directions, the figure indicates the name of the experiment that determines the tolerance – in other words, the experiment which is most sensitive to variations along this specific eigenvector direction. The fact that many different experiments are responsible for determining these tolerances emphasizes the crucial importance of using a very wide dataset in global PDF analyses.

In addition to the PDF uncertainties, it is also possible to compute other statistical estimators such as the correlation coefficients between two PDF flavours or between a PDF set and a specific cross section. If we denote as above by n_{par} the number of eigenvectors (which have both positive and negative variations) and by \mathcal{F} and \mathcal{G} two physical quantities (such as a cross section or the PDFs themselves), their correlation coefficient will be given by

$$\rho(\mathcal{F}, \mathcal{G}) \equiv \cos \phi = \frac{1}{4\sigma_{\mathcal{F}}\sigma_{\mathcal{G}}} \sum_{i=1}^{n_{\text{par}}} (\mathcal{F}(S_i^+) - \mathcal{F}(S_i^-)) (\mathcal{G}(S_i^+) - \mathcal{G}(S_i^-)), \quad (96)$$

where $\sigma_{\mathcal{F}}$ and $\sigma_{\mathcal{G}}$ are computed using Eq. (95). The interpretation of the correlation coefficient $\cos \phi$ is straightforward, with values close to zero implying that the two quantities \mathcal{F} and \mathcal{G} are uncorrelated, and values close to 1 (-1) implying that they are strongly correlated (anti-correlated). This therefore provides a useful tool to investigate the PDF sensitivity of a given measurement; it is only in the kinematic region where the absolute value of the correlation coefficient is large that such a measurement is expected to have significant impact on a PDF analysis.

4.3.2. The Monte Carlo method

In this method, the propagation of the experimental data uncertainties to the PDFs is achieved by constructing a MC representation of the probability distribution associated to the original measurements. This

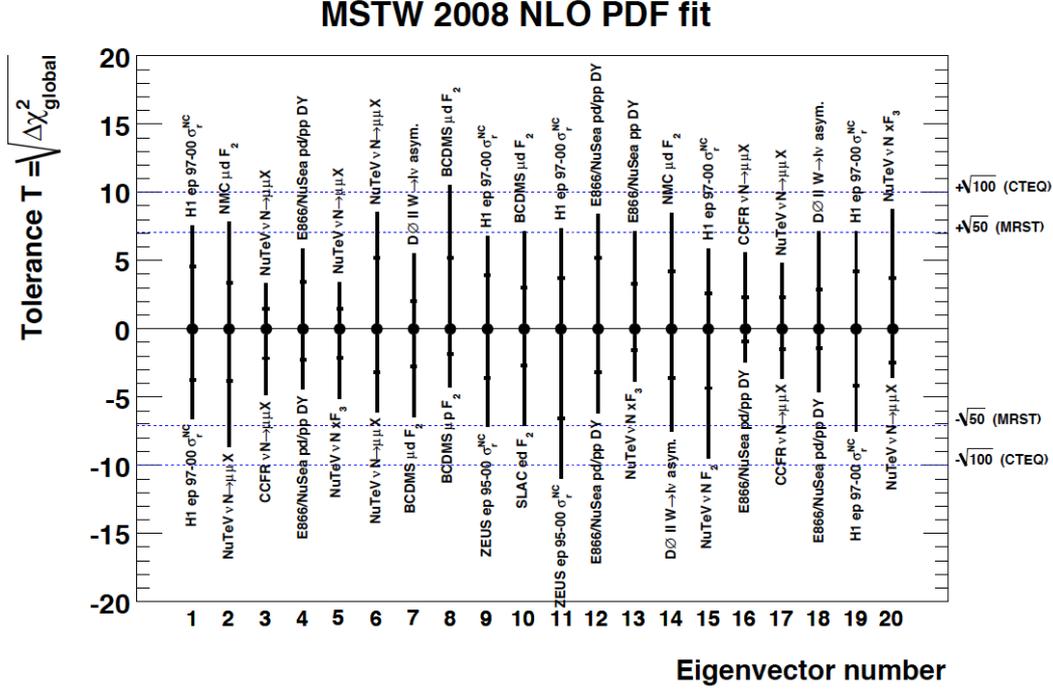


Figure 28: The individual tolerance for each eigenvector, determined by the criterion that each separate experiment should be described within its 68% C.L. (inner error bars) or 90% C.L. (outer error bars) limit. In each case the figure indicates the name of the experiment that determines the 90% C.L. tolerance for the various eigenvector directions.

requires generating a large number N_{rep} of artificial MC replicas of the original data, the so-called ‘pseudo-data’, which encode the same information on central values, variances and correlations as that provided by the experiment. In particular, given an experimental measurement of a hard-scattering cross section denoted generically by $F_I^{(\text{exp})}$ with total uncorrelated uncertainty $\sigma_I^{(\text{stat})}$, N_{sys} fully correlated systematic uncertainties $\sigma_{I,\alpha}^{(c)}$ and N_a (N_r) absolute (relative) normalization uncertainties $\sigma_{I,n}^{(\text{norm})}$, the artificial MC replicas are constructed using the following expression

$$F_I^{(\text{art})^{(k)}} = S_{I,N}^{(k)} F_I^{(\text{exp})} \left(1 + \sum_{\alpha=1}^{N_{\text{sys}}} r_{I,\alpha}^{(k)} \sigma_{I,\alpha}^{(c)} + r_I^{(k)} \sigma_I^{(\text{stat})} \right), \quad k = 1, \dots, N_{\text{rep}}, \quad (97)$$

where the normalization prefactor is given by

$$S_{I,N}^{(k)} = \prod_{n=1}^{N_a} (1 + r_{I,n}^{(k)} \sigma_{I,n}^{(\text{norm})}) \prod_{n'=1}^{N_r} \sqrt{1 + r_{I,n'}^{(k)} \sigma_{I,n'}^{(\text{norm})}}. \quad (98)$$

Here the variables $r_{I,c}^{(k)}$, $r_I^{(k)}$, $r_{p,n}^{(k)}$ are univariate Gaussian random numbers. Eq. (97) represents the fluctuations of the pseudo-data replicas around the measured central values by the amount allowed by the experimental uncertainties. Note that for each replica the random fluctuations associated to a given fully correlated systematic uncertainty will be the same for all data points, that is, $r_{I,\alpha}^{(k)} = r_{I',\alpha}^{(k)}$. The same condition holds for the normalization uncertainties. As discussed in [131], the second term in Eq. (98) (with the square

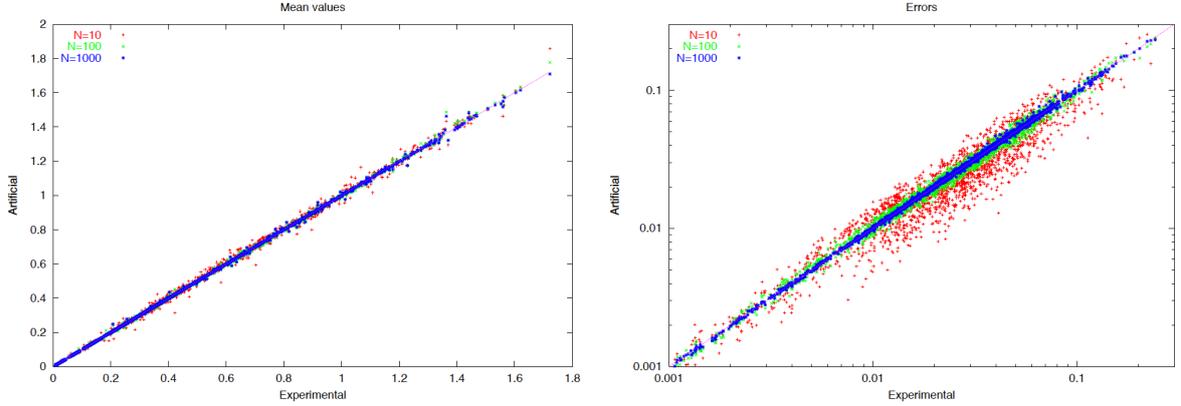


Figure 29: The scatter between the mean values (left) and variances (right plot) of all the data points included in the analysis of [428], comparing the original experimental values with the results obtained from the MC representation for different number N_{rep} of replicas.

root) correspond to the contribution of the N_r relative normalization uncertainties, which needs to be treated differently from the one associated to the N_a absolute normalization uncertainties.

An important question in the MC method is how many replicas N_{rep} need to be generated in order to achieve a faithful representation of the underlying probability density in the space of data. To this purpose, a number of statistical estimators were constructed in Ref. [427]. It was found that $N_{\text{rep}} = 10$ replicas are enough to reproduce central values, $N_{\text{rep}} = 100$ for the variances and that $N_{\text{rep}} = 1000$ is required to satisfactorily reproduce the data correlations. Subsequent analysis have shown that this statement holds for a generic input dataset. To illustrate this point, in Fig. 29 we show the scatter between the mean values and variances of all the data points included in the analysis of [428], comparing the original experimental values with the results obtained from the MC representation for different number N_{rep} of replicas. We find that indeed for central values, the averages computed using only $N_{\text{rep}} = 10$ replicas agree with the original data, but that for variances this is not the case, where $N_{\text{rep}} = 100$ replicas are required. A similar study shows that $N_{\text{rep}} = 1000$ replicas are required to faithfully reproduce the experimental data correlations.

Once the MC sampling of the experimental data has been achieved, a separate PDF fit is performed for each replica. This can be done using either traditional polynomial functional forms or other interpolators such as artificial neural networks – the MC method works in both cases. The resulting sample of N_{rep} PDF replicas realises the concept of the probability density in the space of parton distributions. The calculation of the resulting PDF uncertainties and their propagation to generic cross sections can then be performed using textbook methods. Note that in this approach the PDF uncertainty propagation is fully general, and in particular is not restricted to the Gaussian approximation.

For instance, in the MC method the expectation function of a generic cross section $\mathcal{F}[\{q\}]$ is evaluated as an average over the replica sample,

$$\langle \mathcal{F}[\{q\}] \rangle = \frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} \mathcal{F}[\{q^{(k)}\}], \quad (99)$$

and the corresponding uncertainty is then determined as the variance of the MC sample,

$$\sigma_{\mathcal{F}} = \left(\frac{1}{N_{\text{rep}} - 1} \sum_{k=1}^{N_{\text{rep}}} (\mathcal{F}[\{q^{(k)}\}] - \langle \mathcal{F}[\{q\}] \rangle)^2 \right)^{1/2}. \quad (100)$$

These formulae may also be used for the determination of central values and uncertainties of the parton distribution themselves, in which case the functional \mathcal{F} is trivially identified with the parton distribution q , that is, $\mathcal{F}[\{q\}] \equiv q$.

As in the Hessian case, see Eq. (96), the correlation coefficient between two quantities \mathcal{F} and \mathcal{G} can be straightforwardly computed, using the standard expression

$$\rho(\mathcal{F}, \mathcal{G}) = \frac{1}{\sigma_{\mathcal{F}} \sigma_{\mathcal{G}}} \left[\frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} (\mathcal{F}[\{q^{(k)}\}] \mathcal{G}[\{q^{(k)}\}]) - \left(\frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} \mathcal{F}[\{q^{(k)}\}] \right) \left(\frac{1}{N_{\text{rep}}} \sum_{l=1}^{N_{\text{rep}}} \mathcal{G}[\{q^{(l)}\}] \right) \right], \quad (101)$$

where we have used Eq. (100) to evaluate the standard deviation of \mathcal{F} and \mathcal{G} over the MC replica sample. The statistical interpretation of Eq. (101) is the same as before, with values of $\rho(\mathcal{F}, \mathcal{G})$ close to 1 (-1) indicating that there exists a strong correlation (anti-correlation) between \mathcal{F} and \mathcal{G} .

In the case of a fully consistent dataset, the MC method to estimate the PDF uncertainties is expected to coincide with the Hessian method described in Sect. 4.3.1 for a standard tolerance $\Delta\chi^2 = 1$. This equivalence was explicitly demonstrated in the HERA–LHC workshop proceedings [148]. In Fig. 30 we show the gluon PDF at $Q = 2$ GeV in this benchmark fit, based on HERA inclusive structure function data and a HERAPDF–like parametrization. The one–sigma PDF uncertainties are computed with the Hessian method (black lines) and compared to those of the Monte Carlo method (red lines), finding good agreement. In this figure, each of the green curves corresponds to an individual MC replica. In the left fit, the normalization and systematic uncertainties in the MC replicas from Eq. (97) fluctuate according to a multi–Gaussian distribution, while in the right fit they fluctuate instead according to a log–normal distribution, showing that the resulting PDFs depend only weakly on the specific assumptions about the specific probability distribution of the experimental systematic uncertainties.

Finally, we note that a Monte Carlo representation of a Hessian PDF set can be constructed following the strategy of Ref. [37], and conversely, that a Hessian representation of a Monte Carlo PDF set can be accurately constructed using the mc2h algorithm developed in Refs. [39, 429]. These techniques are discussed in more detail below in Sect. 4.4.

4.3.3. The Lagrange multiplier method

This method was originally developed in Refs. [422, 430] as a generalization of the χ^2 minimization procedure. As in the Hessian case, the first step is to find the PDF parameters $\{a_i^0\}$ that minimize the global $\chi^2(\{a_i\})$. Then one has to select a specific physical quantity that depends on the PDFs, such a DIS structure function or a hadron collider cross section, that we denote generically by $\mathcal{F}(\{a_i\})$, and which takes the value $\mathcal{F}_0 = \mathcal{F}(\{a_i^0\})$ at the global fit minimum. The goal of the Lagrange multiplier method is then to determine the PDF uncertainty associated to \mathcal{F}_0 without making any assumption about the specific behaviour of the χ^2 around the global minimum, and in particular avoiding the Gaussian assumption which is at the core of the Hessian method.

In order to achieve this, the global fit χ^2 is modified by introducing the physical quantity \mathcal{F} as a Lagrange multiplier, so that the new function that needs to be minimized is now given by the sum of two contributions,

$$\Psi(\lambda, \{a_i\}) \equiv \chi^2(\{a_i\}) + \lambda \mathcal{F}(\{a_i\}). \quad (102)$$

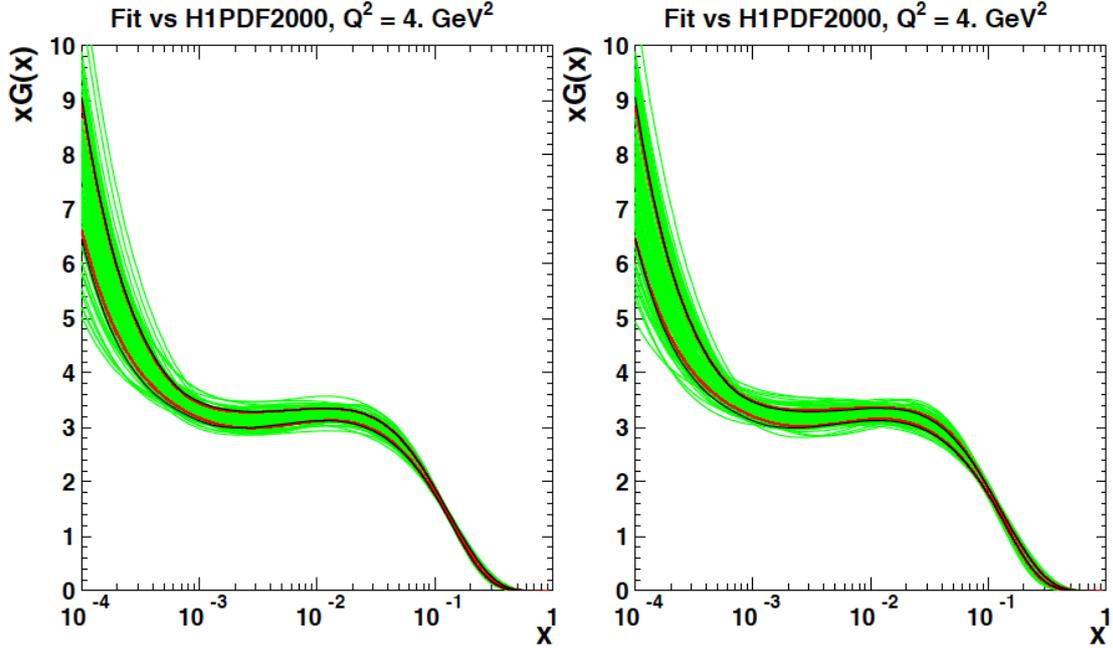


Figure 30: The gluon PDF at $Q = 2$ GeV in the HERA–LHC benchmark fit of Ref. [148], where the one–sigma PDF uncertainties computed with the Hessian method (black lines) are compared to those of the Monte Carlo method (red lines), finding good agreement. Each of the green curves corresponds to an individual MC replica. In the left fit, the normalization and systematic uncertainties in the MC replicas from Eq. (97) fluctuate according to a multi–Gaussian distribution, while in the right fit they fluctuate instead according to a log–normal distribution.

Now for each specific value of λ , denoted by λ_α , the minimization of Eq. (102) will lead to a different set of best–fit PDF parameters, which we indicate by $\{a_i^{(\min)}(\lambda_\alpha)\}$. Mathematically, these parameters are the result of a constrained PDF fit where the value of the physical observable has been fixed to $\mathcal{F}_\alpha = \mathcal{F}(\{a_i^{(\min)}(\lambda_\alpha)\})$. The resulting PDF set of this constrained fit is now indicated by S_α .

The main result of this procedure is to establish a parametric relation between the value of the physical quantity \mathcal{F} and the global fit χ^2 by means of the Lagrange multiplier λ . This means that we can determine the PDF uncertainty associated to \mathcal{F} imposing that the χ^2 satisfies $\chi^2 = \chi^2_{\min} + \Delta\chi^2$ with $\Delta\chi^2 = T^2$ representing the tolerance introduced in Sect. 4.3.1. It is clear that the main advantage of the Lagrange multiplier method in comparison to the Hessian method is that one does not need to restrict to the quadratic expansion or linear error propagation, since the PDF uncertainties in this method are determined only by the values of the χ^2 and not by its specific shape. On the other hand, an important restriction of the method is that the PDF error analysis for each specific physical quantity \mathcal{F} requires redoing a large number of new PDF fits. This is an important limitation not only because this is a CPU time intensive operation, but also because this cannot be carried out outside the PDF fitting collaborations themselves.

The Lagrange multiplier method is schematically illustrated in Fig. 31. In the left plot we show a two–dimensional projection of the PDF parameter space, indicating the contours in χ^2 for fixed values of the physical quantity \mathcal{F} . The parametric relation is provided by the value of the Lagrange multiplier λ . In the right plot we show how the PDF uncertainty associated to \mathcal{F} for a given confidence interval is determined by the condition that the global χ^2 should not grow beyond the tolerance $\Delta\chi^2$. As in the case of the Hessian

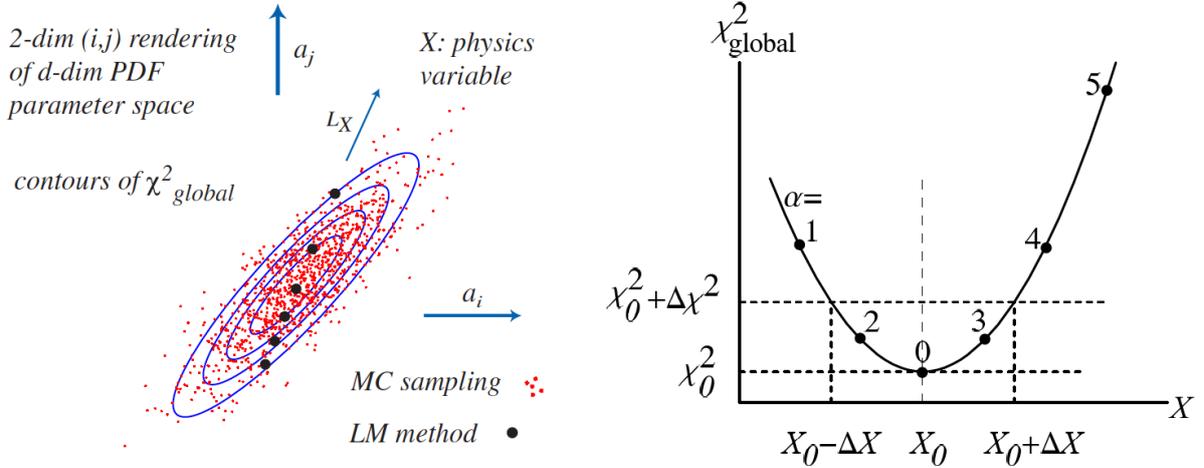


Figure 31: Schematic representation of the Lagrange Multiplier method, from Refs. [422, 430]. In the left plot we show a two-dimensional projection of the PDF parameter space, indicating the contours in χ^2 for fixed values of the physical quantity \mathcal{F} . In the right plot we show how the PDF uncertainty associated to \mathcal{F} for a given confidence interval is determined by the condition that the global χ^2 should not grow beyond the tolerance $\Delta\chi^2$.

method, the specific value of the tolerance $T = \sqrt{\Delta\chi^2}$ is an input to the method and must be determined independently.

4.4. Combined and reduced PDF sets

Individual PDF sets from different groups are widely used when comparing precision theoretical predictions with LHC measurements and in the assessment of the accuracy of PDF sets themselves. However, for several LHC applications an assessment of the *total* PDF uncertainty for certain observables, by taking into account predictions from all applicable PDF sets, is advantageous. This will for example be the case in the extraction of the couplings of the Higgs boson from the experimental cross sections, or the calculation of signal and background rates in searches for BSM physics.

For such purposes, a statistical procedure is needed with which to combine the results from different PDF sets. However, a statistical combination of PDF sets is rather more complicated than in the case of for example the world average of the strong coupling constant or heavy-quark masses, since it combines functions which have in principle an infinite number of degrees of freedom. In addition, a suitable prescription must also accommodate the fact that the individual PDF sets are not identical either in their central values or in their uncertainties, and it should account for possible correlations between PDF sets from different groups. Moreover, it would be desirable to achieve a final representation of this combined PDF set in terms of a relatively small number of Hessian eigenvectors or MC replicas.

The PDF4LHC Working Group 2010 recommendation proposed the use of a simple envelope prescription [3, 4, 137]. That is, the PDF determinations from different groups are treated as instances of a probability distribution affected by unknown sources of systematics, rather than statistically distributed instances of an underlying probability distribution. This envelope prescription can also only be applied at the level of individual observables, thereby losing the information on PDF-induced correlations. Given the better understanding of current PDF determinations, the relatively good agreement between the current releases of

the three main global PDF sets, and the high precision demands for LHC Run II studies, such a prescription is therefore certainly inadequate.

With the above considerations in mind, the updated 2015 PDF4LHC recommendation was proposed as a replacement [2]. There, a number of criteria for the individual PDF sets to be considered for the combination were adopted. First of all, the individual PDF sets should be based on a global determination with a large number of datasets from a variety of experiments, that is, DIS and hadron–hadron scattering in fixed–target and collider experiments. Second, the hard cross sections for DIS and hadron–hadron scattering processes used in the extraction should be evaluated up to two loops in QCD in a GM–VFN scheme, with a maximum number of $n_f = 5$ active quark flavours. Third, all known experimental and procedural sources of uncertainties should be properly accounted for, including the experimental uncertainties propagated from data, uncertainties due to the incompatibility of different data sets, and uncertainties due to the functional form of PDFs.

It was also decided that the combination should be carried out with a central value of $\alpha_s(m_Z) = 0.118$ at both NNLO and NLO, with the corresponding total uncertainty taken to be $\delta\alpha_s(m_Z) = 0.0015$, consistent with the most updated PDG world–average [431]. The heavy quark masses used in individual PDF sets are not currently required to be the same⁸ but should be compatible with their world–average values. The existing PDF sets satisfying all of the above requirements at present have been identified as CT14 [31], MMHT2014 [32], and NNPDF3.0 [30]. The PDF4LHC 2015 PDF sets are therefore statistical combination of these three sets, although the method can be straightforwardly generalised to include additional PDF sets satisfying the same quality conditions.

An important point here is that this statistical combination can only be carried out efficiently using the MC method, as different PDF determinations adopt different forms for the PDF parametrizations. In the first step, the CT14 and MMHT2014 PDFs, which are originally in their Hessian form, are converted into their MC representations by applying the Watt–Thorne method with symmetric formula [37]. It has been validated that a MC ensemble with $N_{\text{rep}} = 300$ replicas is sufficient to reproduce the central value and uncertainties of the original Hessian PDFs to high precision. The NNPDF3.0 PDF set is already in a MC form with 1000 replicas. Following the idea of individual PDF determinations as equally likely representations of an underlying probability distribution, a combined PDF set is then build by taking 300 MC replicas from each input PDF sets and merging them equally.

The resulting combined PDF set, an ensemble of 900 MC replicas, is referred to as the MC900 or PDF4LHC15_prior, and represents the combined probability distribution of the PDFs. However, such a set of 900 PDFs would be unmanageably large for most applications, in particular given the time and storage cost required for complicated NNLO calculations and experimental simulations. Therefore, various methods have been developed to reduce the size of the combined sets, while minimizing the information loss according to various statistical measures.

The first method applies the META–PDFs framework [38]. Here, a flexible functional form with Bernstein polynomials is chosen to parametrize the PDFs at an initial scale. Each replica in the MC900 ensemble is then represented by a group of PDF parameters through a fit to the chosen parametrizations, by minimizing a metric function. The prior probability distribution of PDFs is thus transformed into the probability distributions in the PDF parameter space. The covariance matrix of the PDF parameters is calculated as usual,

$$\text{cov}(a_l, a_m) = \frac{1}{N_{\text{rep}} - 1} \sum_{k=1}^{N_{\text{rep}}} (a_l^{(k)} - a_l^{(0)})(a_m^{(k)} - a_m^{(0)}), \quad (103)$$

⁸It would be desirable in the future for all PDF groups provide error sets with common choices of heavy–quark masses and furthermore to include the uncertainties due to the mass inputs, similarly to the case of α_s .

where $a_l^{(0)}$ and $a_l^{(k)}$ denote the fitted PDF parameters from the central set and the k -th MC replica respectively, and N_{rep} is the total number of MC replicas. The covariance matrix can be diagonalized by an orthogonal transformation. Eigenvectors are then calculated and ordered according to their impact on the PDF uncertainties with a designed error metric; the eigenvectors with smaller contributions can be dropped according to the accuracy required.

In the final step of the META–PDF procedure, a central PDF set and a group of orthogonal error PDF sets are generated under the assumption of a multi–Gaussian distribution, which can then be used in a similar way to the conventional Hessian PDF sets. For example, the 68% CL uncertainty or 1σ error on a generic cross section \mathcal{F} is given by

$$\delta^{\text{PDF}}\mathcal{F} = \sqrt{\sum_{i=1}^{N_{\text{eig}}} (\mathcal{F}_i - \mathcal{F}_0)^2}, \quad (104)$$

where \mathcal{F}_0 is the prediction on observable X given by the central set and \mathcal{F}_i is the prediction given by the i -th error set. Note there is only one error set along each eigenvector/orthogonal direction since symmetric Gaussian distributions are assumed by construction.

The second method that can be used to compress the MC900 prior set is the mc2h algorithm with Singular Value Decomposition, followed by the Principle Component Analysis [39]. The idea underlying this method is to first discretise the PDFs at $N_x N_{\text{pdf}}$ grid nodes, where N_x denotes the total number of grid points in the momentum fraction x and N_{pdf} is the number of total independent flavours. Then, a $N_x N_{\text{pdf}} \times N_x N_{\text{pdf}}$ covariance matrix on all those PDF values can be constructed from the MC replicas,

$$\text{cov}_{ll'} = \frac{1}{N_{\text{rep}} - 1} \sum_{k=1}^{N_{\text{rep}}} X_{lk} X_{kl'}^T = \frac{1}{N_{\text{rep}} - 1} X X^T, \quad (105)$$

where X_{lk} is the PDF value on l -th grid point given by k -th MC replica subtracted with the corresponding value from central PDF set, N_{rep} is the total number of MC replicas. The above covariance matrix can be rewritten in its singular value decomposition form as follows,

$$\text{cov}_{ll'} = \frac{1}{N_{\text{rep}} - 1} (U S V^T)(U S V^T)^T, \quad (106)$$

where S is a diagonal matrix constructed out from singular values of X , V is an orthogonal $N_{\text{rep}} \times N_{\text{rep}}$ matrix of coefficients, and U is a $N_x N_{\text{pdf}} \times N_{\text{rep}}$ matrix containing orthogonal eigenvectors of the covariance matrix with nonzero eigenvalues. Indeed, the matrix V gives a Hessian basis built upon linear combinations of original MC replicas, which reproduces fully the covariance matrix given by the original MC replicas. This basis can be further truncated using Principle Component Analysis with a certain error metric, resulting in a smaller Hessian PDF set, an approach similar to the case of the META–PDF method.

A third option is provided by the compressed MC (CMC) method [40]. In this case, an ensemble of pseudo–MC replica PDFs (CMC–PDFs) are generated. The CMC–PDFs have a different statistical interpretation in comparison to the native MC PDFs. However, certain statistical measures, such as the mean, covariance matrix, skewness, kurtosis and the Kolmogorov distance can be reconstructed in a similar way to the native MC PDFs. The CMC–PDFs aim to preserve some of the non–Gaussian features in the prior given by MC900 in addition to the Gaussian features, for which the Hessian form is more adequate. Note that in the current prescription some of the non–Gaussian behaviors from individual PDF sets have

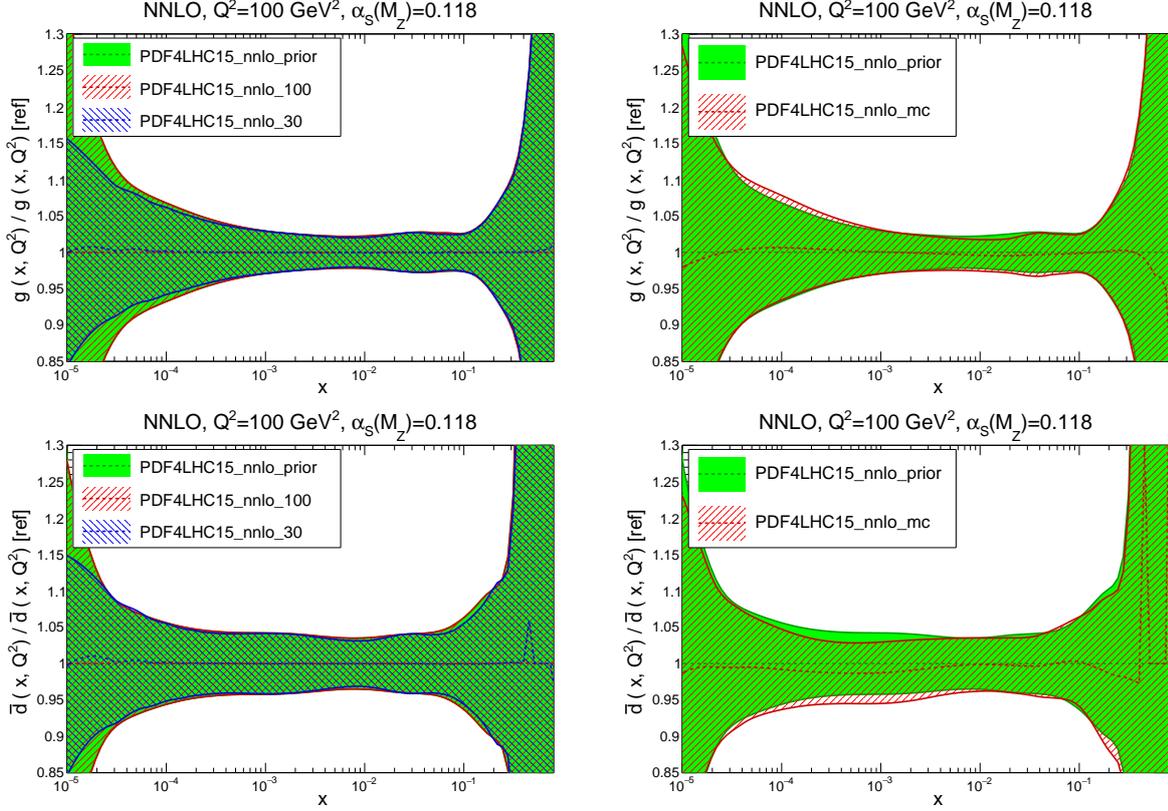


Figure 32: Comparison of the gluon and \bar{d} quark PDFs at a scale of $Q^2 = 100 \text{ GeV}^2$ between the PDF4LHC15 prior and the two reduced Hessian sets, and between the prior and the compressed MC set, normalized to the central value of the prior [2].

been smoothed out due to the symmetric formula used in converting Hessian PDFs to MC replicas. The compression starts with a figure of merit or error function ERF,

$$\text{ERF} = \sum_k \frac{1}{N_k} \sum_i \left(\frac{C_i^{(k)} - O_i^{(k)}}{O_i^{(k)}} \right), \quad (107)$$

where k runs over the number of chosen statistical estimators, N_k is a normalization factor, $O_i^{(k)}$ is the value of the k -th estimator calculated at the generic point (x_i, Q_i) from the prior and $C_i^{(k)}$ is the corresponding value of the same estimator in the compressed set. The compressed set is simply a subset of the MC900 ensemble. For any given number of total MC replicas, the compressed set is chosen by minimization of the above error function using a genetic algorithm.

These three compression and reduction techniques were used to construct the three PDF sets that are the main output of the 2015 PDF4LHC recommendation: a Hessian set with $N_{\text{eig}} = 30$ eigenvectors (PDF4LHC15_30), a Hessian set with $N_{\text{eig}} = 100$ eigenvectors (PDF4LHC15_100), and a compressed MC set with $N_{\text{rep}} = 100$ replicas (PDF4LHC15_mc). All of these are constructed from the same prior (MC900), but have a slightly different focus in each case. The symmetric PDF uncertainties of any observables can be calculated using Eq. (104) for Hessian sets and the usual master formula for MC PDFs, Eq. (100).

Fig. 32 shows the comparison of the central value and the uncertainties of the gluon and \bar{d} quark PDFs for the PDF4LHC15 prior and the three reduced sets. The agreement between the Hessian set with 100 eigenvectors and the prior is good for all PDF combinations in the complete range of x . The Hessian set with 30 eigenvectors also shows good agreement with the prior in the x range related to precision physics measurements, but gives a slightly smaller uncertainty in the extrapolation regions at small and large x as a tradeoff for the smaller number of PDF eigenvectors. The compressed MC set also agrees well with the prior in most of the region for the mean and uncertainty except for small overall fluctuations.

The PDF4LHC Working Group recommendation for the usage of these different PDF4LHC15 sets depends on the particular case under consideration:

- Use individual PDF sets, and, in particular, as many of the modern PDF sets as possible, for comparisons between data and theory for Standard Model measurements.
- Use the PDF4LHC15_mc sets for searches for BSM phenomena where non-Gaussian behaviour could be important.
- Use the PDF4LHC15_30 sets for calculation of PDF uncertainties in situations when computational speed is needed, or a more limited number of error PDFs may be desirable.
- Use the PDF4LHC15_100 sets for calculation of PDF uncertainties in precision observables.

The cases listed above are not exclusive, and one or the other will be more appropriate depending on the theoretical interpretation of a given experimental measurement.

In addition, two further combined PDF4LHC15 sets are provided with $\alpha_s(m_Z) = 0.1165$ and 0.1195 , in order to be able to estimate the uncertainty due to $\delta\alpha_s$. The corresponding uncertainty at 68% CL for a generic cross section \mathcal{F} is given by

$$\delta^{\alpha_s}\mathcal{F} = \frac{\mathcal{F}(\alpha_s = 0.1195) - \mathcal{F}(\alpha_s = 0.1165)}{2}, \quad (108)$$

where $\mathcal{F}(\alpha_s)$ is the value calculated using the PDF together with the hard matrix elements evaluated at that α_s value. The combined PDF+ α_s error is then computed by adding in quadrature the two sources of uncertainty,

$$\delta^{\text{PDF}+\alpha_s}\mathcal{F} = \sqrt{(\delta^{\text{PDF}}\mathcal{F})^2 + (r \cdot \delta^{\alpha_s}\mathcal{F})^2}, \quad (109)$$

where the rescaling factor $r = 1$ is recommended but can be varied according to user's choice of uncertainty on $\alpha_s(m_Z)$, as compared to the range $\delta\alpha_s = 0.0015$ adopted as default.

It is also noted that the PDF4LHC15 PDF sets can be further reduced to compact sets with around ten eigenvectors or less if the applications are restricted to a certain group of observables, e.g. the cross section and distributions for Higgs boson production at the LHC. This can be achieved either through the data set diagonalization method [432, 433] or the singular value decomposition method [434], the latter leading to the so-called Specialised Minimal PDF (SM-PDF) sets.

4.5. Treatment of theoretical parametric uncertainties

The theoretical calculations that enter PDF determinations depend on a number of external parameters, such as the value of the strong coupling constant $\alpha_s(m_Z)$ and that of the heavy quark masses. The choice of these input parameters therefore represents as an additional component of the total PDF uncertainty, known as the theoretical parametric uncertainty. As there can be strong correlations between the fitted PDFs and the values of these inputs, a robust evaluation of the PDF errors requires a consistent combination of these with the PDF parametric uncertainties. On the other hand, the global analysis can also provide an independent determination of those QCD parameters, which can in turn contribute to the world average values.

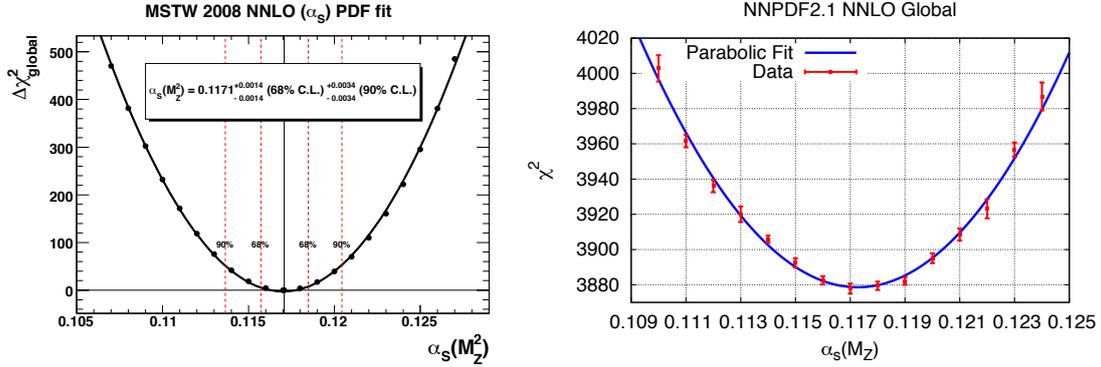


Figure 33: Left plot: the profile of global χ^2 as in a scan of $\alpha_s(m_Z)$ in MSTW2008 NNLO analysis [436]. Right plot: the profile of global χ^2 as in a scan of $\alpha_s(m_Z)$ in NNPDF2.1 NNLO analysis [437], where the error bars indicate fluctuations of the χ^2 due to finite number of MC replicas used for each α_s value.

4.5.1. The strong coupling constant α_s

The current world average value for the strong coupling constant is $\alpha_s(m_Z) = 0.1181 \pm 0.0011$. This average is extracted from six families of determinations, namely τ decay, lattice results, e^+e^- jets and event shapes, structure functions, EW precision fits, and $t\bar{t}$ cross sections at LHC. These are obtained with theoretical predictions at NNLO or higher orders and combined using the χ^2 averaging method [435]. The 2015 PDF4LHC combined PDF sets are based on a slightly different value of $\alpha_s(m_Z) = 0.1180 \pm 0.0015$ [2], that is, rounded to a value that is often used in global fits and with a somewhat more conservative uncertainty band. Individual PDF groups also extract values of $\alpha_s(m_Z)$ including uncertainties solely from their global analyses.

The choice of strong coupling constant affects a global PDF analysis in two principle ways, through the DGLAP evolution of the PDFs themselves, and through the perturbative QCD predictions for the processes that enter the fit. To study these effects, a scan over different values of $\alpha_s(m_Z)$ is typically performed. For each choice of α_s , the best-fit PDFs are found and the χ^2 profile is constructed. The best-fit value of $\alpha_s(m_Z)$ is then identified and the uncertainty on this can be determined in a similar way to the standard PDF uncertainties, using either a ‘ $\Delta\chi^2 = 1$ ’ or a tolerance criteria.

As an illustrative example, Fig. 33 shows the χ^2 profiles from MSTW and NNPDF NNLO global analyses. The extracted $\alpha_s(m_Z)$ values at NNLO are 0.1171 ± 0.0014 [436] and 0.1173 ± 0.0007 [437, 438] respectively. The CT and ABMP groups have also extracted values of the strong coupling, finding at NNLO the values 0.115 ± 0.003 [31] and 0.1147 ± 0.0008 [33], that is, with rather lower central values than the MSTW and NNPDF determinations. The error reported by the CT group is much larger than the other groups due to the weaker tolerance condition used, in a way that is consistent within uncertainties with the results from the other groups. There is therefore a large spread in the best-fit values from the different PDF groups, and so the combined 0.1156 ± 0.0021 which enters the world average has a much larger error than those reported by the individual groups. At NLO, the global analyses prefers a $\alpha_s(m_Z)$ value that is about $0.002 \sim 0.003$ higher than at NNLO, compensating for the missing higher-order corrections.

The choice of a value for the strong coupling constant $\alpha_s(m_Z)$ obviously has a significant impact on the predictions for various important processes at hadron colliders, such as the Higgs boson production via gluon fusion and top quark pair production, both of which are proportional to α_s^2 at LO. As mentioned above, here it is crucial to account for the correlations between α_s and PDFs when evaluating the full uncertainties

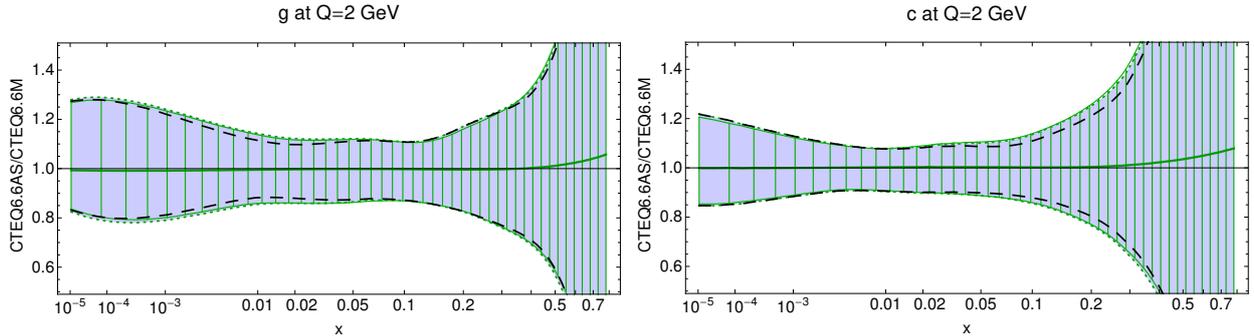


Figure 34: Comparison of the combined PDF+ α_s uncertainties for the gluon and charm quark PDFs using the full eigenvectors including $\alpha_s(m_Z)$ in the Hessian matrix (filled error band with dotted borders) and with the separate PDF and α_s uncertainties added in quadrature (hatched band with solid borders) from [439]. Here the dashed lines indicate the PDF uncertainties only.

of observables at hadron colliders. For example, it is well known that the gluon PDF is anti-correlated with α_s in the small and intermediate x regions due to the constraints from scaling violations of inclusive structure functions, which can partly compensate the change of cross sections due to change of α_s in the matrix elements.

In principle, it is possible within a global PDF analysis to treat $\alpha_s(m_Z)$ in exactly the same way as other PDF parameters, *e.g.*, in the Hessian method, by calculating the full Hessian matrix and then determining the eigenvector directions and the uncertainties along each direction. In this way the PDF+ α_s uncertainty on any observable can be evaluated using the Hessian error PDFs, in exactly the same way as for the standard case where only PDF uncertainties are included. The ABM and later ABMP group follows precisely this procedure. The downside of this approach is that it is not possible to separate the PDF and α_s uncertainties, and each error PDF will be associated with a different value of α_s .

A more convenient but completely equivalent method has been proposed in [439]. There, it has been shown that, under the quadratic approximation for the χ^2 , the full PDF+ α_s uncertainty can be calculated by simply adding the usual PDF uncertainty and the α_s uncertainty in quadrature, with the eigenvectors for the PDF uncertainties constructed with α_s fixed to its best-fit value. The α_s uncertainty is then calculated through one additional eigenvector (with two directions) constructed by fixing $\alpha_s(m_Z)$ to its upper and lower limits and then fitting the remaining PDF parameters in the usual way. The equivalence of the above two approaches is shown in Fig. 34 for the gluon and charm quark PDFs. This latter approach is now adopted by CT, MMHT, and NNPDF collaborations due to its simple form and ease of use. Note that the upper and lower limits on $\alpha_s(m_Z)$ can come either from the fit itself, as in the case of MMHT 2014, or can be chosen according to the world average, as in CT14 and NNPDF3.1. Modified ranges for the α_s uncertainty as compared to the initially assumed range $\delta\alpha_s$ can be easily obtained by means of a linear rescaling [2].

4.5.2. Heavy quark masses

Global PDF fits also rely on the input values of the heavy quark (charm, bottom and top) masses. For instance, in an analysis based on the GM-VFNS (see Sect. 2.5), the charm and bottom quark masses enter through the running of α_s , the boundary conditions for the switching of active flavours, as well as the predictions for the inclusive DIS structure functions and for open charm/bottom production in DIS or hadron-hadron collisions. On the other hand, the dependence on the top quark mass is less pronounced unless top quark production data itself are included in the analysis. In this case, since processes such as

$gg \rightarrow t\bar{t}$ are strongly dependent on both α_s , the top quark mass, and the gluon PDF, it is important to account for the non-trivial correlation between the values of m_t and α_s used in the fit.

For the charm and bottom masses, the world averages and the corresponding uncertainties for m_c and m_b in $\overline{\text{MS}}$ scheme are [435]

$$m_c(m_c) = 1.27 \pm 0.03 \text{ GeV} \text{ and } m_b(m_b) = 4.18 \pm 0.035 \text{ GeV}, \quad (110)$$

which can be translated into the following pole masses as described in [440]

$$m_c^{\text{pole}} = 1.5 \pm 0.2 \text{ GeV} \text{ and } m_b^{\text{pole}} = 4.9 \pm 0.2 \text{ GeV}, \quad (111)$$

by using the 3-loop conversion the between pole and $\overline{\text{MS}}$ values for the bottom quark mass, together with a known relation between the bottom and charm masses [441]. The larger uncertainties arise here due to the fact that the pole mass is in fact not well defined due to the diverging series, i.e. there is a renormalon ambiguity of order $\sim 0.1 - 0.2$ GeV. On the other hand, this effect largely cancels in the difference of the two masses, and therefore the above uncertainties are highly correlated.

The majority of PDF groups use the pole mass as input, as the relevant coefficient functions and matrix elements are calculated in the on-shell scheme. In particular, CT14 takes a default value of $m_{c(b)} = 1.3$ (4.75) GeV, MMHT14 takes 1.4 (4.75) GeV and NNPDF3.1 takes 1.51 (4.92) GeV. Both the CT14 and MMHT 2014 NNLO analyses prefer a lower charm quark mass of about 1.3 GeV [31, 440] if it is treated as a free parameter in the fit, which is consistent with the converted value from the $\overline{\text{MS}}$ world average value. In ABMP16 the $\overline{\text{MS}}$ masses are extracted directly from the fit, giving $m_c(m_c) = 1.252 \pm 0.018$ GeV and $m_b(m_b) = 3.84 \pm 0.12$ GeV [33].

As in the case of the strong coupling constant, the uncertainty due to the heavy quark masses can be calculated by constructing an additional eigenvector from fits with alternative mass values. The full uncertainty can then be obtained by adding it in quadrature to the PDF uncertainty obtained with the default choice of heavy quark masses. For example, the CT14, MMHT 2014 and NNPDF3.1 analyses all provide a series of best-fit PDFs with m_c or m_b fixed to alternative values around their default choices. However, there has so far not been an agreement on a common choice of the heavy quark masses and their errors in global PDF analyses, although this is foreseen for the next PDF4LHC recommendation.

The dependence of the predicted total cross sections for electroweak gauge boson and Higgs boson production at the 13 TeV LHC on the choice of charm quark pole mass used in the CT14 NNLO analysis [55] is shown in Fig. 35. It is found that varying m_c by 0.2 GeV has a negligible effect on the Higgs boson cross section and induces at most a 2% change in the weak boson cross sections. This is well within the PDF uncertainties. Similar conclusion has been reached in the MMHT 2014 analysis [440]. In addition, in both the CT14 and NNPDF3.1 analyses, it is observed that the effect of varying m_c can be partly canceled by changes in the non-perturbative component of the charm PDFs, since it is re-absorbed into the fitted charm boundary condition.

4.6. Approximate methods

In this section we discuss two approximate methods that can, under certain circumstances, be used instead of a full-fledged fit to estimate the impact on the PDFs of new experimental measurements. These two methods go under the name of the Bayesian reweighting of MC replicas [442, 443] and the profiling of Hessian sets [444]. The main advantage of these two techniques is that they can be used to quantify the impact of new experiments on a pre-existing fit based only on publicly available information, in particular the LHAPDF PDF grids, and thus such studies can be carried out outside the PDF fitting collaborations.

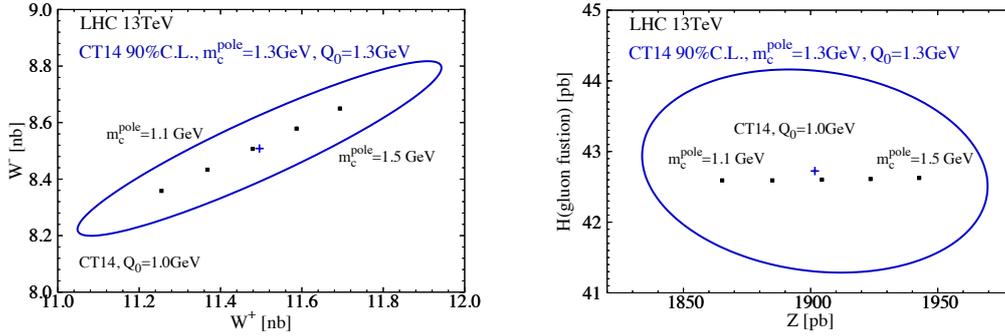


Figure 35: Dependence of the total cross sections for weak boson and Higgs boson production at the 13 TeV LHC on the choice of charm quark mass in the CT14 NNLO analysis [55]. The dots correspond to CT14 NNLO fits with charm quark mass of 1.1-1.5 GeV and an initial parametrization scale of $Q_0 = 1.0$ GeV. The ellipse indicates the PDF uncertainties at 90% C.L.

On the other hand, these methods have a number of limitations. In particular, they can fail if the impact of the new experiments is too large, and moreover are not able to account for the effect of methodological changes, such as the input PDF parametrization, or of modifications in the theoretical calculations. Therefore, they represent an important addition to the PDF fitting toolbox, but some care should be taken when using them and when interpreting their results.

4.6.1. Bayesian Monte Carlo reweighting

The Bayesian reweighting method developed in Refs. [442, 443] can be applied to a native Monte Carlo PDF set to quantify the impact at the PDF level of a new experimental measurement. The basic idea is that, starting from a sample of N_{rep} MC replicas with each of them carrying equal weight, the compatibility of each replica with the new experimental dataset can be quantified by computing a series of new weights for each replica. In this approach, the weight of the k -th replica is given by

$$\omega_k = \frac{(\chi_k^2)^{(n-1)/2} e^{-\chi_k^2/2} / N_{\text{rep}}}{\frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} (\chi_k^2)^{(n-1)/2} e^{-\chi_k^2/2}}, \quad k = 1, \dots, N_{\text{rep}}, \quad (112)$$

where χ_k^2 is the goodness-of-fit estimator between the replica k and the new experimental measurement. For instance, if for a given replica the agreement with the new experiment is very poor, its χ_k^2 will be large, and thus the weight of this specific replica will be exponentially suppressed. Note that by definition these new weights ω_k are appropriately normalized, and from the statistical point of view, they can be interpreted as the probability of the replicas f_k , given the χ_k^2 for the new experimental measurement. The validity of the Bayesian MC method has been explicitly demonstrated by comparing the reweighted results with those of a direct refit, finding good agreement in all cases.

One of the limitations of the Bayesian reweighting method is that it entails a loss of information as compared to the initial prior, because some of the original N_{rep} MC replicas will carry a very small weight, meaning that they have been effectively discarded. One suitable estimator to quantify this efficiency loss is the so-called Shannon entropy, which allows the effective number of replicas left out after the reweighting to be evaluated. Following standard information theory, the Shannon entropy of the reweighted set, or more

precisely, the effective number of replicas, is given by

$$N_{\text{eff}} \equiv \exp \left[\frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} \omega_k \ln (N_{\text{rep}} / \omega_k) \right], \quad (113)$$

where by construction, $0 \leq N_{\text{eff}} \leq N_{\text{rep}}$. The interpretation of this effective number of replicas is that a reweighted PDF set carries the same amount of information as a direct refit based on N_{eff} replicas. Clearly, the smaller N_{eff} is, the more the new dataset constrains the PDFs, but on the other hand if N_{eff} becomes small enough, the reweighting method loses validity and a full refit becomes necessary. The Bayesian reweighting approach has been applied in a variety of PDF studies, both for proton (unpolarized and polarized) PDFs, see for instance Refs. [27, 66, 261, 343, 380], and for nuclear PDFs [445].

An advantage of this Bayesian reweighting method is that it provides a way to estimate if the experimental uncertainties have been either underestimated or overestimated, assuming that theoretical uncertainties are under control. To achieve this, we can rescale the total experimental uncertainties of the data by a factor α , and then use inverse probability in order to evaluate the probability density associated to the rescaling parameter α , namely

$$\mathcal{P}(\alpha) \propto \frac{1}{\alpha} \sum_{k=1}^{N_{\text{rep}}} \omega_k(\alpha), \quad (114)$$

where the weights $\omega_k(\alpha)$ are computed using Eq. (112) but replacing χ_k^2 by χ_k^2/α^2 , and therefore represent the probability of f_k given the new data with rescaling error. If this probability density Eq. (114) peaks far above (below) one, then this suggest that the uncertainties in the data have been under (over) estimated, providing a useful handle to assess the compatibility of a new measurement with a prior PDF analysis.

It is important to mention here that for native Hessian PDF sets, as well as for Monte Carlo sets converted from Hessian sets using the procedure of [37], the weighting Eq. (112) might not be the optimal one. Indeed, other authors have advocated a different functional form for the weights, see for instance [115, 446, 447] and the discussions in [71, 448]. For example, in the nCTEQ analysis of nuclear parton distributions [71], it has been shown that using a weight of the form,

$$\omega_k = \frac{e^{-\chi_k^2/(2T)}}{\sum_k e^{-\chi_k^2/(2T)}/N_{\text{rep}}}, \quad k = 1, \dots, N_{\text{rep}}, \quad (115)$$

namely a modified Giele-Keller expression, leads to a reasonable agreement with the corresponding fit results, where T is the tolerance criterion used when defining the Hessian error PDFs. Note also that the application of the naive GK reweighting method would fail in this case, since they do not account for the tolerance $T \neq 1$ adopted in the original Hessian fit.

In order to provide an illustrative example of the Bayesian reweighting method applied to native Monte Carlo sets, in Fig. 36 we show the gluon PDF in the NNPDF3.0 closure tests [30], estimating the impact of the collider inclusive jet data and comparing the results of the Bayesian reweighting with those of a direct refit. In this study, the prior was a set of $N_{\text{rep}} = 1000$ replicas obtained with NNPDF2.3-like dataset but without any collider inclusive jet production data included. The pseudo-data were generated using the MSTW08 NLO set, though similar results were obtained with other priors. We observe that there is good agreement between the approximate Bayesian reweighting method and the exact refit results.

4.6.2. Hessian profiling

For a Hessian PDF set, the so-called profiling technique provides a method to approximately quantify the impact of a new experimental measurement. This method is based on the minimization of a χ^2 estimator

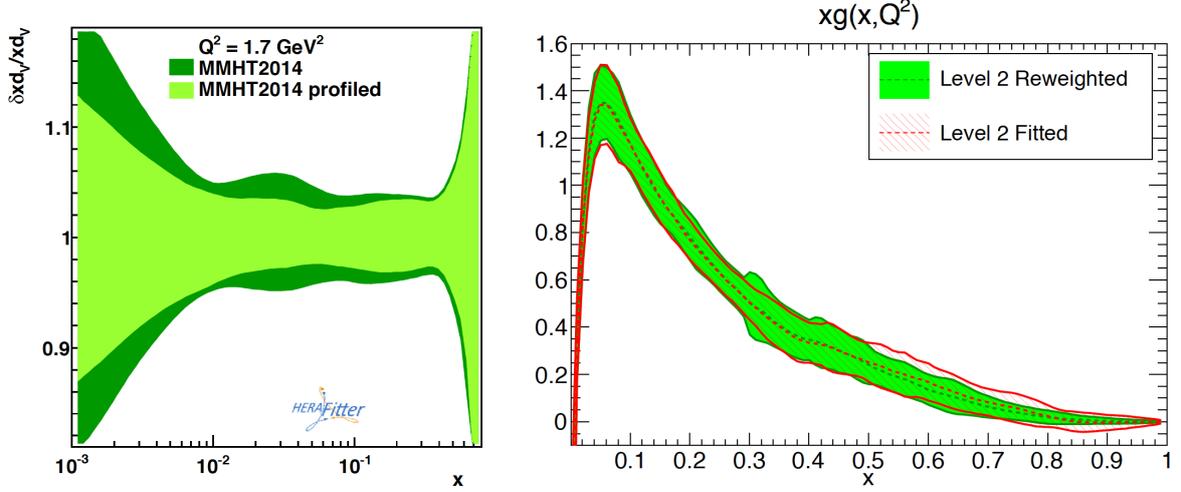


Figure 36: Left plot: the impact of the Tevatron W and Z data on the MMHT2014 NLO fit, estimated by the Hessian profiling method of Ref. [444]. Right plot: the gluon PDF in the NNPDF3.0 closure tests, quantifying the impact of the collider inclusive jet data, and comparing the results of the Bayesian reweighting with those of a direct refit.

that compares the theoretical predictions obtained with a given input Hessian PDF set with the new experimental measurements. This estimator takes into account both the experimental uncertainties and the effects from the PDF variations (as encoded by the Hessian eigenvectors) and is defined as follows:

$$\chi^2(\beta_{\text{exp}}, \beta_{\text{th}}) = \frac{1}{\Delta_i^2} \sum_{i=1}^{N_{\text{dat}}} \left(\sigma_i^{\text{exp}} + \sum_j \Gamma_{ij}^{\text{exp}} \beta_{j,\text{exp}} - \sigma_i^{\text{th}} + \sum_k \Gamma_{ik}^{\text{th}} \beta_{k,\text{th}} \right)^2 + \sum_j \beta_{j,\text{exp}}^2 + \sum_k \beta_{k,\text{th}}^2, \quad (116)$$

where $\beta_{j,\text{exp}}$ are the nuisance parameters corresponding to the set of fully correlated experimental systematic uncertainties, and $\beta_{k,\text{th}}$ are the nuisance parameters corresponding to the PDF Hessian eigenvectors. Δ_i is the total experimental uncorrelated uncertainty, and N_{dat} is the number of data points of the measurement which is being added into the PDF fit. Finally, the matrices Γ_{ij}^{exp} and Γ_{ik}^{th} encode the effects of the corresponding nuisance parameters on the experimental data and on the theory predictions, respectively.

Upon minimization of the χ^2 estimator Eq. (116), the corresponding values of the theoretical nuisance parameters, denoted by $\beta_{k,\text{th}}^{\text{min}}$, can be interpreted as leading to PDFs that have been optimized (hence the name ‘‘profiled’’) to describe this new specific measurement. Note also that in general the profiling will modify both the central value and the total PDF uncertainty. For example, the new measurement might reduce the allowed range of variation of a given eigenvector, if the original variation leads to large values of the figure of merit Eq. (116).

As in the case of the Bayesian reweighting method described in the previous section, there are a number of limitations of the the Hessian profiling method that limit the cases where it can be used to replace a complete refit. First, it assumes that the optimal PDF parametrization will not be modified by the addition of the new data. It is well known that this condition does not necessarily holds, as new experiments might require the use of more flexible input PDF parametrizations in order to achieve an optimal description, and this effect cannot be accounted for with the profiling method. Secondly, the standard version of the Hessian profiling method assumes that the PDF uncertainties are defined by the $\Delta\chi^2 = 1$ criterion, which is generally not the case for global Hessian PDF fits, see Sect. 4.3. For this reason, the impact of the data as estimated

by Hessian profiling will in general differ in comparison to the result of a full refit. However, this limitation can be eliminated by using a tolerance criterion that mimics the one used in the prior Hessian PDF set, see for example Ref. [71].

As an example of the applications of the profiling method, in Fig. 36 we show the impact of the Tevatron W and Z data on the MMHT2014 NLO set from Ref. [444], estimated by Hessian profiling. This comparison shows that these measurements lead to a reduction of the PDF uncertainties in the down valence PDF $d_V(x, Q)$. An important point to emphasise here is that this exercise was performed using completely public tools, in this case the MMHT2014 LHAPDF grids, and the experimental information of the Tevatron W and Z measurements, without any additional input from the authors of the original MMHT2014 analysis. This is the main advantage of both the Hessian profiling and Bayesian reweighting techniques. That is, they allow for the possibility of carrying out PDF studies without the need of first having to produce a complete baseline global PDF fit.

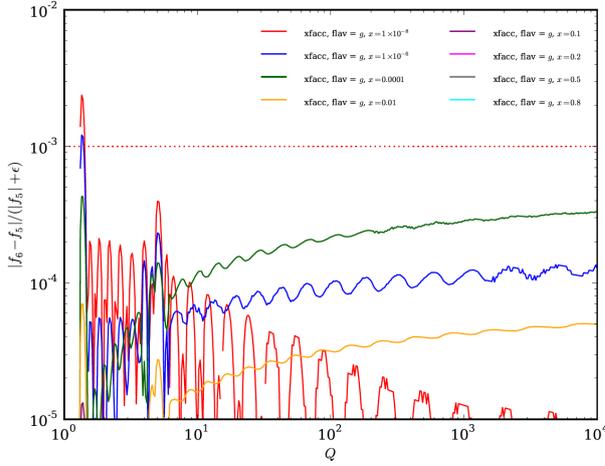
4.7. Public delivery: LHAPDF

The last step of any PDF analysis is of course to make it publicly available. In the earlier days of PDF fitting, this was achieved by means of (x, Q) interpolation tables and corresponding driver codes. Back then, the specific format of these tables and of the driver codes were specific to each PDF group. This was however far from optimal, since standardization was rather difficult, with programs requiring PDFs as input having to be adapted each time a new PDF set was released. A first step towards PDF access standardization was achieved with the release in 1993 of PDFLIB [449] as part of the CERN Program Library software. This allowed a unique interface for calling PDFs to be used without the need to add external files on a case by case basis. In addition to the PDFs, the value of $\alpha_s(Q)$ used in each specific fit could also be accessed.

The next step in this standardisation process came with the release in 2005 of LHAPDF, the *Les Houches Accord on PDFs* [147, 450], which was developed as a functional replacement for PDFLIB. In order to ensure backwards compatibility, LHAPDF included LHAPDFGlue, a PDFLIB-like interface. One of the main motivations to release LHAPDF was the realisation that dealing with the large number of error PDF sets that had then recently become available was extremely cumbersome with PDFLIB. In particular, LHAPDF was organised around the concept of *PDF set*, which was constituted by the central (average) member as well as the corresponding error PDF sets. As in the case of PDFLIB, LHAPDF was written in Fortran 77, although later a C/C++ interface was also developed.

While the Fortran incarnation of LHAPDF was very popular and widely used, at some point its further development became very challenging in particular due to the intrinsic limitations of Fortran 77 as its native language. In particular, since Fortran 77 required tspace for all available PDFs to be allocated at compilation time, the memory footprint eventually become impossible to handle and LHAPDF v5.9.1 was the last release. To overcome these limitations, a complete rewriting of LHAPDF from scratch in C++ was completed in 2014, dubbed LHAPDF6 [451]. In addition to reducing static memory requirements by orders of magnitude, this C++ incarnation of LHAPDF offered improved CPU performance and improved interpolation and extrapolation functionalities. Moreover, its cascading meta-data system ensured that software releases are completely decoupled from the availability of novel PDF sets. To ensure backwards compatibility, Fortran 77 interfaces, still very popular within many MC programs, were also provided.

In terms of interpolation accuracy, LHAPDF6 reproduces the v5 results down to residual differences of at most 0.1%. This is illustrated in Fig. 37, where we show the relative difference between LHAPDF v5 and v6 for the gluon PDF at different values of x as a function of Q , using CT10 as the input PDF. In Ref. [451] it was also shown that LHAPDF6 improves also the CPU timings as compared to v5 by a factor of between 2 and 6. This is seen in the right table in Fig. 37, which represents the timing improvements in v6 as compared to v5, t_5/t_6 , for a cross-section integration of 1M phase space points with Sherpa [452] and for



Process/PDF	t_5	t_6	t_5/t_6
Cross-section integrations, 1M phase space points			
CT10			
$pp \rightarrow jj$	23'10"	9'17"	2.5
$pp \rightarrow \ell\ell$	4'12"	2'02"	2.1
$pp \rightarrow H$ (ggF)	0'20"	0'15"	1.3
NNPDF23nlo			
$pp \rightarrow jj$	54'40"	9'28"	5.8
$pp \rightarrow \ell\ell$	8'06"	2'33"	3.2
$pp \rightarrow H$ (ggF)	0'25"	0'11"	2.3
CKKW event generation, 100k $pp \rightarrow \leq 4$ jet events			
CT10			
Weighted	43'02"	35'47"	1.2
Unweighted	5h04'39"	4h30'26"	1.1
NNPDF23nlo			
Weighted	47'47"	27'20"	1.7
Unweighted	6h44'47"	4h48'26"	1.4

Figure 37: Left plot: the relative difference between LHAPDF v5 and v6 for $g(x, Q)$ for different values of x as a function of Q , using CT10 as input PDF. Right plot: the timing improvement in v6 as compared to v5, t_5/t_6 , for a cross-section integration of 1M phase space points with Sherpa and for CKKW event generation of 100k $pp \rightarrow 4$ jet events.

CKKW event generation of 100k $pp \rightarrow 4$ jet events. The reason for this improvement is the adoption of a more efficient (x, Q^2) interpolation algorithm. Indeed, as opposed to v5 where each group had to provide their own interpolation code, in v6 the interpolation settings are universal.

Currently LHAPDF6 has established itself as the almost universal software with which to access PDFs. Its current version is 6.2 and more than 700 PDF sets can be accessed. In addition to unpolarized parton distributions, the flexibility of the LHAPDF6 framework makes it suitable to release other types of non-perturbative QCD objects, and indeed polarized PDFs [66], nuclear PDFs [70], and even fragmentation functions [453], are also available.

5. PDF analyses: state of the art

In this section we review the latest developments from the main PDF fitting groups, as well as the PDF efforts carried out within the LHC collaborations.

5.1. CT

The CTEQ-TEA global analysis was established by Wu-Ki Tung and collaborators in the early 1990s with the CTEQ1 PDFs [105]. The most recent release of general purpose PDFs from this collaboration are the CT14 PDF sets [31], which include the nominal sets as well as alternative sets with different choices of α_s and the maximum number of active flavours n_f^{\max} . In the CT approach, the PDFs are parameterised at the starting scale $Q_0 = 1.3$ GeV using the functional form Eq. (70) described in Sect. 4.1.1. In pre-CT14 analyses, the interpolating function $I_f(x)$ was chosen as an exponential of a polynomial in x or \sqrt{x} , such that positivity conditions on the PDFs at the initial scale were enforced. In the recent CT14 analysis, on the other hand, an improved parametrization choice was introduced, with for example for the u -valence quark

$$I_u(y) = d_0 p_0(y) + d_1 p_1(y) + d_2 p_2(y) + d_3 p_3(y) + d_4 p_4(y), \quad (117)$$

where $y = \sqrt{x}$ and p_n are the fourth order Bernstein polynomials, given by

$$p_0(y) = (1 - y)^4, \quad p_1(y) = 4y(1 - y)^3, \quad p_2(y) = 6y^2(1 - y)^2, \quad p_3(y) = 4y^3(1 - y), \quad p_4(y) = y^4. \quad (118)$$

Namely, the interpolating function is chosen as a fourth-order polynomial in y with an expansion in the basis of Bernstein polynomials. As discussed in Sect. 4.1.1, this greatly increases the stability of the fit within the Hessian approach. Moreover, in the CT14 case the positivity of PDFs at $Q_0 = 1.3$ GeV is not imposed *a priori* but rather emerges automatically as a consequence of the fit to data. The CT14 PDFs have a total number of 28 free parameters; using a yet more flexible parametrizations, by adding higher-order polynomials, is found to have a small effect on both the best-fit and the estimated PDF uncertainties in the region that is well constrained by the experimental data.

The CT14 global analysis includes a wide variety of experimental data. The majority comes from the inclusive DIS and semi-inclusive DIS measurements of the structure functions and the reduced cross section measurements from fixed-target experiments (BCDMS [178, 179], NMC [181], CCFR [187, 188, 191], NuTeV [192], CDHSW [454]) or HERA experiments [129, 195, 199, 455]. A Q cut of 2 GeV and W cut of 3.5 GeV are adopted in the selection of DIS data to minimize non-perturbative effects from either nuclear corrections or higher-twists corrections. Thus no further nuclear or higher twists corrections are included in theory predictions in CT14 except for those already applied in the unfolding of experimental data. For the NC DIS process, the CT14 analysis uses a treatment of heavy quark mass effects up to NNLO, through a variant of the GM-VFNS known as S-ACOT- χ [167]. For CC DIS, the theoretical calculations are only implement at NLO, which is judged to be sufficient given the relatively small number of data points and their large experimental errors. Drell-Yan production data from fixed-target experiments (E605 [456], E866 [266]) and W/Z boson production data from Tevatron [269, 457–460] including the new D0 electron charge asymmetry data [271], are also fit.

The Tevatron W, Z data provide further discriminations on quark flavours in large- x region, with the W asymmetry data probing the average slope of d/u ratio at $x \gtrsim 0.1$, see Sect. 3.4 for a more detailed discussion. Predictions from ResBos [309, 318, 461, 462] are used for the W/Z boson production data, with a p_T cut imposed on the charged leptons, and incorporating soft gluon resummation effects at small p_T of the vector boson. These resummed predictions provide a better description of the p_T spectrum of the charged leptons. The updated D0 electron charge asymmetry data [271] shows a large impact on the d/u PDF ratio at large x compared to CT10 and CT10W [125, 463]. In the CT10 fits the D0 lepton charge asymmetry data resulted in larger asymptotic value of d/u though tensions were found between different subsets of the data or the D0 data and other DIS experiments. As shown in Fig. 38, for CT14 the updated D0 electron charge asymmetry data shows better agreements with other data sets in the global analysis and drives the d/u ratio to a lower value, close to CTEQ6.6 [263] at large x . The d/u ratio in CT14 also shows good agreement with the extraction from CJ12 [464], which is based on independent large- x and low- W DIS data, including power corrections and deuteron corrections.

Similar data on W and Z boson production from LHC Run I are also included from the ATLAS [275], CMS [280, 281] and LHCb [465] experiments, which further extend the coverage to the intermediate and small- x region. In addition, single inclusive jet production from the Tevatron [233, 234] and the LHC [230, 239] are fitted, providing the dominant constraint on the gluon PDF at large x , with the latter data also extending the coverage to the intermediate x region. For inclusive jet production at hadron colliders only the NLO predictions were available at the time of the CT14 fit, and therefore this is used in the NNLO fit. This will be updated with the recent NNLO calculations [42] in future instalments of the CTEQ-TEA global analysis.

The CTEQ-TEA group uses the Hessian method with certain tolerance conditions for the nominal fits in the determination of PDF uncertainties at 90% CL, see Sect. 4.3. In addition, the Lagrange Multiplier

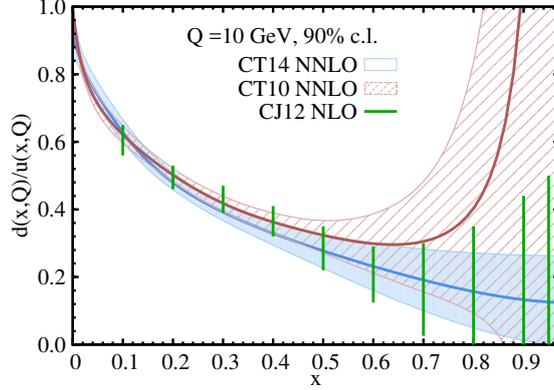


Figure 38: A comparison of the d/u ratio at $Q = 10$ GeV for the CT14 NNLO (solid blue), CT10 NNLO (dashed red), and CJ12 NLO (green lines) analyses, together with the corresponding 90% CL PDF uncertainties. From Ref. [31].

(LM) scan method discussed in Sect. 4.3.3 is also adopted for specific important observables or for PDFs in region poorly constrained by data. In the pre-CT10 analyses, it was found that within a global χ^2 tolerance of $\Delta\chi^2 = 100$ (for more than 2000 data points) the fits agree with all experiments at 90% CL. In later CTEQ-TEA analyses, a more efficient dynamic tolerance criteria is adopted to account for the agreement with individual data set. It is constructed from an equivalent Gaussian variable, e.g.,

$$S_n = \sqrt{2\chi^2(N_n)} - \sqrt{2N_n - 1}, \quad (119)$$

where N_n is the total number of data points in data set n and $\chi^2(N_n)$ represents the χ^2 of the fit to that data set. S_n follows a normal distribution given the number of data points is large enough. Thus a value of S_n greater than 1.3 will be excluded at 90% CL. Subsequently, a second layer of penalty is added to the global χ^2 when determining the boundaries of confidence intervals, called a Tier-2 penalty, defined as

$$P = \sum_{n=1}^{N_{\text{exp}}} (S_n/S_{n,\text{best}})^{16}, \quad (120)$$

where the sum runs over all data sets included and we normalize S_n to its value in the best-fit $S_{n,\text{best}}$ to account for the poor fit to certain experiments. The power of 16 is introduced so that the penalty will reach the tolerance of 100 as soon as any data set shows disagreement at 90% CL. The tolerance criteria then changes to $\Delta\chi^2 + P = 100$. Fig. 39 shows the distribution of $S_{n,\text{best}}$ for all 33 experiments included in CT14 analysis. The distribution is wider than a normal distribution, indicating the presence of disagreement, or tensions, between some of the included experiments, a well-established fact in global PDF fits.

With the best-fit and $2N_{\text{eig}}$ uncertainty eigenvectors sets, the asymmetric errors for any QCD observable \mathcal{F} can be calculated through the master formula:

$$\begin{aligned} (\delta\mathcal{F})_+ &= \sqrt{\sum_{i=1}^{N_{\text{eig}}} [\max(\mathcal{F}_{+i} - \mathcal{F}_0, \mathcal{F}_{-i} - \mathcal{F}_0, 0)]^2}, \\ (\delta\mathcal{F})_- &= -\sqrt{\sum_{i=1}^{N_{\text{eig}}} [\max(\mathcal{F}_0 - \mathcal{F}_{+i}, \mathcal{F}_0 - \mathcal{F}_{-i}, 0)]^2}, \end{aligned} \quad (121)$$

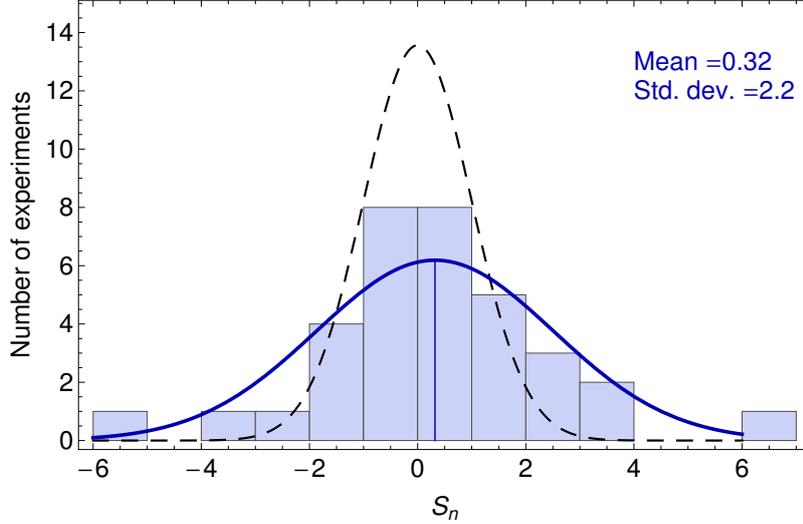


Figure 39: Best-fit values of the equivalent Gaussian variable S_n , Eq. (119), for all 33 experiments in the CT14 NNLO global analysis [31].

where \mathcal{F}_0 is the prediction from central set, and \mathcal{F}_{+i} and \mathcal{F}_{-i} are from two error sets in the direction of i -th eigenvector. These PDF errors can be scaled down to 68% CL with a factor of 1.64 assuming Gaussian distributions.

As mentioned above, the CTEQ-TEA analysis also uses the Lagrange multiplier method [430] to cross-check the error estimation from nominal Hessian sets. In the CT14 analysis, Lagrange multiplier scans have been performed for the cross sections of Higgs boson production via gluon fusion the top quark pair production at the LHC. In such scans the best-fits and the associated χ^2 are found for each fixed value of the observable studied. Then the PDF uncertainties on the observable are determined from the χ^2 profile obtained using the same tolerance criteria as in the Hessian method. Fig. 40 shows the good agreement of the 90% CL uncertainties for the Higgs cross sections from the CT14 Hessian PDFs and the CT14 LM scans. The latter can be read off from the intersection of the horizontal line $\Delta\chi^2 = 100$ and the various curves. As the LM method does not rely on the linear approximation, it serves as a robust check of the Hessian results.

Other dedicated applications of the CTEQ-TEA global analysis have been presented recently. The CTEQ-TEA analyses uses the world average of strong coupling constant $\alpha_s(m_Z)$ as an input. Usually the nominal fit is performed with $\alpha_s(m_Z) = 0.118$ at both NLO and NNLO, but additional fits with alternative α_s choices are also provided, which can be used to compute the combined PDF+ α_s uncertainties. The fit itself provides a much weaker constraint on α_s than the world average, see also the discussion in Sect. 4.5.1. Similarly, the pole mass of the charm and bottom quarks are chosen to be close to the world average values, with $m_c = 1.3$ GeV and $m_b = 4.75$ GeV.

The CTEQ-TEA group also provides specialized fits with non-perturbative charm quark PDFs [55, 466]. For these studies, a sea-like or valence-like charm distribution is added to the nominal parametrization and then fitted to the data. Limits on the momentum fraction carried by the fitted charm at the initial scale $Q_0 = 1.3$ GeV are then derived. In the most recent CT14 analysis, the limits are 1.6% for the SEA model and 2.1% for the BHPS model both at 90% CL [467]. Finally, there are also CT14 QED PDFs [50]

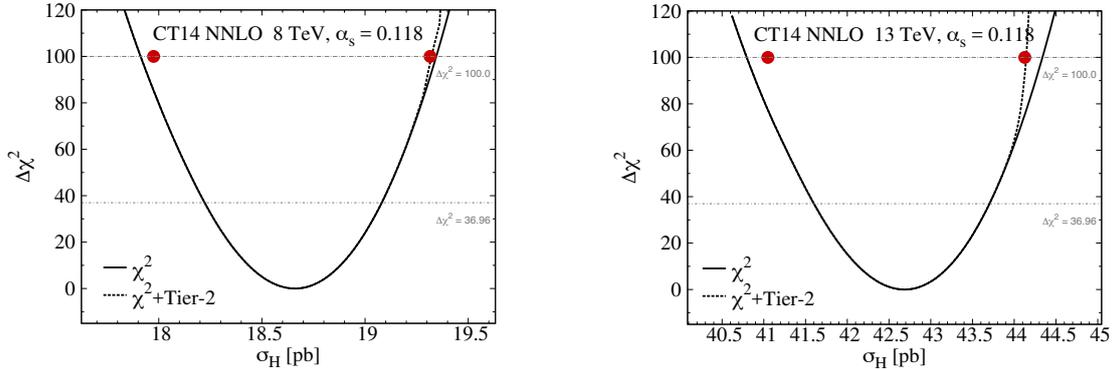


Figure 40: Dependence of the increase in χ^2 with respect to the global minimum χ_{\min}^2 , defined as $\Delta\chi^2 = \chi^2 - \chi_{\min}^2$, for the constrained CT14 fit on the expected cross section σ_H at the LHC 8 and 13 TeV [31], for $\alpha_s(m_Z)=0.118$. The solid and dashed curves are for the constrained fits without and with the Tier-2 penalties, respectively. The red dots correspond to the upper and lower 90% CL limits calculated by means of the Hessian method.

based on a radiative ansatz for the inelastic component of the photon PDFs. The 90% CL limit on the momentum fraction of the proton carried by the photon is found to be 0.11% at $Q_0 = 1.3$ GeV, as derived from a fit to the ZEUS measurement on isolated photon production [468]. Recent progress in PDF fits with QED corrections and the role of the photon PDF is discussed in Sect. 7.

5.2. MMHT

The MMHT14 PDFs [32] are the successor to the MSTW08 [124] set, which in turn derives from the earlier MRST and MRS studies. The first NLO fit [101] to DIS data was performed in the late 80s, while in the mid 90s the MRS(A) [107] fit was released, including data from HERA and the Tevatron for the first time. This corresponded to a truly global analysis, fitting to fixed target, DIS and hadroproduction data to constrain the PDFs as precisely as possible. Subsequent releases have all built on this approach, but with significant advances achieved over the years due to improvements in both theory and experiment, as well as with significant methodological developments. The MRST98 release [469] was the first set to include a full treatment of heavy flavours within the GM-VFNS developed in [120], and discussed further in Sect. 2.5. This was motivated by the new HERA measurements of the charm structure function, which demonstrated the importance of a consistent treatment of charm production at low and high scales; indeed, the introduction of this flavour scheme resulted in an improved description of such data. The MRST02 release [118] included a full treatment of PDF errors for the first time, described further below (see also Sect. 4.3). The MRST04 [470] set went to NNLO for the first time, although earlier fits based on approximate NNLO treatments in fact gave quite similar results. A proper treatment of the PDF discontinuities which occur at NNLO across heavy flavour thresholds was presented later in [471].

These elements were all incorporated in the major MSTW08 [124] release. This presented a global fit to a range of DIS data from HERA and fixed proton and nuclear targets, fixed target Drell-Yan and dimuon production and W , Z and jet production at the Tevatron, with $O(2500)$ data points in total. Fits were performed up to NNLO in the strong coupling, with an improved dynamical error treatment, and with an up to date heavy flavour scheme applied. This release aimed to provide a general-purpose PDF set for use at the LHC, which began operation soon after the release, and was subsequently very widely used in LHC phenomenological studies and experimental analyses. This fit was updated in the latest MMHT14 [32]

set, which includes a number of theoretical and experimental updates. In particular for the first time LHC data on W , Z , $t\bar{t}$ and jet production are included, as well as updated HERA data on the charged, neutral, charm and longitudinal structure functions, and updated Tevatron W and Z measurements. As in earlier fits, for DIS data a Q^2 cut of 2 GeV^2 and W^2 cut of 15 GeV^2 are imposed to avoid sensitivity to higher twist corrections.

In the case of Tevatron jet production, in the absence of a full NNLO calculation at the time, an approximation to the NNLO corrections based on the threshold corrections of [257] was applied in the NNLO MSTW08 fit, with the judgment being made that the difference between this and the full NNLO result would be expected to be smaller than the systematic uncertainties on the data, which itself provided the only direct constraint on the gluon at high x . At the LHC much of the jet data are quite far from threshold, while those that are not probe a kinematically similar region to the Tevatron data, and so at NNLO these are not included in the MMHT14 fit. For the $t\bar{t}$ data the top mass is allowed to be determined from the fit, with the pole mass value of $m_t = 172.5 \pm 1 \text{ GeV}$ taken as an input. This gives a value at NNLO that is consistent with the world average, while at NLO it is somewhat lower.

The MSTW PDFs were parameterised in terms of simple polynomials in x , with 29 free parameters. However, in [265] it was shown that this parameterisation was not sufficiently adaptive to describe the Tevatron and LHC W asymmetry data. In particular, it was necessary to introduce a more flexible basis for the interpolating function described in Sect. 4.1.1, with

$$I_f(x) = \sum_i^n \alpha_{f,i} T_i(y(x)), \quad (122)$$

where T_i is a Chebyshev polynomial of order i and $y(x) = 1 - 2\sqrt{x}$ is chosen so as to sample a wide range of x , and has the additional advantage that this provides a half-integer separation in powers of x , as expected on Regge theory grounds.

In order to determine how many parameters n were needed in Eq. (122), in [265] pseudo-data points with a constant percentage error were generated for the required distributions, in terms of a very large order polynomial with additional smoothness constraints applied. The fractional deviation from the true PDF, as well as the decrease in χ^2 , were then determined as the number of parameters were increased, until no further significant improvement was observed and the level of agreement was well below the PDF uncertainty for the set. In this way $n = 4$ was arrived at as a good choice with which to parameterise the u_V , d_V , $s + \bar{s}$ and light quark sea S distributions. Fitting to the same MSTW08 data set, these resulted in some improvement in the fit quality, but with the only significant change in the PDF being in the u_V at lower x . This was found to lie outside the previous PDF uncertainty band, and the additional flexibility provided a greatly improved description of W asymmetry data from the LHC Run I.

In the MMHT14 set, this Chebyshev parameterisation is used at $Q_0^2 = 1 \text{ GeV}^2$ for the u_V , d_V , $s + \bar{s}$ and light quark sea S distributions, while for the gluon a term with $n = 2$ Chebyshev is included, but with a second term still present, as in MSTW08, which has a different low x power and provides the additional flexibility at low x that is required by the HERA data; this has the effect that the gluon at NLO and NNLO can become negative at low x and Q^2 . Standard polynomial parameterisations are taken for the less constrained $s - \bar{s}$ and $\bar{d} - \bar{u}$ distributions, although as the data becomes more precise we can expect this to change.

A further improvement described in [265] that is included in the MMHT14 set is in the treatment of the non-perturbative corrections that should in general be applied when considering DIS data on deuteron targets, to account for the binding of the proton and neutron within the deuteron. While in MSTW08 and earlier fits, a fixed shadowing correction at small x was applied, a more flexible approach is now

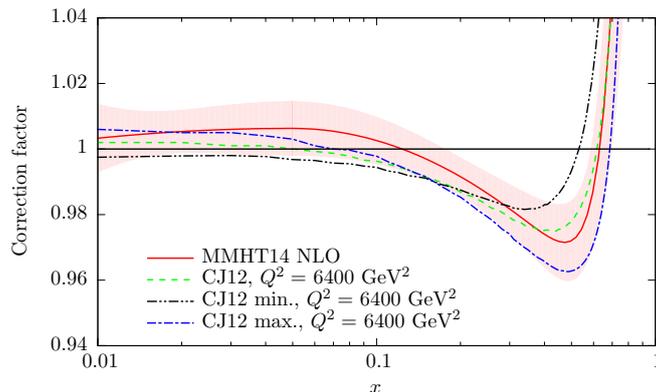


Figure 41: Fitted MMHT14 deuteron correction factors with uncertainty, compared to the CJ12 [464] predictions. Taken from [32].

taken. In particular the deuteron corrections are freely parameterised in terms of a function $c(x)$, which is determined along with its corresponding uncertainties from the PDF fit. This results in a significantly improved description of the BCDMS deuteron structure function data, the E866 Drell–Yan asymmetry and the Tevatron lepton asymmetry data, with some significant changes in d_V . The result of the MMHT14 fit is shown in Fig. 41 and compared against different model predictions used in the CJ12 [464] analysis. Interestingly, very good agreement is found with the CJ12mid prediction, demonstrating the power of global PDF fits to extract additional physical information beyond the PDFs themselves.

In MMHT14, the Hessian approach is applied to calculate the PDF errors, with the ‘dynamical’ tolerance criteria described in Sect. 4.3.1 taken. For MMHT14 the 68% uncertainties are calculated using this procedure. In the fit there are 37 free PDF parameters in total, however in the error determination certain parameter directions are found to be largely degenerate, leading to departures from quadratic χ^2 behaviour. This is corrected by fixing some parameters when calculating the error eigenvectors, reducing the number to 25, that is 50 directions.

Other theoretical updates in the MMHT14 include the treatment of the $D \rightarrow \mu$ branching ratio, which is required in the fit to dimuon production in DIS. This is now determined from the fit but with the measurement of [472], which is not determined from dimuon production data, included as a data point. The result is somewhat lower than that taken in MSTW08, corresponding to a larger strangeness, but the most dramatic effect is that the $\sim 10\%$ uncertainty on the branching ratio allows for a much larger strangeness uncertainty when fitting to the same data. Other smaller improvements include an updated treatment of nuclear corrections and a multiplicative, rather than additive, treatment of systematic uncertainties where appropriate.

In MMHT the value of the strong coupling is allowed to be determined by the fit, it being argued that valuable information can be provided from global PDF fits about this object. This in addition serves as a consistency test on the overall fit; if the extracted value is in strong tension with the world average then this might indicate that further work is needed. In the fit the preferred values at NLO and NNLO are indeed found to be consistent with the world average, and including this as an additional data point is not found to affect the fit significantly. In particular, the detailed study of [473] found best fit values of $\alpha_S(m_Z^2) = 0.1201 \pm 0.0015$ at NLO and $\alpha_S(m_Z^2) = 0.1172 \pm 0.0013$, to be compared with the world average value of $\alpha_S(m_Z^2) = 0.1181 \pm 0.0013$. The NNLO χ^2 profile for $\alpha_S(m_Z^2)$ and the corresponding individual constraints from the most constraining data sets are shown in Fig. 42. The PDF sets for a range of α_S values,

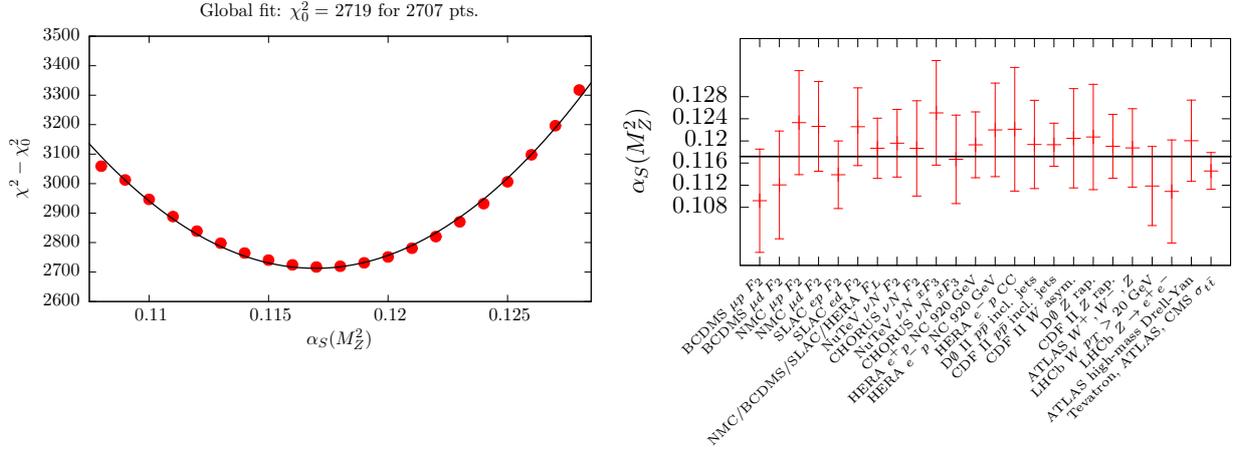


Figure 42: Left: global χ^2 values as a function of $\alpha_S(m_Z^2)$. Right: best fit $\alpha_S(m_Z^2)$ together with the upper and lower 1σ constraints from the most constraining data sets. Both figures correspond to NNLO fits and are taken from [473].

from 0.108 to 0.128 in steps of 0.001, are publicly available.

Concerning the treatment of the heavy quarks, the GM–VFNS is taken, applying the ‘optimal’ TR scheme of [474] which improves the smoothness of the transition region where the number of active flavours is increased by one. The charm and bottom pole masses are fixed to $m_c(m_b) = 1.4(4.75)$ GeV, but a detailed study of the heavy quark mass dependence was performed in [440]. In particular, when m_c and m_b were taken as free parameters in the fit, the preferred values were found to be somewhat lower than these defaults, possibly suggesting a $\overline{\text{MS}}$ scheme may be more appropriate, although the impact of such a choice is expected to be fairly small. Sets with a range of charm and bottom masses, as well as with a maximum number of active flavours $n_f^{\text{max}} = 3$ or 4, are also made available.

A subsequent study within the MMHT framework on the PDF impact of the final HERA I + II combined data [34], which was released following MMHT14, was performed in [475]. This was found to lead to some reduction in the PDF uncertainties, principally in the gluon, with little change in the central values. It was therefore decided not to release an updated set but rather to wait until more precise and varied LHC data became available, as well as theoretical calculations such as NNLO jet production. Subsequent work towards a new PDF set has been presented in for example [476], where a first fit at NNLO to jet data is presented, while the impact of new LHC data in the fit is seen to be significant. In addition, the first steps towards the inclusion of the photon PDF within the MMHT framework are presented, see Sect. 7 for more discussion. A further public release is therefore anticipated in the near future.

5.3. NNPDF

As discussed in Sect. 4, the NNPDF fitting methodology is based on the combination of three main components: i) the use of artificial neural networks as universal unbiased interpolants, ii) the Monte Carlo method to estimate and propagate PDF uncertainties, and iii) Genetic Algorithms (GA) minimization for the training of the neural networks combined with a look–back cross–validation stopping to avoid over–fitting. Here we review the historical development of the NNPDF family of PDF fits, highlighting recent progress.

The NNPDF methodology was first presented in [131], where it was used to produce a neural network based determination of the proton, deuteron and non–singlet DIS structure functions from the NMC and BCDMS fixed target data. As a first phenomenological application, this determination was used to

extract the strong coupling constant $\alpha_S(m_Z)$ from scaling violations of truncated moments of structure functions [477]. This analysis was subsequently extended [428] to include as well the F_2^p measurements from the H1 and ZEUS experiments at the HERA collider. Note that a determination of structure functions is purely data-driven, with no theoretical input required at this point.

When moving from fitting structure functions to PDFs, there are a number of simplifications, for instance one needs to fit only a 1D function $q_i(x, Q_0)$ as opposed to a 2D function $F_2^p(x, Q^2)$, but also technical complications, the most important one being the need to compute DIS structure functions starting from a neural-network based parametrization of $q_i(x, Q_0)$. First of all, the usual ANN training algorithm of back-propagation cannot be used in this case, due to the convolution of the PDFs with the DGLAP evolution kernels and the DIS coefficients functions. To overcome this limitation, it was demonstrated how Genetic Algorithms can be efficiently used for ANN training under a non-trivial mapping between the latter and the experimental data, and used to extract the QCD vacuum condensates from hadronic tau decay data [423]. An efficient method to solve the DGLAP evolution equations in N -space was also developed, called the Fast Kernel method. With these ingredients at hand, it became possible for the first time to apply the NNPDF methodology to a determination of the PDFs, starting from a fit of the non-singlet combination $q_{NS}(x, Q_0)$ [427] and then moving to a first full-fledged NLO PDF fit based on NC DIS structure function data [478] in the NNPDF1.0 analysis.

Subsequently, the global NNPDF fits were improved both by adding new experimental data, updating the theoretical calculations and/or refining the fitting methodology. The NNPDF1.2 analysis [216] relaxed the previous assumption that the strange sea was proportional to the light quark sea, $s = \bar{s} = \kappa(\bar{u} + \bar{d})$, and parametrized both s^+ and s^- using neural networks, exploiting the constraints from the NuTeV dimuon CC neutrino scattering data. Two important phenomenological consequences of this analysis were, first of all, the demonstration that the PDF uncertainties associated with s^- were enough to remove the NuTeV anomaly [479] in the determination of the weak mixing angle $\sin^2 \theta_W$; and second, a direct extraction of the CKM matrix element V_{cs} with a precision compatible with that of the PDG average.

In 2010, the NNPDF2.0 set was released [132], which constituted the first truly global PDF fit from the NNPDF collaboration. In addition to the NC and CC DIS structure function data included in previous releases, this fit in addition fit to fixed-target Drell-Yan cross sections from the Fermilab E605 and E866 experiments, inclusive jet production measurements from CDF and D0 at the Tevatron, as well as the differential rapidity distributions of the Z boson also from the Tevatron. From the theoretical point of view, NNPDF2.0 was still based on the ZM-VFNS, and thus charm and bottom structure function data from HERA were not included. A good overall description of all experiments in the global fit was found. NNPDF2.0 was also the first global PDF set to include the recently release HERA combination of H1 and ZEUS structure function data for the Run I data period [129].

While NNPDF2.0 demonstrated that the NNPDF methodology could be successfully applied to a global determination of parton distributions, there were still a number of important limitations from the theoretical point of view. First, the use of a ZM-VFNS neglected heavy quark mass effects in the DIS structure functions, which were known to be important for the description of the low- x , low- Q^2 HERA data. Second, all NNPDF fits so far were based on NLO theory, and NNLO accuracy was essential to match the corresponding precision of important partonic hard-scattering cross sections such as Higgs production in gluon fusion. The first of these theory limitations was removed with the release of NNPDF2.1 [133], which was based on the FONLL GM-VFNS for the calculation of DIS structure functions, which allowed the HERA charm and bottom structure function data to be fit. This analysis also showed that the impact of heavy quark mass effects was less drastic than previously reported, with the cross section predictions between NNPDF2.0 and 2.1 typically agreeing at the one-sigma level. The NNPDF2.1 fit was also used to produce a determination

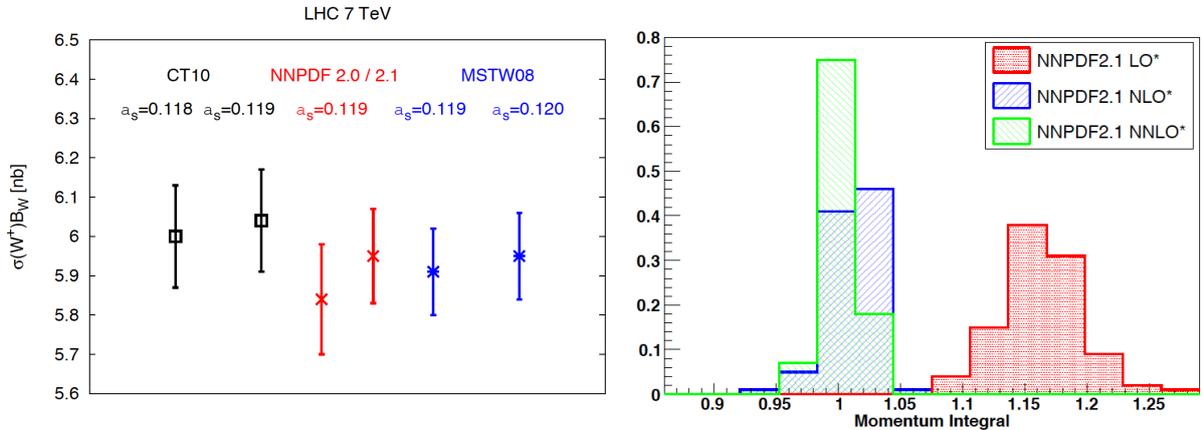


Figure 43: Left plot: comparison between the NNPDF2.0 and 2.1 predictions for the inclusive W^+ production cross section at the LHC 7 TeV, which illustrates the phenomenological impact of heavy quark mass effects. Right plot: the distribution of the momentum integral Eq. (123) among the MC replicas for the variants of the NNPDF2.1 LO, NLO, and NNLO fits that do not impose the momentum sum rule.

of the strong coupling $\alpha_s(m_Z)$ from the global dataset [438].

The second of these theoretical limitations was removed shortly after, with the release of a NNLO version of NNPDF2.1 [134]. This PDF set was based on the same dataset as its NLO counterpart, but with the DIS and hadronic cross sections computed at NNLO, in the former case using the FONLL–C GM–VFNS. In the same publication, the first NNPDF LO sets were also presented. The availability of NNPDF2.1 fits at LO, NLO and NNLO allowed a systematic study of the perturbative convergence of the global fit, finding in particular reasonable agreement at the one–sigma level between the NLO and NNLO versions. The consistency of the global QCD analysis framework was also tested by performing fits without imposing the momentum sum rule and then verifying a posteriori that the global fit result was consistent with QCD expectations, finding indeed that at NNLO

$$[M] \equiv \int_0^1 dx (g(x, Q^2) + \Sigma(x, Q^2)) = 1.002 \pm 0.014. \quad (123)$$

The NNPDF2.1 NNLO analysis was also used to perform a determination of the strong coupling constant [437], finding a value $\alpha_s(m_Z) = 0.1173 \pm 0.0007^{\text{stat}} \pm 0.0009^{\text{pert}}$, a result which is included in the current PDG global average of α_s [435].

Similarly to the early 90s, when the availability of the HERA structure function measurements became a game–changer for global fits, from 2010 onwards the LHC experiments have started producing a wealth of PDF–sensitive information for global fits. With this motivation, in 2012 the NNPDF2.3 set was released [135], and was the first PDF fit to include LHC data, in particular electroweak gauge boson production from ATLAS, CMS and LHCb as well as jet production from ATLAS. As with all subsequent releases, NNPDF2.3 was available at LO, NLO and NNLO and was based on the FONLL general mass scheme. The NNPDF2.3 set became the baseline PDF set in several popular MC event generators, such as Pythia8 [480] and aMC@NLO [414]. Based on NNPDF2.3, an study of theoretical uncertainties affecting the PDF from sources such as deuteron corrections and higher twists was presented in [481].

Following the release of NNPDF2.3, it was realised that the increase in complexity required to include the many new experiments that were either available or about to be released could not be satisfactory tackled

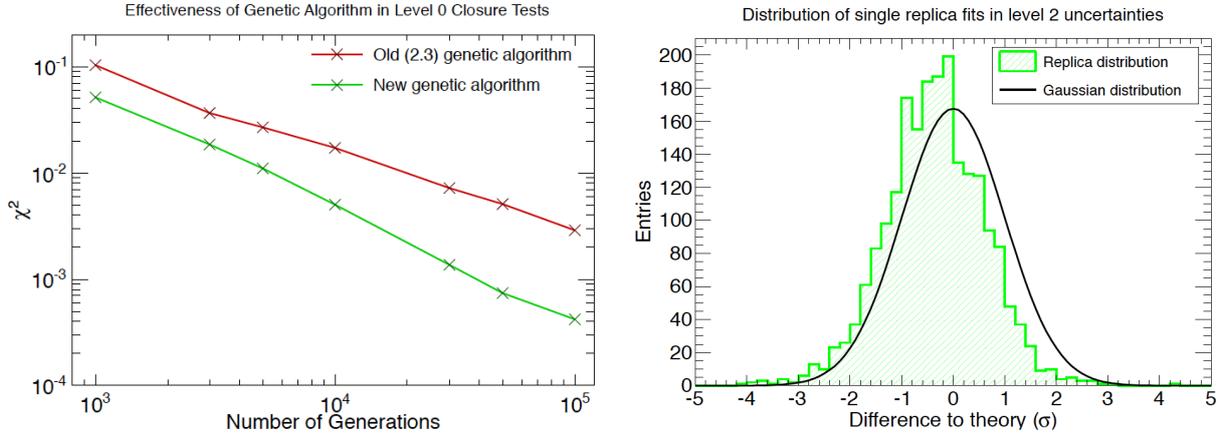


Figure 44: Representative results of the closure tests presented in the NNPDF3.0 analysis. Left plot: in a level 0 closure tests, where the pseudo–data is generated without any statistical fluctuations, the χ^2 should decrease monotonically as a function of the number of GA iterations, down to arbitrarily small values. Right plot: the distribution of the difference between theory and data in units of the error of the latter among each of the Monte Carlo replicas. This distribution is consistent with the Gaussian expectations.

with the current, FORTRAN77–based code. With this motivation, a complete rewriting of the NNPDF global analysis framework into C++ and Python was undertaken, a two–year long effort that culminated with the release of the NNPDF3.0 set [30]. In addition to including many new LHC experiments on jets, vector boson production, W +charm and top production, the main result of NNPDF3.0 was the systematic validation of the complete fitting methodology based on statistically robust closure tests.

In these closure tests, pseudo–data was generated based on some “truth” PDFs, and then a PDF fit was performed: if the resulting PDF central values and uncertainties were consistent with the (known) input PDFs, then the closure test can be considered successful. In Fig. 44 we show representative results of the closure tests presented in the NNPDF3.0 analysis. In the left plot we show the results of a level 0 closure test, where the pseudo–data is generated without any statistical fluctuations, and as expected the χ^2 is found decrease monotonically as a function of the number of GA iterations, down to arbitrarily small values. And in the right plot we show the distribution of the difference between theory and data, in units of the error of the latter among each of the MC replicas. This is consistent with the expectations from an underlying Gaussian distribution.

A recent development in the NNPDF family of global analyses concerns the treatment of the charm PDF. In previous PDF sets, it was assumed that the charm PDF $c(x, Q)$ was generated dynamically from the gluons and light quarks, as dictated by the DGLAP evolution starting from the charm mass threshold $\mu_c \simeq m_c$. However, a possible non–perturbative component of the charm PDF would invalidate this assumption, a hypothesis which can ultimately be tested against experimental data. In addition, treating the charm PDF on an equal footing with the gluon and light quark PDFs offers other potential advantages, such as a reduced dependence on the value of m_c and an improved agreement between data and theory from the more flexible input PDF parametrization.

With this motivation, a variant of the NNPDF3.0 analysis with a fitted charm PDF was studied in [56]. By parametrising the charm PDF with an artificial neural network with 37 free parameters, it was found that fitting charm leads to an improved χ^2 for several experiments and stabilises the dependence of the fit with respect to the value of m_c . It was also observed that fitting charm allows a reasonable description of

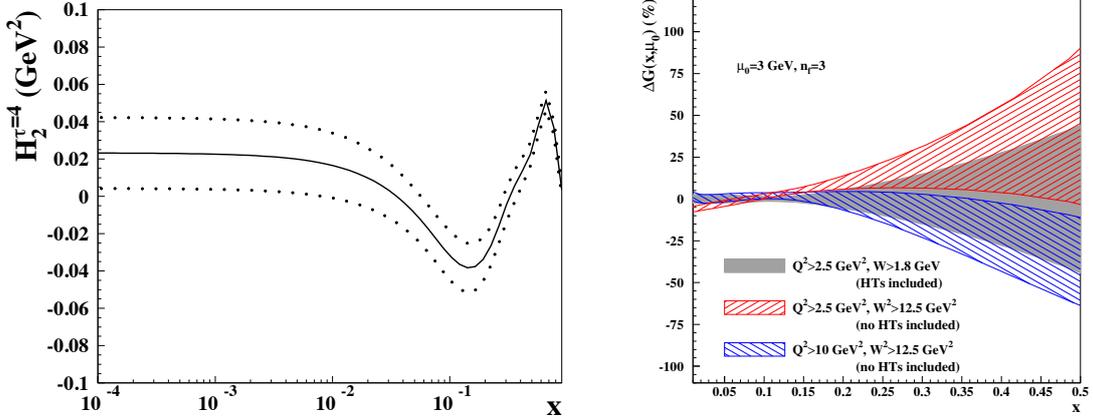


Figure 45: Left: the higher twist coefficient for the F_2 structure function from the ABMP16 fit, including the associated $1\text{-}\sigma$ uncertainty. Right: percentage difference in ABMP $n_f = 3$ gluon distribution between the default result and fits performed with higher W^2 cuts but without the higher twist corrections. The $1\text{-}\sigma$ uncertainty bands are also shown. Taken from [33].

the EMC charm structure function data, something which is not possible if perturbative charm is used. This result was shown to be robust with respect to variations of the treatment of the EMC data, such as increasing some of its systematic errors or introducing additional kinematical cuts [482].

The resulting charm PDF can then be compared with non-perturbative models [483], and some intriguing evidence for a large- x non-perturbative charm component in the proton was found. Predictions for a number of LHC process such as Z +charm and large- p_T D meson production were also provided, showing the potential of Run II data to disentangle the charm content of the proton. See Sect. 6.4 for more details about the charm content of the proton. As discussed there, the main qualitative features of the comparison between fitted and perturbative charm [56] were confirmed by the subsequent NNPDF3.1 set [169] and shown to be independent of whether or not the EMC data is included in the fit.

The most recent incarnation of the NNPDF global analysis is the NNPDF3.1 set. The main motivation for this release was the availability of a large number of high-precision collider measurements (and the corresponding NNLO calculations) providing PDF-sensitive information, including some that for the first time could be used in a PDF fit, such as differential distributions in top quark pair production and the p_T of Z bosons. The second main motivation was to provide a state-of-the-art PDF set without assuming that charm is generated perturbatively, that is, providing baseline global PDF fits with fitted charm. Some of the new experiments included in NNPDF3.1 were the $t\bar{t}$ distributions from ATLAS and CMS, the legacy LHCb inclusive W and Z measurements from Run II, the D0 W asymmetries in the muon and electron channel, the p_T of Z bosons from ATLAS and CMS at $\sqrt{s} = 8$ TeV, as well as several other inclusive gauge boson and jet production measurements from ATLAS and CMS.

5.4. ABM

The ABMP16 [33] set is the latest PDF fit from the ABM collaboration, following on from the previous ABM11 [484], ABM12 [136], and ABKM09 [126] sets. These PDF sets were in turn based on the earlier fits of Refs. [112, 127, 128] to HERA and fixed proton and deuteron target DIS data, with the ABKM09 fit [126] and those that follow it including in addition fixed target Drell-Yan and dimuon production data

from neutrino DIS on fixed nuclear targets. The ABM sets of parton distributions are parameterised in terms of polynomials in x , with the latest fits including 25 free parameters. In the context of this fit to a somewhat reduced dataset, the use of the classical ‘ $\Delta\chi^2 = 1$ ’ criteria for determination of the PDF errors is applied. All sets from ABKM09 onwards go to NNLO in the strong coupling.

Two notable features of these fits are the use of a purely FFNS for the charm and bottom quark contributions to DIS structure functions using the $\overline{\text{MS}}$ running masses [213], and the treatment of higher-twist effects. In the latter case, no attempt is made to impose a cut to remove the region of sensitivity to such effects. Rather, a lower cut ($W > 1.8$ GeV) is imposed for the DIS data than is typically applied in other PDF fits. The structure functions are then given by

$$F_i(x, Q^2) = F_i^{\text{TMC}}(x, Q^2) + \frac{H_i^{\tau=4}}{Q^2}, \quad (124)$$

where $i = 2, T$. Thus x dependent and Q^2 independent twist-4 corrections H_i^4 are included. While the effect of these dies off with increasing Q^2 , at lower scales they can have a significant effect. These are then parameterised in terms of cubic splines defined at x_k ($k = 1..7$) points roughly linearly spaced between $x = 0$ and 1, which are then determined from the fit. The result for the F_2 correction is shown in Fig. 45 (left), and is found to be inconsistent with zero, in particular at higher x . The effect of these corrections, and of conversely omitting them and including a more stringent W^2 cut on the DIS data is shown in Fig. 45 (right) for the extracted gluon PDF. The fit with the cut of $W^2 > 12.5$ GeV² and no higher twist corrections is found to prefer a somewhat larger gluon at higher x , and in some regions lies outside the 1- σ uncertainty band of the default fit.

In addition to the higher twist corrections included in Eq. (124), the structure function functions also include target mass corrections, that is the impact of terms $\sim M_N^2/Q^2$, where M_N is the nucleon mass. These are taken into account according the Georgi–Politzer prescription [485] (see also [484]), with

$$F_2^{\text{TMC}}(x, Q^2) = \frac{x^2}{\xi^2 \gamma^3} F_2(x, Q^2) + 6 \frac{x^3 M_N^2}{Q^2 \gamma^4} \int_{\xi}^1 \frac{d\xi'}{\xi'} F_2(\xi', Q^2), \quad (125)$$

where $\xi = 2x/(1 + \gamma)$ and $\gamma = (1 + 4x^2 M_N^2/Q^2)^{1/2}$, and a similar result holds for F_T . Thus, as $Q^2 \rightarrow \infty$ the corrected F_i^{TMC} reduce to the regular F_i . The same TMC are also included in other PDF analyses, such as in the NNPDF family [478].

As mentioned above, the ABMP fit uses a purely fixed flavour scheme to describe the DIS structure function data. That is, the fit is performed with $n_f = 3$ light quark PDFs and with the heavy c, b treated as massive final-state partons which can be produced at order NLO and higher, regardless of the value of Q^2 . It is then argued that the bulk of the DIS data can be described within this scheme. The Tevatron and LHC collider as well as fixed-target DY data, on the other hand, for which $\mu_F^2 \gg m_{c,b}^2$, is treated using a $n_f = 5$ flavour set evolved from the same input by means of the NNLO matching conditions [126]. The use of a FFNS as opposed to a GM-VFNS for the DIS structure functions has been shown to be one of the dominant reasons for the differences between the ABM family and the global fits [165, 481]. PDF sets for $n_f = 3, 4$ and 5 active flavours are made publicly available, with the caveat that each set should be used within its specified validity regime, for instance the ABMP16 $n_f = 5$ set can only be used in applications for which $Q \geq m_b$.

A further feature of note in the ABM family of PDF sets is that the strong coupling α_S is always determined from the fit together with the PDFs. In the ABM11 analysis, this was found to be $\alpha_S(m_Z^2) = 0.1134 \pm 0.0011$ at NNLO, that is in some tension with the PDG world average value [435] of $\alpha_S(m_Z^2) = 0.1181 \pm 0.0013$ (the dominant uncertainty in which is determined by the lattice QCD input) used by the CT

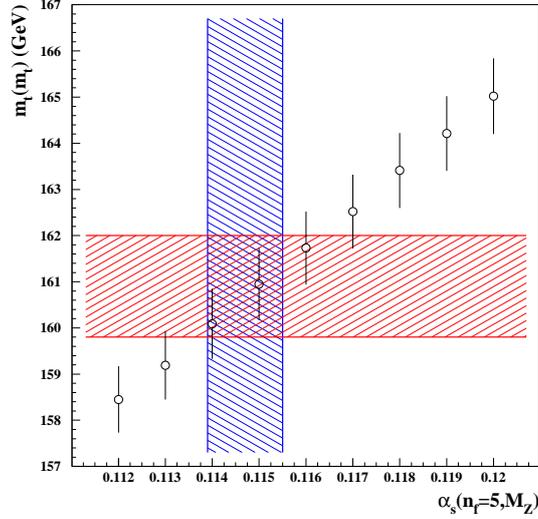


Figure 46: The $\overline{\text{MS}}$ value of the top quark mass $m_t(m_t)$ obtained in the ABMP16 fit for variants of $\alpha_s(m_Z^2)$ (data points), compared to the corresponding best-fit values (hatched bands). Taken from [33].

and NNPDF collaborations, and the value extracted by MMHT. While in [33, 484] the omission of higher twist corrections is found to lead to a sizeable increase in α_s , in contrast in [165] the use of the FFNS within the MSTW framework is found to lead to a smaller extracted value, consistent with that seen by ABM(P), while higher twist effects are found to have less of an impact.

In the ABM11 fit [484], the heavy quark masses $m_{c,b}$ were included in the $\overline{\text{MS}}$ scheme for the first time, in contrast to other PDF fits which adopt instead the pole mass scheme. This allows the heavy quark mass values to be constrained directly from the PDG results without relying on the perturbative transformation between the $\overline{\text{MS}}$ and pole masses, which is known to be poorly convergent. Thus, the quark masses are left free in the fit but with the PDG values added in as pseudo-data points. For the charm mass $m_c(m_c)$, the DIS data included in this fit are then found to give a comparable error to the PDG average.

The ABM12 fit [136] included hadron collider data for the first time, with a range of LHC W and Z boson measurements fit, while a comparison to top quark pair production data from the LHC and the Tevatron was also presented. The latest ABMP16 fit [33] includes an increased LHC set, with single top for the first time, as well as Tevatron lepton asymmetry data. In addition, the HERA I+II combined data set and updated NOMAD and CHORUS data on dimuon production are fit. For the $t\bar{t}$ data, the top quark mass m_t is treated in the $\overline{\text{MS}}$ scheme and is determined from the fit, giving $m_t(m_t) = 160.9 \pm 1.1$ GeV. This is consistent with the PDG value of $160.0^{+4.8}_{-4.3}$ GeV, although this clearly has quite a large uncertainty, as it based on a single Tevatron measurement. Converting to the pole mass, they find $m_t^{\text{pole}} = (170.4 \pm 1.2)$ GeV, which is in some tension with the PDG world average value of (173.1 ± 0.6) GeV [435]. The result is shown in Fig. 46, with the masses extracted at different $\alpha_s(m_Z^2)$ also given. The correlation between m_t and α_s is clear; as discussed in [33], further information on this correlation can be provided by considering single top production data.

Interestingly, the extracted value of the strong coupling from the ABMP16 analysis, $\alpha_s(m_Z^2) = 0.1147 \pm 0.0008$, is somewhat larger in this fit than in the previous sets, due dominantly to the HERA I+II combined data, as well as by the increased number of LHC experiments, although this is still lower than the world

average value. The ABMP16 PDFs are available at NNLO for 3,4 and 5 fixed flavours, and for a range of $\alpha_S(m_Z^2)$ values, from 0.112–0.120 in steps of 0.001, in the $n_f = 5$ flavour scheme, including the corresponding Hessian eigenvectors. Note that, as explained, in Sect. 4.5, in the ABM fits the uncertainties due to the value of $\alpha_S(m_Z^2)$ and the heavy quark masses are also included when determining the Hessian eigenvectors.

5.5. CTEQ-JLab (CJ)

The CTEQ-Jefferson Lab (CJ) collaboration has performed a series of global PDF analyses [184, 185] with the latest PDF set being CJ15 [36], following the previous CJ12 set [464]. The analyses are carried out at NLO in QCD only, and focus on utilising DIS data at the highest- x values applicable to a perturbative QCD treatment. The kinematic selection cuts are chosen to be $Q^2 > 1.69 \text{ GeV}^2$ and $W^2 > 3 \text{ GeV}^2$ so as to keep data points at low Q and high x , unlike the more restrictive cuts used by other analyses. This results in about 1300 more data points from proton and deuteron targets, roughly a 50% increase in comparison to standard cuts. These additional data points provide useful information on the PDFs at large x , into the $x \gtrsim 0.7$ region where the constraints for most global analyses are indirect or purely from extrapolation. In particular, the deuterium data can improve on the determination of d quark at large x , for which the proton DIS data are less sensitive.

For the treatment of heavy quark mass effects in DIS structure functions, CJ12 uses a ZM-VFNS with heavy quark masses implemented as the flavour thresholds. CJ15 uses a more adequate S-ACOT [150] GM-VFNS to better describe data over a wide kinematic range, including the threshold regions. It is found that the implementation of the GM-VFNS leads to large changes in the gluon PDF at large x . Going to low Q and large x involves further complications to the theoretical predictions for the DIS structure functions, as finite Q^2 corrections to the leading-twist calculation, i.e. power corrections of $\mathcal{O}(1/Q^2)$, must be taken into account. In particular, the CJ analyses adopt the standard Operator Product Expansion (OPE) expression for the target mass corrections (TMCs) which allows structure functions at finite Q^2 be expressed in terms of their massless ($M^2/Q^2 \sim 0$) values through the scaling variable $\rho^2 = 1 + 4x^2M^2/Q^2$ [485, 486]. Other subleading $1/Q^2$ corrections, including higher twists, are parametrized by a phenomenological function form [36].

Another important aspect of the CJ analyses concerns the nuclear corrections for processes with deuteron targets, which become significant in the intermediate and large- x region and are equally important for low and high Q values. The nuclear corrections account for Fermi motion, binding, and nucleon off-shell effects, and can be implemented as convolutions with nuclear smearing functions. In the CJ12 analysis, three PDF fits are provided with different models of deuteron corrections, CJ12min, CJ12mid and CJ12max, corresponding to mild, medium, and strong corrections, respectively. These corrections are only applied at the level of structure functions.

The CJ15 analysis employs a phenomenological parametrization for part of the deuteron corrections with free parameters fitted to data, reducing the model dependence and increasing the flexibility of the fit. Moreover, these deuteron corrections are formulated at the parton level and can therefore also be applied to non-DIS processes. For example, the total quark PDF in deuteron can be written as $q^d = q^{d(\text{on})} + q^{d(\text{off})}$, with the on-shell and off-shell components given by [487, 488],

$$\begin{aligned} q^{d(\text{on})}(x, Q^2) &= \int \frac{dz}{z} f^{(\text{on})}(z) q^N(x/z, Q^2), \\ q^{d(\text{off})}(x, Q^2) &= \int \frac{dz}{z} f^{(\text{off})}(z) \delta f^N(x/z, Q^2) q^N(x/z, Q^2), \end{aligned} \quad (126)$$

where $q^N(x, Q^2)$ corresponds to the quark PDF in a free nucleon. The on-shell and off-shell smearing functions $f^{(\text{on})}$ and $f^{(\text{off})}$ can be calculated systematically within the weak binding approximation, using the

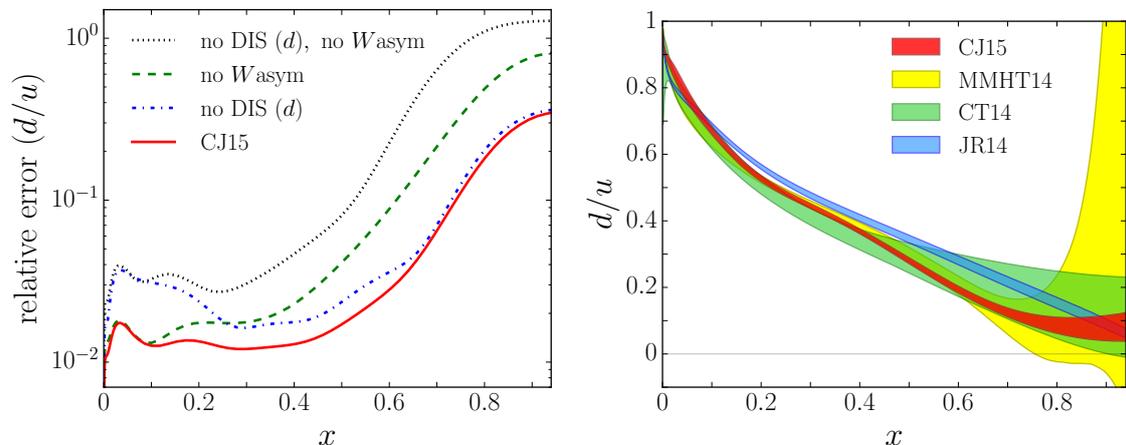


Figure 47: Left plot: relative error (90% C.L.) on the d/u PDF ratio as a function of x at $Q^2 = 10 \text{ GeV}^2$ from the CJ15 fit, compared with uncertainties obtained in fits excluding certain data sets [36]. Right plot: comparison of the d/u ratio at $Q^2 = 10 \text{ GeV}^2$ for different PDF sets with PDF uncertainties shown fat 90% C.L. except for MMHT14, which corresponds to 68% C.L. [36]

deuteron wave functions [489, 490]. The nominal CJ15 PDF fit is based on AV18 wave functions [491], but alternative fits with CD-Bonn [492], WJC-1 and WJC-2 [493] wave functions are also provided. The off-shell correction $\delta f^N(x)$ in Eq. (126) is parametrized as [489]

$$\delta f^N(x) = C(x - x_0)(x - x_1)(1 + x_0 - x). \quad (127)$$

The two zeros x_0 , x_1 and the normalization C are free parameters fitted to data with the constraint of maintaining the total number of valence quarks in the nucleon. It was found that different wave function models give similar quality of fits to the global data set and result in changes in the PDFs that are well within the uncertainties, as their differences can be largely compensated by the parametrization of the off-shell corrections.

In the CJ15 analysis, in addition to the deuteron data, two new data sets are found to place significant constraints on the d quark PDF at large x . These are the measurement of the F_2 structure function of a nearly free neutron inside a deuteron nucleus from the BONuS experiment [183, 494] at Jefferson Lab using a spectator tagging technique, and the lepton and reconstructed W boson charge asymmetry measurements from D0 Run 2 with full luminosities [270, 271]. Fig. 47 (Left) shows the impact of different data sets on the 90% C.L. PDF uncertainty of the d/u ratio in various CJ15 fits. It was found that at $x \lesssim 0.3$ the DIS data from deuteron target can reduce the PDF uncertainty on d/u by almost 50%. For $x \gtrsim 0.3$ the W asymmetry data provides the dominant constraint. The constraint from deuteron DIS data dies out for $x \gtrsim 0.6$, as in this region the fit is mostly sensitive to the deuteron off-shell corrections. In Fig. 47 (Right) the comparison of the d/u ratios for the CJ15, MMHT14, CT14 and JR14 PDFs is shown. These are in good agreements within PDF uncertainties as x goes to 1. The CJ15 PDF set has smaller PDF error on d/u in general, with an extrapolated value

$$d/u \xrightarrow{x \rightarrow 1} 0.09 \pm 0.03, \quad (128)$$

at 90% C.L., due to the new data sets that constrain the d quark at large x . Note that the fact that the d/u ratio has a finite $x \rightarrow 1$ limit occurs by construction in the CJ15 and CT14 parametrizations [416]. MMHT14 and

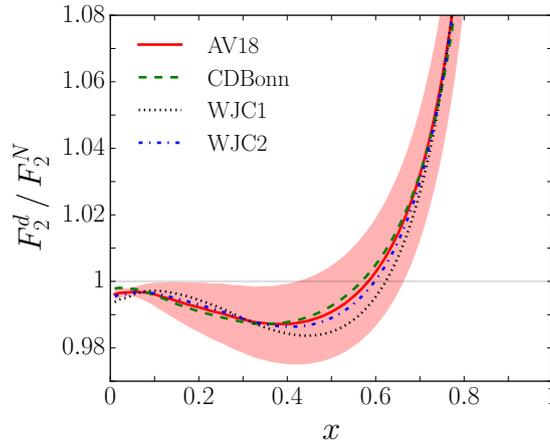


Figure 48: Ratio of deuteron to isoscalar nucleon structure functions F_2^d/F_2^N at $Q^2 = 10 \text{ GeV}^2$ for the CJ15 fits with different models of wave functions [36]; the colored band is the 90% C.L. error for CJ15 nominal fit, obtained using the AV18 nuclear wave functions.

NNPDF3.0 do not impose the exact isospin symmetry when x goes to 1, namely $b_{uv} = b_{dv}$ (see Sect. 4.1.1 for definition), and the predicted d/u can range from zero to infinity [416].

The CJ15 analysis is able to pin down the deuteron corrections through the interplay of additional data sets that are less sensitive to the nuclear corrections, i.e., the D0 W asymmetry data and BONuS data, and the deuteron DIS structure function data. This is illustrated in Fig. 48, which gives the deuteron to isoscalar nucleon ratio F_2^d/F_2^N from the CJ15 fits with different input wave functions together with the 90% C.L. uncertainty of the CJ15 nominal fit shown in the coloured bands. Significant corrections are found for $x \gtrsim 0.7$. The ratio turns out to be insensitive to the choice of wave functions, since it is only the combination of the wave function and the off-shell corrections that are constrained by current data.

5.6. HERAFitter/xFitter

For many years, the H1 and ZEUS collaborations performed QCD analyses of their structure function data, first separately and then together based on the H1+ZEUS combined datasets. The backbone of these analyses was the neutral- and charged-current inclusive structure function measurements, in some cases supplemented by the charm production structure functions and DIS jet cross sections. The main results from these studies were the HERAPDF family of PDF fits, which include HERAPDF1.0 [129], based on the Run I data, and HERAPDF2.0 [34], based on the final combination of inclusive measurements from Runs I+II. In Fig. 49 (Left) we show the results of the HERAPDF2.0 analysis for the u_V , d_V and S quarks and the gluon. The central value of the HERAPDF2.0A.G set, corresponding to an alternative gluon parameterisation, is also shown.

In the HERAPDF methodology, the total PDF uncertainty is divided into three types of errors: experimental uncertainties, propagated from the statistical and systematic uncertainties in the fitted data, model uncertainties, for instance due to α_s and Q_0 variations, and parametrization uncertainties, reflecting the spread from different comparable choices of input functional form for the PDFs. Note that in principle, for a sufficiently flexible input parametrization, it should be possible to neglect the model uncertainties related to the choice of Q_0 as well as the parametrization uncertainties.

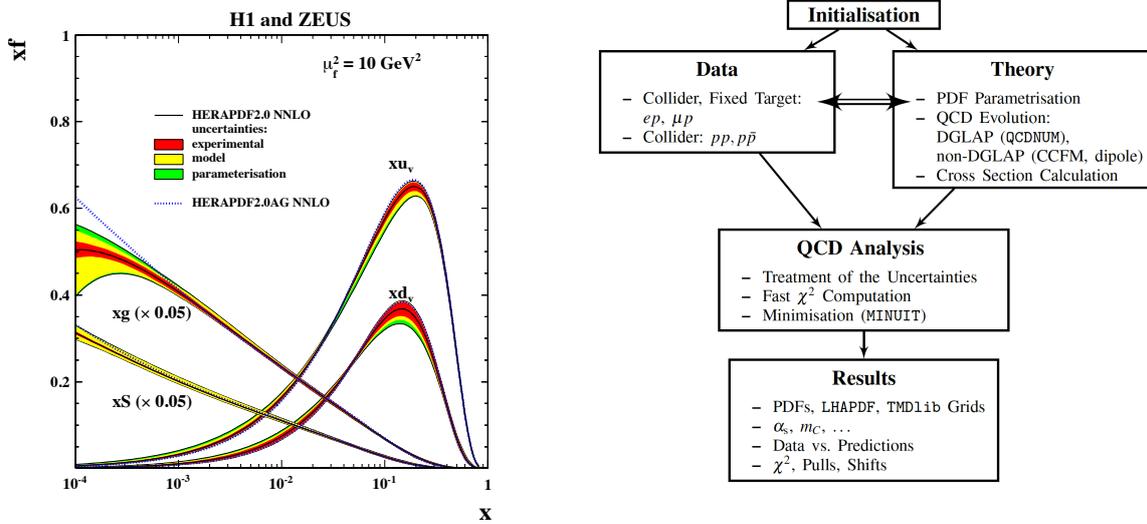


Figure 49: Left plot: the HERAPDF1.0 determination of parton distributions, based on the analysis of the combined HERA structure functions from Run I. Right plot: schematic representation of the xFitter code structure, see text for more details.

The expertise developed by these QCD analyses of the HERA structure function data lead to the development and release of HERAFitter [57], a publicly available open-source PDF fitting toolbox. This was developed as an extension of the H1 and ZEUS internal PDF fitting codes, that were extensively tested and applied in various QCD analyses of HERA inclusive and charm data, including the HERAPDF sets. Despite its name, HERAFitter was not restricted to the analysis of HERA data, and could also be used for the PDF interpretation of measurements from fixed-target DIS and proton-proton collisions. The flexibility of this open-source software tool also allows QCD analyses beyond unpolarized fixed-order PDF fits to be performed, such as fits of transverse momentum dependent (TMD) parton distributions and fragmentation functions. Recently, HERAFitter was renamed xFitter, to emphasise that this code is a general fitting toolbox, not necessarily related to or involving the analysis of HERA inclusive structure function data.

The xFitter framework includes modules that allow for various theoretical and methodological options, and is capable of fitting to a large number of data sets from HERA, Tevatron and LHC. For example, polarized and unpolarized PDF evolution can be performed using either APFEL or QCDNUM, and a number of fixed and variable flavour number schemes are implemented, such as the FFN scheme from OpenQCDrad and the S-ACOT, TR and FONLL general-mass schemes. In addition to PDF fitting, a large number of other functionalities are available, such as the approximate inclusion of new datasets in existing PDF sets by means of either Bayesian reweighting and Hessian profiling (see Sect. 4.6), and a wide variety of PDF plotting options.

In Fig. 49 (Right) we show a schematic representation of the xFitter code structure. The first part is the initialization, where the fit settings are specified in the steering file. This involves a number of choices, in particular selecting the fitted data sets and the theory and methodology settings such as the specific PDF parametrisation or the scheme for heavy quark structure functions. Then the PDF fit is performed, where the fit parameters are determined by means of MINUIT-based minimization including the propagation of experimental results. The final result is the LHAPDF6 grid file, together with various PDF and data/theory comparison plots.

The HERAfitter/xFitter framework has been used in many ATLAS and CMS PDF interpretation studies, discussed in Sect. 5.7.2. In addition, the HERAfitter/xFitter developer’s team has released a number of dedicated PDF studies, including:

- A QCD analysis of the legacy W and Z boson production measurements at the Tevatron [444], based on the precise W asymmetries in the electron and muon channels by the D0 collaboration, together with the HERA structure function data. This analysis demonstrated that these measurements, which are now included in most global PDF fits, provide useful information on quark flavour separation at medium and large x .
- A determination of the running charm quark mass $m_c(m_c)$ from HERA structure function data within the framework of the FONLL GM–VFNS [175]. This study demonstrated that the best–fit value of $m_c(m_c)$ was consistent to the result of a fit using the FFNS, as expected from general theoretical considerations. This value was also consistent with previous determinations of the running mass from HERA data and with the global PDG average.
- A determination of the photon PDF $x\gamma(x, Q^2)$ [52] from the measurement of high mass Drell–Yan cross sections at 8 TeV from the ATLAS collaboration. This was the first analysis where LHC data was included in a QED fit of the photon PDF directly, rather than using reweighting methods, by means of an extension of the aMCfast interface to account for the photon–initiated contributions in MadGraph5_aMC@NLO. The results of this analysis showed that the high–mass DY data indeed allowed important constraints on the photon PDF at intermediate x , although the resulting PDF uncertainties are not competitive with those from the more recent determinations of $\gamma(x, Q^2)$, discussed in Sect. 7.1.

Additional work based on xFitter include studies of PDFs with correlated uncertainties between different perturbative orders [495], non–DGLAP evolution equations [496], a study of the impact of the heavy quark matching scales in PDF fits [497], and the determination of transverse momentum dependent PDFs [72]. In Fig. 50 we show two representative PDF–related analyses performed by the xFitter Developer’s Team. First, we show the impact on the d_V PDF of the Tevatron W and Z data when added to an HERA–only fit, comparing the impact of the lepton–level measurements with that of the boson–level measurements, from the xFitter analysis of Ref. [444]. We also show the χ^2 profile of a xFitter fit based on the inclusive HERA and charm data, as a function of the running mass $m_c(m_c)$ from Ref. [175]. In this analysis charm structure functions were computed with APFEL in the FONLL–C general mass scheme. As discussed above, this analysis finds a value of the running charm mass $m_c(m_c) = 1.335 \pm 0.043$, which is consistent with the PDG average as well as with previous determinations based on HERA data, such as those in [33, 176, 177].

Concerning future developments, the xFitter code is now being rewritten from Fortran to C++, to ensure modularity and to facilitate its maintenance and the addition of novel theoretical ingredients. Several new external codes and additional features are being implemented, such as the possibility of new parametrization options like Chebyshev polynomials, the fast convolution option for hadronic cross sections as realised in APFELgrid [62], more flexible PDF parametrizations including the charm and the photon PDF, and improvements in the QED evolution interface.

5.7. PDF efforts by the LHC collaborations

As discussed in Sect. 3, the LHC experiments have provided a large number of experimental measurements with important PDF sensitivity, most of which are now part of the toolbox of global PDF fits. In

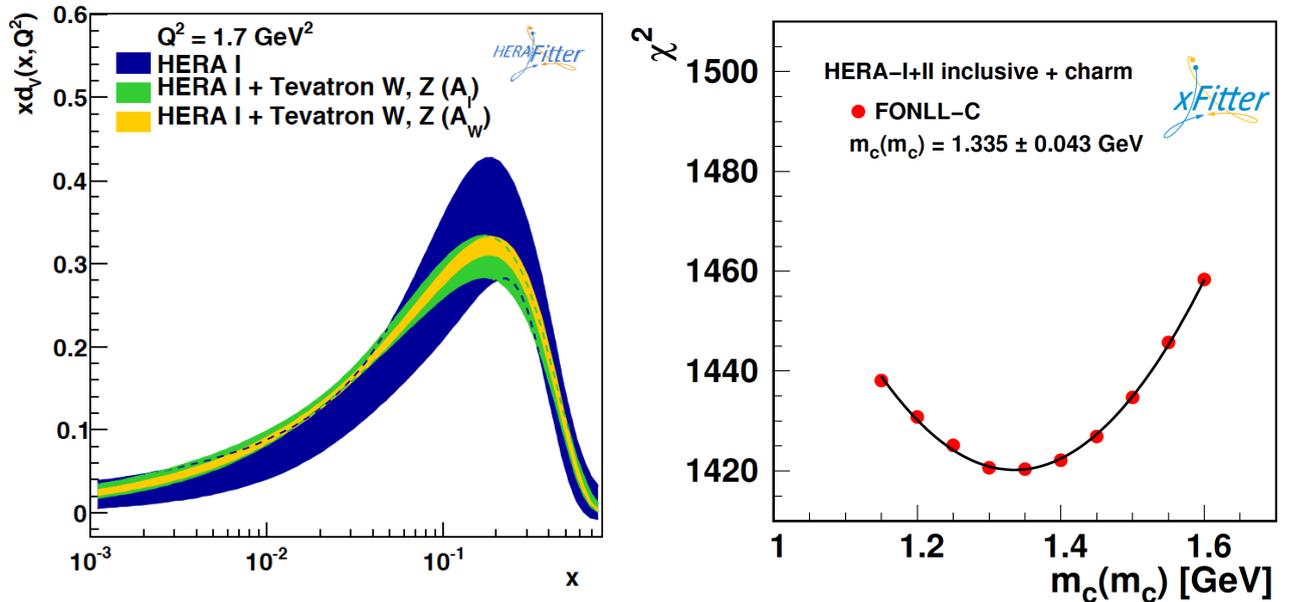


Figure 50: Two representative analyses of PDF-related studies performed by the xFitter Developer’s Team. Left plot: the impact on d_V of the Tevatron W and Z data on a HERA-only fit, comparing the impact of the lepton-level measurements with that of the boson-level measurements, from Ref. [444]. Right plot: the χ^2 profile of a fit based on the inclusive HERA and charm data, as a function of the running mass $m_c(m_c)$ from Ref. [175]. In this analysis, the charm structure functions were computed with APFEL in the FONLL-C general mass scheme.

addition to presenting the results of such measurements, the ATLAS and CMS collaborations have developed an active program of PDF interpretation studies, aimed at quantifying the constraints of their data on the proton structure.

In most cases, these studies have been performed using the HERAFitter/xFitter frameworks, described in Sect 5.6. Thus the PDFs are parameterised at $Q_0^2 = 1.9 \text{ GeV}^2$ in terms of relatively simple polynomials in x . Fits are then performed with an increasing number of free parameters introduced up to the point when no further improvement in χ^2 is observed. This generally leads to ~ 15 free parameters in the fit, with the precise number depending on the particular analysis. Experimental uncertainties are calculated using the standard ‘ $\Delta\chi^2 = 1$ ’ criteria, and additional model and parameterisation uncertainties are included, as in the HERAPDF approach. In all cases, either the HERA I DIS [129], or in later studies the I + II combination [34] are included in a baseline fit, before assessing the impact of the corresponding LHC measurements, which are then fit in addition.

However, it should be emphasised that such studies can only give an indication of the effect of the considered data on the PDFs. In particular, as they compare to a restricted DIS-only baseline set, the impact will inevitably be exaggerated in comparison to its inclusion within a global fit, and indeed if some dominant constraint from a global analysis is omitted, then the presented results may be misleading. Moreover, due to the relatively restricted parameterisation that is used, the PDF impact in x regions, for example at high x , that are not directly probed by the data can have non-negligible parameterisation dependence.

Nonetheless, such studies represent an important contribution to the PDF fitting community, not only because they give an indication the PDF impact of specific measurements, but also because they provide an internal cross-checks that the information required for PDF fits, in particular the full experimental co-

variance matrix, is ready to be used. In the following, we describe the individual efforts from ATLAS and CMS, highlighting a representative sample of their PDF studies.

5.7.1. ATLAS

A representative selection of ATLAS PDF interpretation studies is given below:

- The ATLAS measurements of W^+ , W^- and Z rapidity distributions at 7 TeV from the 2010 dataset were used in Ref. [312], to determine the strange content of the proton. The full cross-correlations between the three rapidity distributions were accounted for, and while W^+ and W^- measurements constrained the up and down quarks and antiquarks, the Z measurement constrained the strangeness. This analysis found that the strange sea was not suppressed relative to the up and down quark sea.
- The recent study [282], based on the updated W^+ , W^- and Z rapidity distributions at 7 TeV from the 2011 dataset, corresponding to a greatly improved precision in comparison to 2010. An analysis of this data was found to prefer a strange quark sea that is symmetric with respect to the light quark sea, consistent with the previous findings of [312]. The PDF uncertainties in the strangeness determination were significantly reduced in comparison to the PDFs determined from the analysis of the 2010 data. The issue of the strange content of the proton will be discussed in more detail in Sect. 6.3.
- PDF fits based on jet production measurements have also been performed by ATLAS. For instance, in [237] an analysis of the inclusive jet cross sections at $\sqrt{s} = 2.76$ TeV and $\sqrt{s} = 7$ TeV (from the 2010 run) was performed. It was shown that an improved constraint could be achieved by considering the ratio $R_{7/2.76}$ of jet cross-sections is used, given that many experimental and theoretical uncertainties cancel when taking such ratios between different centre-of-mass energies [245].

In Fig. 51 we show some representative results of PDF interpretation studies performed within the ATLAS collaboration. In the left plot, we show the results of a PDF fit quantifying the effect on the gluon of the ATLAS inclusive jet measurements at $\sqrt{s} = 2.76$ TeV and 7 TeV, in comparison to a HERA-only fit. In the right plot we show the determination of the strangeness ratio $R_s(x = 0.023, Q^2 = 1.9 \text{ GeV}^2)$ for a HERA-only fit and including the 2011 ATLAS measurements of the W^\pm and Z rapidity distributions at 7 TeV. The results of the ATLAS analysis, denoted by ATLAS-epWZ16, are compared with the predictions from various PDF fits.

5.7.2. CMS

A representative selection of PDF interpretation studies from CMS is given below:

- In [281] the 7 TeV measurement of the W charge asymmetry, as well as $W + c$ production [404], is fit at NLO, and improvements in the determination of the up and down valence quark PDFs due to the W asymmetry, and the strange quark PDFs due to the $W + c$ data, are demonstrated.
- In [311] a fit to the 8 TeV differential W boson production data [311] is performed at NNLO, again showing improvements in the determination of the up and down valence quark PDFs.
- In [16] the 7 TeV inclusive jet measurement is fit at NLO, and the significant impact of these data on the gluon PDF in particular is demonstrated. A study is also performed here using the MC method for PDF determination, allowing for a more flexible PDF parameterisation, and consistent results are found but with larger PDF uncertainties.

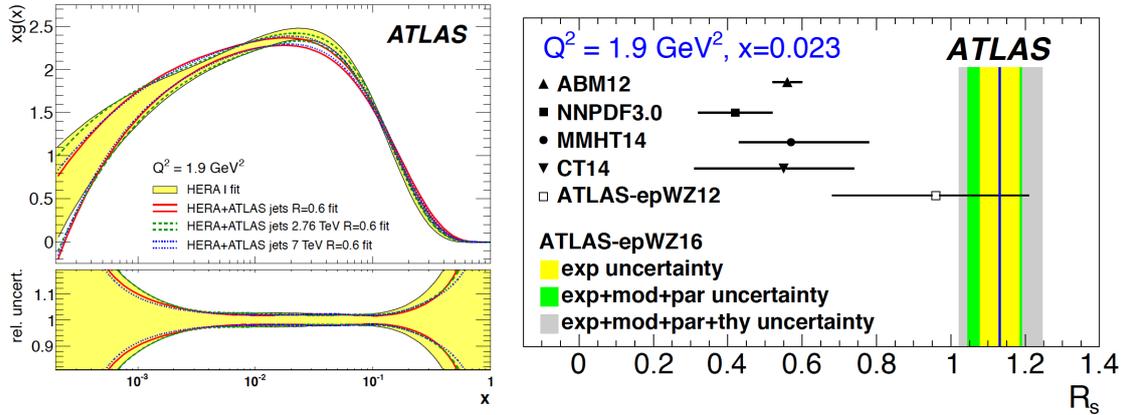


Figure 51: Two representative results of the PDF fitting efforts performed within the ATLAS collaboration. Left plot: a PDF fit quantifying the effect on the gluon from the HERA-only fit of the ATLAS inclusive jet measurements at $\sqrt{s} = 2.76$ TeV and 7 TeV, from Ref. [237]. Right plot: the determination of the strangeness ratio $R_s(x = 0.023, Q^2 = 1.9 \text{ GeV}^2)$ for a fit to HERA data and the 2011 ATLAS measurements of the W^\pm and Z rapidity distributions at 7 TeV, where the results of the xFitter analysis, denoted by ATLAS-epWZ16, are compared with the predictions from various PDF fits. Taken from [282].

- A NLO fit to the 8 TeV jet data is performed in [227], and a direct comparison to the 7 TeV case is shown, with the impact found to be very similar.
- In [369] a NLO fit to the 8 TeV double-differential top pair production data is compared to a baseline fit that includes the 8 TeV W boson production data [311]. The impact of including the data differential in different kinematic variables is assessed, and a sizeable reduction in the uncertainty on the gluon PDF in particular is found for $x > 0.01$, with the largest constraint coming from the rapidity, $y_{\bar{t}t}$, and invariant mass, $M_{\bar{t}t}$ of the top pair.
- More recent preliminary results including PDF fits to triple-differential dijet production at 8 TeV [498] and the top pair production cross section at 5.02 TeV [499] have been presented.

To illustrate these CMS PDF studies, in Fig. 52 we show the impact on the gluon PDF of the CMS W^\pm data [311], the double-differential top pair production data [369] and the inclusive jet production data [227], in all cases at 8 TeV. This is seen to lead to a sizeable reduction in the uncertainty at higher x , in a way that is consistent between the data sets in the probed x region. The $\bar{t}t$ differential data are competitive with the jet measurement. In Fig. 52 (Right) we show the impact of the CMS W^\pm boson charge asymmetry measurement at 8 TeV on the down valence distribution, in comparison to a HERA-only fit. The impact on the shape, and reduction in uncertainties, achieved by the asymmetry data is clear.

6. The proton structure

In this section we compare the results of a number of state-of-the-art PDF analysis. This comparison is organised in terms of specific PDF flavour combinations relevant for phenomenology. We begin by discussing the gluon PDF, before considering the quark flavour separation, and subsequently studying the

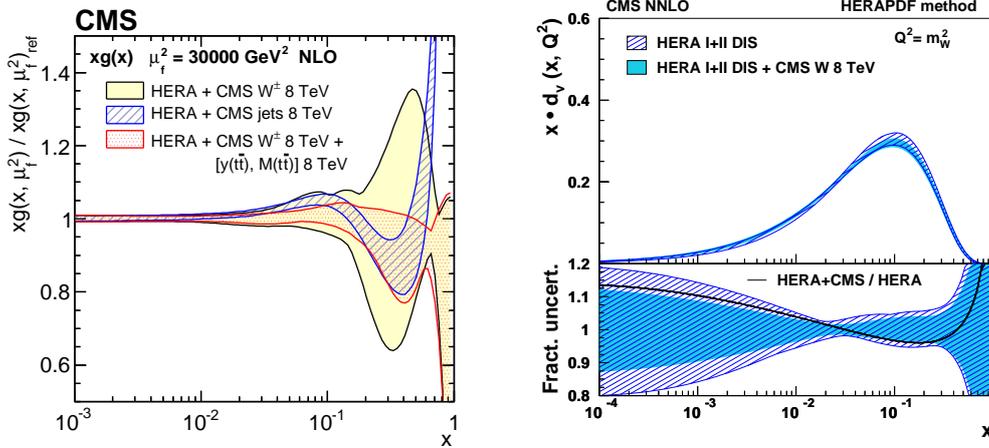


Figure 52: Left: the gluon distribution at $\mu^2 = 30000 \text{ GeV}^2$, as obtained from a PDF fit to HERA DIS data and CMS W^\pm boson charge asymmetry measurements [311], the CMS inclusive jet production cross sections [227], and the W^\pm boson charge asymmetry plus the double-differential $t\bar{t}$ cross section [369], in all cases at 8 TeV. All presented PDFs are normalized to the results from the fit using the DIS and W^\pm boson charge asymmetry measurements, and the total uncertainty band in each fit is shown. Taken from [369]. Right: the down valence distribution at $\mu^2 = M_W^2$, as obtained from a PDF fit to the HERA DIS data and CMS W^\pm boson charge asymmetry measurement at 8 TeV, with the total PDF uncertainties shown. In the lower panel the distributions are normalized to 1. Taken from [311].

strange and charm content of the proton. In the following, we compare the ABMP16, CT14, MMHT14, and NNPDF3.1 NNLO sets, all with $\alpha_S(m_Z) = 0.118$. We will present results both at low ($Q = 1.7 \text{ GeV}$) and high ($Q = 100 \text{ GeV}$) scales. While we will only show a representative selection of PDF comparisons, other results, including with PDF sets not shown here, can be straightforwardly produced using the APFEL-WEB online PDF plotting interface [500].

We emphasize that in order to ensure a consistent comparison between PDFs, we use the same value of $\alpha_S(m_Z) = 0.118$ for all the sets everywhere. This is however not the best-fit value for ABMP16, which is instead $\alpha_S(m_Z) = 0.1147$. We will briefly comment on the main differences in ABMP16 with respect to the two different values of $\alpha_S(m_Z)$. Note also that in the comparisons at low scales Q the ABMP16 cannot be included, since their set with $\alpha_S(m_Z) = 0.118$ is only available in the $n_f = 5$ scheme, and therefore can only be used for $Q > m_b$.

6.1. The gluon PDF

In Fig. 53 we show the gluon PDF, $g(x, Q^2)$, at a low scale $Q = 1.7 \text{ GeV}$ (left) and at a typical LHC scale of $Q = 100 \text{ GeV}$ (right). We find that in general there is reasonable agreement between the four considered sets within PDF uncertainties. This remains true at small x and small Q , where the PDF uncertainties become rather large, due to the limited experimental constraints. In this case, while the central value of the MMHT14 gluon becomes negative for $x \lesssim 5 \times 10^{-5}$, the CT14 result exhibits a flat behaviour, and the NNPDF3.1 gluon increases rapidly.

Perhaps the most important discrepancy between the gluon PDFs from the four groups arises in the large- x region, where the NNPDF3.1 result, and even more markedly ABMP16, is rather softer in comparison to CT14 and MMHT14. For example, at $x \approx 0.2$ the differences between the NNPDF3.1 and CT14 central values are at the $2\text{-}\sigma$ level. This may be related to the fact that different datasets are used to constrain

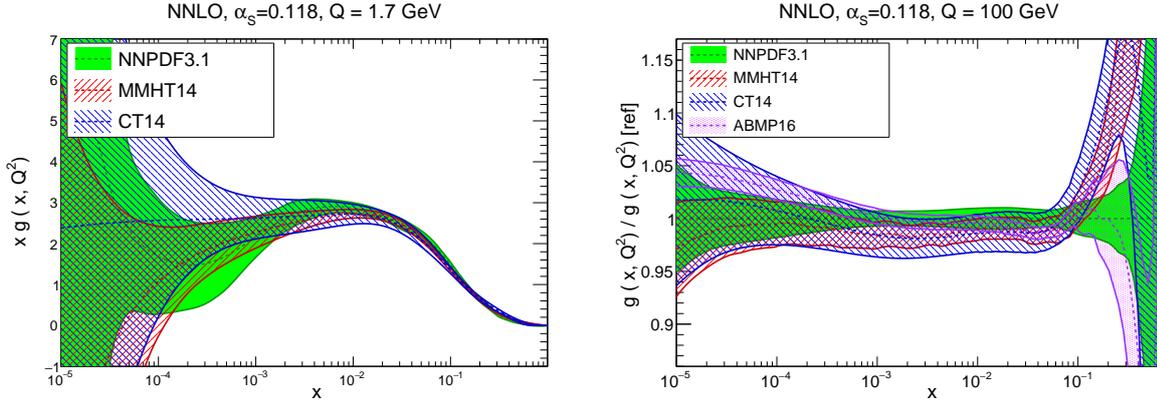


Figure 53: The gluon PDF $xg(x, Q^2)$ at $Q = 1.7$ GeV (left) and $Q = 100$ GeV (right) comparing the ABMP16, CT14, MMHT14, and NNPdf3.1 NNLO sets with $\alpha_s(m_Z) = 0.118$. In the right plot, results are normalized to the central value of NNPdf3.1.

the large- x gluon, with NNPdf3.1 in particular fitting to top-quark differential distributions, in contrast to the other sets. This has been shown to lead to a softer large- x gluon in comparison to the same fit without any $t\bar{t}$ data included [379], see also Sect. 3.7. Indeed, the CT14 and the previous NNPdf3.0 sets (which do not fit such data) are in better agreement within uncertainties. It will therefore be interesting to compare the large- x gluon PDF once other groups also include the $t\bar{t}$ differential measurements. We also note that in the large- x region PDF uncertainties are quite large, leading to significant theoretical uncertainties for the production of new BSM heavy particles, as will be discussed in Sect. 8.2.

Another interesting feature of the comparison in Fig. 53 is the fact that the gluon PDF uncertainties in ABMP16 turn out to be much smaller than those of the other three groups. As discussed in Sect. 5.4, the underlying reason for this is the use of a $\Delta\chi^2 = 1$ criterion to define the Hessian PDF uncertainties, instead of using a tolerance $\Delta\chi^2 > 1$ as done in the cases of CT14 and MMHT14. We also note that the agreement between ABMP16 and the other three sets becomes worse if the PDF set corresponding to their best-fit $\alpha_s(m_Z) = 0.1147$ value is used. To illustrate this, in Fig. 54 we show the gluon PDF in ABMP16 at $Q = 100$ GeV, comparing the results obtained with their best-fit value $\alpha_s(m_Z) = 0.1147$ with those with $\alpha_s(m_Z) = 0.118$ used to compare with the other PDF sets. When the former is used, the gluon is suppressed at large x , resulting in an increased disagreement with the results of the other three fits shown in Fig. 53 in the large- x region.

It is worth emphasising that until recently, the gluon at large- x was only constrained in the PDF fit by inclusive jet cross-sections, and to a lesser extent by DIS data via scaling violations. However, there are now at least three datasets available with which constrain the large- x gluon, namely inclusive jets, the p_T distribution of Z bosons, and top quark differential distributions. In all cases, NNLO calculations are now available. To illustrate the robustness of the resulting gluon, in Fig. 55 (Left) we show a comparison of the NNPdf3.1 NNLO global fit at $Q = 100$ GeV with the corresponding fits where the Z p_T , top quark, or inclusive jet data have been removed. We can observe that the four fits agree within PDF uncertainties, highlighting that these three families of processes have statistically consistent pulls on the large- x gluon.

Another consideration that is relevant for the determination of the large- x gluon in a PDF analysis are the settings for the theoretical calculations used for the inclusive jet cross sections. Until 2016, only the NLO calculation was available, and different groups treated jet data in different ways, either adding the

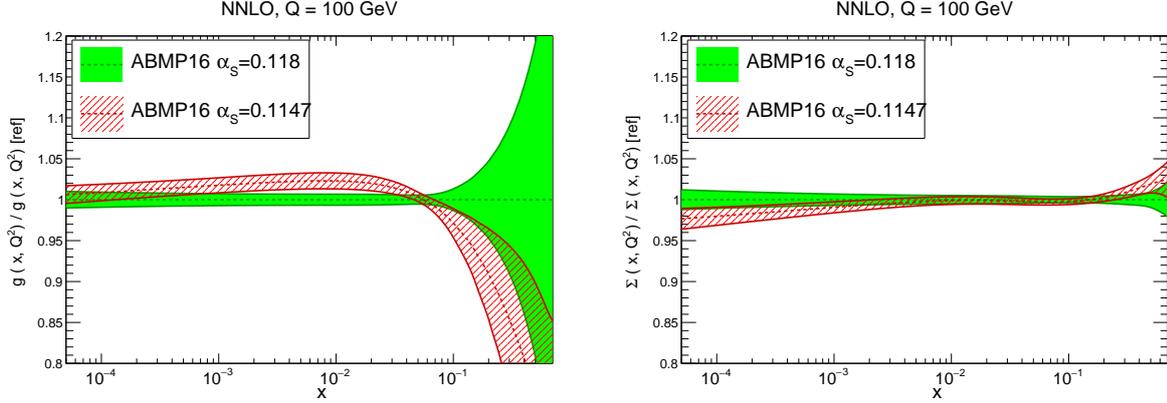


Figure 54: The gluon (left) and quark singlet (right) PDFs in ABMP16 at $Q = 100$ GeV, comparing the results obtained with their best-fit $\alpha_s(m_Z) = 0.1147$ with those with $\alpha_s(m_Z) = 0.118$ used to compare with the other PDF sets.

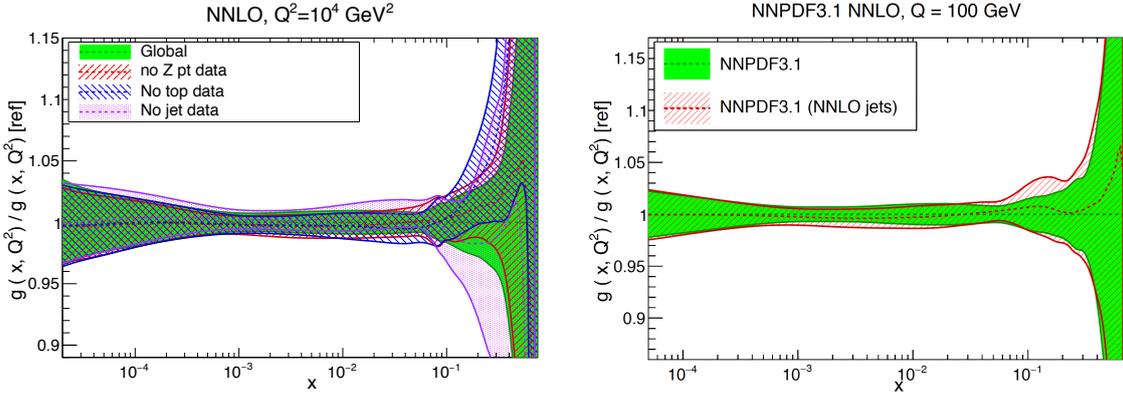


Figure 55: Left: comparison of the NNPDF3.1 NNLO global fit at $Q = 100$ GeV with the corresponding fits where the $Z p_T$, top quark, or inclusive jet data have been removed. Right: same, now comparing with the NNPDF3.1 NNLO fit where the ATLAS and CMS 7 TeV inclusive jet data have been treated using exact NNLO theory, from [169].

NLO scale errors as additional systematic uncertainties as in CT14 and NNPDF3.1, using the threshold approximation to the full NNLO result as in MMHT14, or excluding jet data altogether as advocated by ABMP16 (see Sect. 5). The availability of the complete NNLO calculation, discussed in Sect. 3.3, means that future PDF analyses will be able to fit to inclusive jet data using the exact NNLO theory.

There are some indications that for specific settings of the NLO calculation, the use of the exact NNLO theory may have moderate phenomenological impact. In particular, if the jet p_T^{jet} is adopted as the central renormalization and factorization scale, and a larger value of R is used, the NNLO/NLO K -factor is $O(5\%)$ or less [42, 229]. To illustrate this, in Fig. 55 we compare the baseline NNPDF3.1 NNLO fit, where the jet production theory is treated at NLO using p_T^{jet} as central scale, with scale variations included as an additional systematic error, to the same fit where the exact NNLO theory has been used, for the ATLAS and CMS 7 TeV data [238, 239]. We see that the resulting differences are small at the PDF level, and at the χ^2 level one finds [169] a small but non-negligible improvement once NNLO theory is used. On the other hand, for lower choices of jet radius and/or other scale choices, in particular the jet p_T^{max} per event, the NNLO/NLO

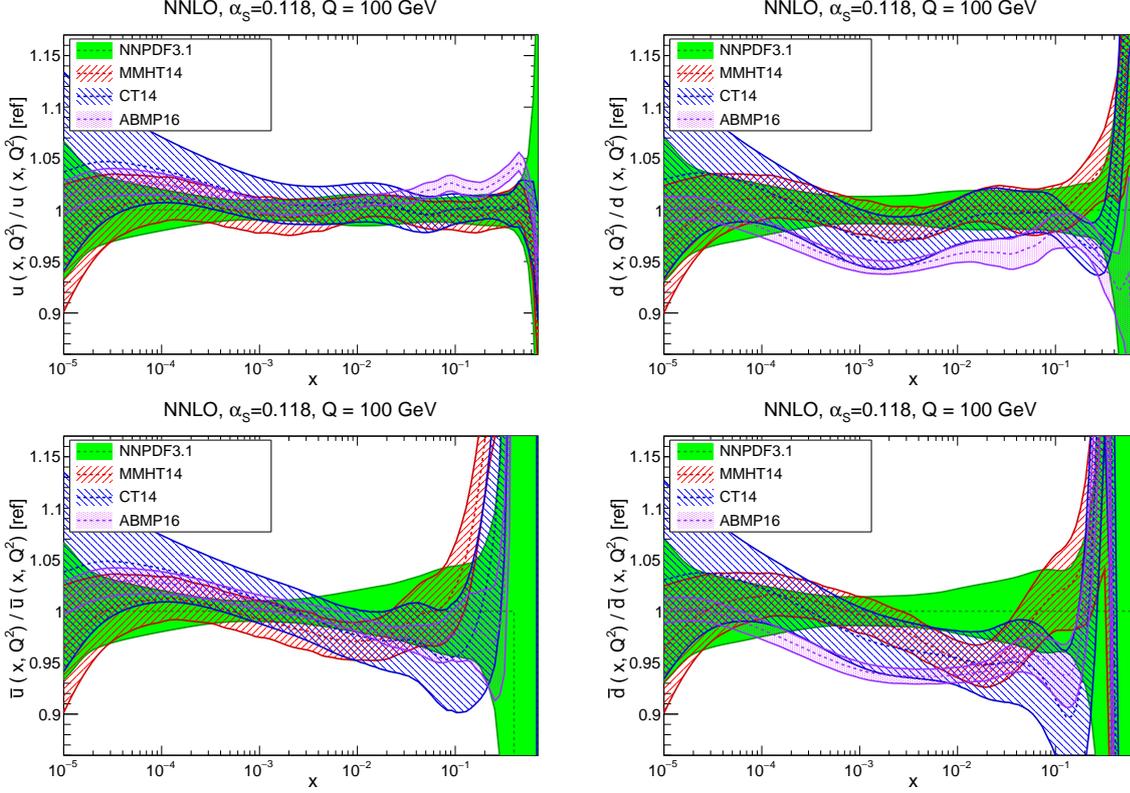


Figure 56: Same as Fig. 53 (right), now comparing the up, down, anti-up, and anti-down quark PDFs.

K-factor can be larger and the impact of the exact NNLO theory may be more significant. A full comparison of these choices within the context of a NNLO global fit will therefore clearly be desirable in the future.

6.2. Quark flavour separation

In this section we discuss the light quark content of the proton, and in particular quantify the degree with which current fits achieve quark flavour separation, i.e. the ability to separate u from d quarks, and quarks from anti-quarks. First, in Fig. 56 we show the up, down, anti-up, and anti-down quark PDFs at $Q = 100$ GeV. The up quark, $u(x, Q)$, is one of the better constrained PDFs, in particular at large x , due to fixed-target DIS data. For this PDF, we find good agreement within uncertainties in the entire range of x , with the only exception being ABMP16, which overshoots the other three sets in the large- x region.

For the down quark, $d(x, Q)$, the spread between the central values is larger, and the PDF uncertainties are also comparatively increased. Here we find good agreement between CT14, MMHT14, and NNPDF3.1 within uncertainties for the entire range of x , while ABMP16 is around 5% lower than the central NNPDF3.1 value at intermediate values of x . The PDF uncertainties are the largest at high x , with CT14, MMHT14 and ABMP16 on the other hand pointing in different directions, and the NNPDF3.1 central value lying somewhere in the middle. One possible source of difference between the groups is the treatment of deuteron nuclear corrections in the fitting of the deuteron structure functions [32], though this effect is known to be localised in the region around $x \simeq 0.1$ [481].

For the light anti-quark PDFs, \bar{u} and \bar{d} , there is reasonable agreement between the various sets within uncertainties for \bar{u} , while this agreement is only marginal for \bar{d} . In the latter case, the ABMP16 result is again

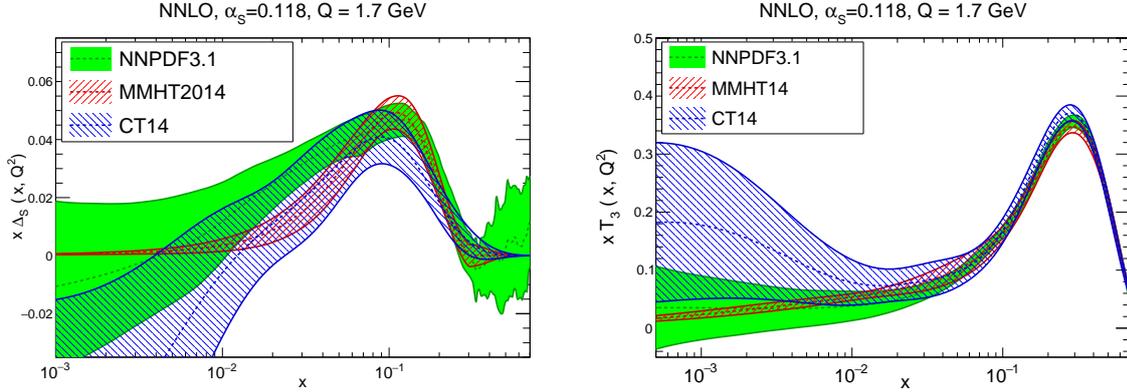


Figure 57: Same as Fig. 53 (left), now comparing the sea quark asymmetry $\Delta_S = \bar{d} - \bar{u}$ (left) and the quark isotriplet $T_3 = u + \bar{u} - d - \bar{d}$ (right plot).

around $\approx 5\%$ smaller than the NNPDF3.1 central value. As in the case of the quark PDFs, we see significant differences at large x ; in this region there are limited experimental constraints, and thus the methodological differences in each PDF fit can have a rather more marked impact. Similarly to the gluon, these large PDF uncertainties at high x have phenomenological consequences, for instance for the production of a heavy W' or Z' boson, or the pair production of a squark–anti–squark pair $\tilde{q}\tilde{q}^*$, both processes being driven by the quark–anti–quark luminosity, see Sect. 8.2.

Another useful way to compare the quark flavour separation between the various PDF groups is to plot flavour combinations that can be more directly related to physical cross sections. In Fig. 57 we compare the sea quark asymmetry $\Delta_S = \bar{d} - \bar{u}$ and the quark isotriplet $T_3 = u + \bar{u} - d - \bar{d}$ at $Q = 1.7$ GeV for CT14, MMHT14, and NNPDF3.1. The former quark flavour combination is closely related to the W asymmetries in collider Drell–Yan production (see Sect. 3.4), while the latter is directly sensitive to the difference between the proton and deuteron DIS structure functions, the non–singlet structure function $F_2^{\text{NS}} \equiv F_2^p - F_2^d$. From this comparison, we see that for the sea quark asymmetry Δ_S the general shape is similar between the three groups, although there are large differences in the estimate of the PDF uncertainties, both at small and large x , which in some cases can be traced back to the PDF parametrization assumptions. The agreement is also reasonably good both in terms of central values and of uncertainties for the quark non–singlet combination T_3 , although here again the small– x behaviour does differ among the three groups, due to the different parametrization choices.

From these comparisons, we can see that the differences in the quark flavour separation between the various groups are mostly localised in the large– x region. With this in mind, in Fig. 58 we again show the up and down quark PDFs, focusing now on the large– x region and adopting a linear scale in the x axis. From this comparison we can see that PDF uncertainties are the largest in NNPDF3.1. In terms of central values, there is reasonable agreement for the up quark, less so for the down quark. Note that as discussed in Sect. 4.1, in the NNPDF framework, the PDFs themselves are not forced to be positive (although the physical cross sections are indeed positive–definite) and therefore the down PDF can become negative at large x , although its central value remains always positive. An alternative approach that can be used to compare the behaviour of PDF sets at large x , and in doing so comparing with non–perturbative models such as the quark counting rules, is the effective exponent method discussed in Ref. [416].

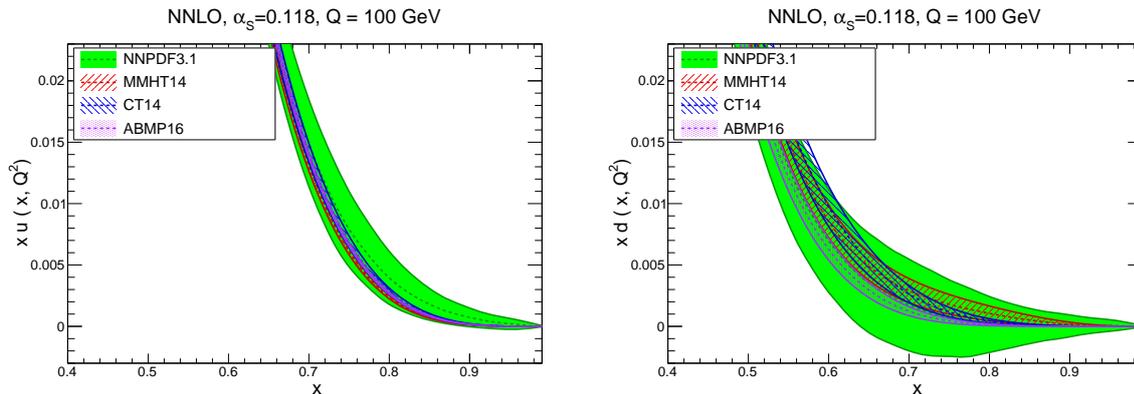


Figure 58: Same as Fig. 56, now focusing on the large- x region of the up quark (left) and down quark (right plot) PDFs.

6.3. Strangeness

The size and shape of the strange PDF has recently attracted a lot of debate, see [214, 215] for recent summaries. On the one hand, most PDF analyses find a suppressed strangeness relative to the non-strange light quark sea, a pull driven mostly by the deep-inelastic neutrino inclusive and charm production (‘dimuon’) data (see Sect. 3.2). On the other hand, high-precision collider data from the LHC have instead exhibited the opposite trend, with a recent QCD analysis from ATLAS based on the W, Z 7 TeV rapidity distributions from the 2011 dataset finding a strange sea that is in fact larger than the non-strange sea. Given the importance of strangeness for many phenomenological applications, such as the measurement of the W mass, it will be important to resolve this issue in the future. Alternative probes of $s(x, Q)$ are therefore particularly useful, such as the recent suggestion of using semi-inclusive DIS measurements with kaons in the final state [501].

In Fig. 59 we show the total strange PDF, $x s^+(x, Q^2)$, at $Q = 100$ GeV, in the same format as that of Fig. 53. The strangeness-sensitive datasets included in the four analysis are rather different, both in terms of the neutrino fixed-target data and the LHC collider data. For example, only the ABMP16 fit includes the NOMAD dimuon data [194], while only NNPDF3.1 includes the ATLAS W, Z 2011 rapidity distributions. We can see that there is reasonable agreement within uncertainties between the four groups except for ABMP16 for $x \lesssim 10^{-3}$, which has a much harder strangeness than the other fits. We also note that the differences in the size of the strange PDF uncertainty can vary by up to a factor ~ 5 , with ABMP16 exhibiting the smallest uncertainties.

A more physically transparent method to quantify the strange content of the proton is given by the ratio of the strange to the non-strange sea quark PDFs, defined as

$$R_s(x, Q^2) \equiv \frac{s(x, Q^2) + \bar{s}(x, Q^2)}{\bar{u}(x, Q^2) + \bar{d}(x, Q^2)}. \quad (129)$$

In this ratio, a symmetric strange sea would correspond to $R_s \simeq 1$, while a suppressed strangeness instead leads to $R_s \ll 1$. Note that this strangeness ratio $R_s(x, Q^2)$ depends on the specific scale Q^2 and thus when providing predictions (or measurements) for R_s it is necessary to specify the value of Q^2 which is being assumed, since results can be quite different depending on the value of Q^2 . In addition, from general considerations one expects that $R_s \rightarrow 1$ for small- x and large Q^2 , since in this region all sea quarks become very similar to each other due to perturbative evolution that becomes dominant in this region.

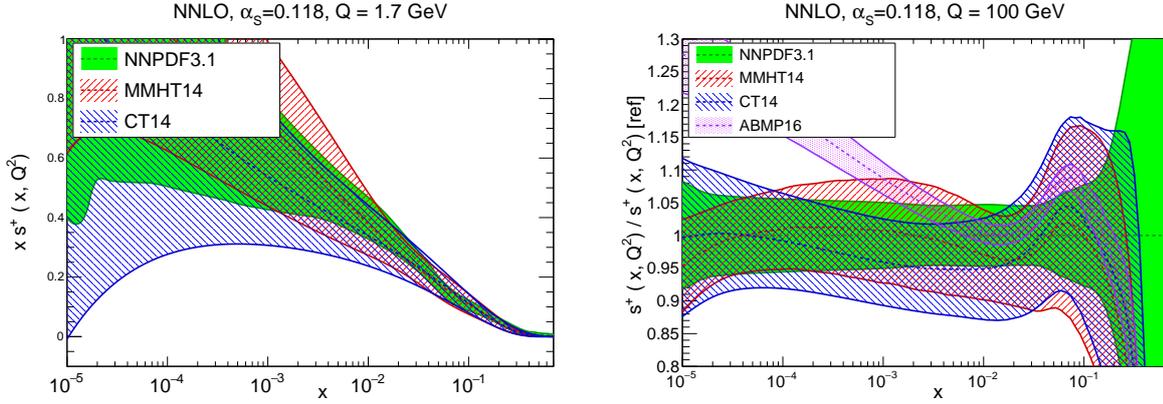


Figure 59: Same as Fig. 53 for the total strangeness $x s^+(x, Q^2)$.

Traditionally, the constraints from neutrino dimuon production have suggested a value $R_s \sim 0.5$ in most global fits. In Fig. 60 we show the ratio of strange to non-strange sea quarks $R_s(x, Q^2)$, Eq. (129), for $Q^2 = 1.9 \text{ GeV}^2$ and $x = 0.023$. We compare the results of the various global PDF fits with those of the ATLAS/xFitter study [282], which includes the recent ATLAS W, Z high precision data, as well as with those of a NNPDF3.1 fit based on the same dataset as the ATLAS study. The vertical lines indicate two possible scenarios for the strange PDFs, namely a suppression of size $R_s \approx 0.5$ and a strange sea which is symmetric with the non-strange one, $R_s \approx 1$. We can observe that the four PDF sets considered here exhibit a preference for a suppressed strangeness. On the other hand, this comparison also shows that if only the HERA and ATLASWZ11 data are considered, the NNPDF3.1 analysis yields an unsuppressed strangeness, although with PDF uncertainties rather larger than those from the xFitter analysis. This comparison demonstrates that the opposite pull between the low-energy neutrino data and the high-energy collider data is genuine, although the tension is only at the 1- or 2- σ level, as indicated by the fact that the NNPDF3.1 global and HERA+ATLASWZ11 results agree within PDF uncertainties.

One restriction of the comparison summarised in Fig. 60 is that it is restricted to a specific point $x \approx 0.023$. To bypass this limitation, in Fig. 61 we therefore show the $R_s(x, Q^2)$ ratio as a function of x both at low and at high scales. There are a number of interesting features that can be observed from this comparison. First, we can see that DGLAP evolution automatically increases the value of R_s as we go to higher scales, in particular at medium and small x ; the reason being that, as explained above, in this region all sea quarks become similar to each other. Second, we find a consistent strangeness content for the four groups in most of the range of x , although the size of the corresponding uncertainties in each case can vary by a significant amount. Another important point from this comparison is that clearly any statement about whether or not strangeness is suppressed depends on the region of x that is being considered. For instance, in the MMHT14 case for $Q = 1.9 \text{ GeV}^2$ the value of R_s changes from around 0.4 at $x \approx 0.1$ to around 0.8 for $x \approx 0.007$. So different x regions exhibit different amounts of suppression with respect to the light sea quarks, and therefore the question of the suppression (or lack thereof) of the strange PDF is a more nuanced issue than is sometimes stated. In any case, it is clear from the comparison of Fig. 61 that a symmetric strange sea in the entire range of x is not favoured by any of the four fits shown here, in particular in the region around $x \approx 0.1$ and above. We can expect future data from the LHC and elsewhere to help shed some light on this important issue.

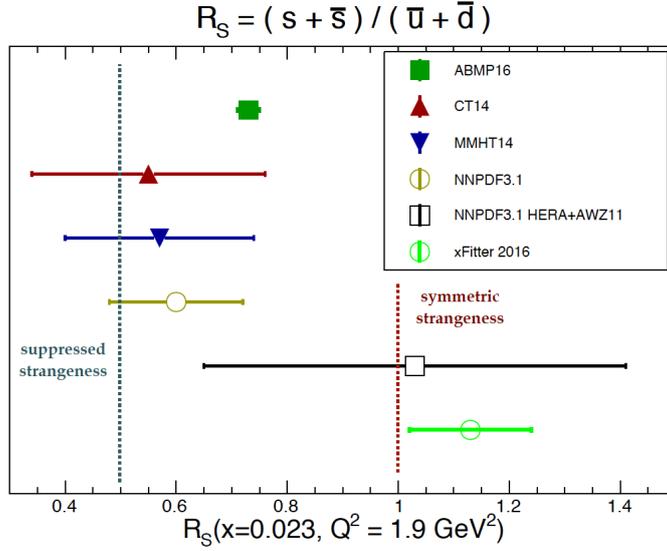


Figure 60: The ratio of strange to non–strange sea quarks $R_S(x, Q^2)$, Eq. (129) for $x = 0.023$ and $Q^2 = 1.9 \text{ GeV}^2$. We compare the results of various global PDF fits with those of the ATLAS/xFitter interpretation study as well as with those of a NNPDF3.1 fit based on the same dataset as the ATLAS study. The vertical lines indicate two possible scenarios for the strange PDFs, namely a suppression of size $R_S \approx 0.5$ and a strange sea which is symmetric with the non–strange light quark sea, $R_S \approx 1$.

6.4. The charm content of the proton

The charm content of the proton (see [502] for a recent review) is a topic that has recently received quite a lot of attention, in particular with the development of more sophisticated theoretical and methodological treatments of the charm quark PDF, moving beyond the assumption where charm is entirely generated by perturbative DGLAP evolution [55, 56]. In this respect, the NNPDF collaboration advocates that the charm PDF should be treated on an equal footing to the light quarks in global fits [56], and this leads to a reduced m_c dependence of important collider cross-sections as well as improving the agreement with high–precision data. On the other hand, the CT group [55] argues instead that a non-zero charm PDF at the initial parametrization scale Q_0 , under certain assumptions, may represent twist-4 contributions that are not suppressed by the usual Q^2 and W^2 kinematic cuts, but could also arise due to missing radiative corrections which are however not universal but rather process-dependent. From the phenomenological point of view, these and related approaches, reviewed in this section, allow for improved comparisons with non–perturbative models of the charm content of the proton [502].

As discussed in Sect. 4.1, there are two different approaches to treating the charm PDF within a PDF analysis. On the one hand, one can assume that the charm PDF is generated entirely from perturbative evolution, and thus compute the charm PDF from the gluon and light quark PDFs starting from the charm threshold $\mu_c \approx m_c$ by means of the DGLAP evolution equations. On the other hand, it is also possible to release this assumption and treat the charm PDF on an equal footing to the light quark PDFs, namely introducing a functional form for $c(x, Q_0)$ with parameters to be determined by experimental data.

Until recently, in most global fits the charm PDF was generated using perturbative evolution, and then separately, in dedicated intrinsic charm studies, variants of these global fits were performed with specific models for the charm PDF. In these studies, the parameters of the model charm PDF, typically its overall normalization, were constrained by experimental data, see for instance Refs. [55, 467, 503, 504]. An alternative

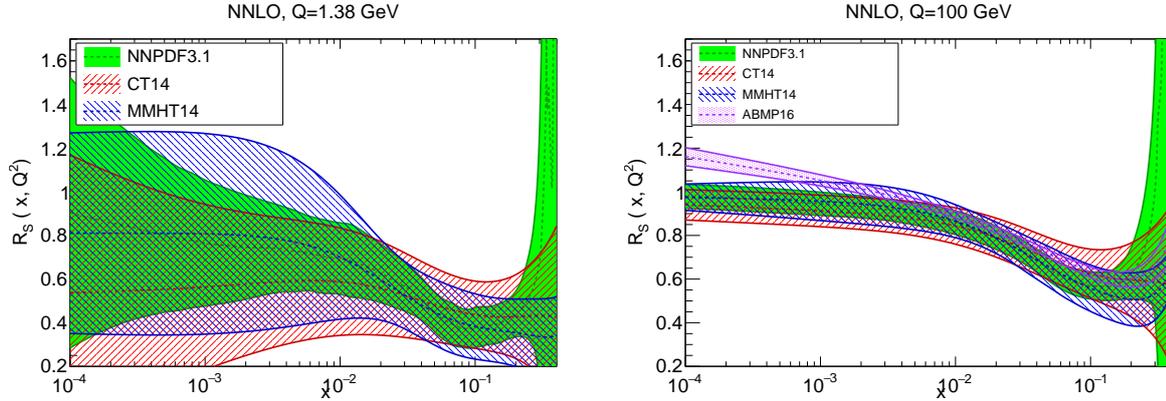


Figure 61: The ratio of strange to non–strange sea quarks $R_s(x, Q^2)$, Eq. (129), as a function of x for $Q = 1.38$ GeV (left plot) and for $Q = 100$ GeV (right plot).

approach is taken by the NNPDF3.1 global analysis, which fits the charm PDF using the same parametrization as for the light quarks. When the charm PDF is fitted, the dominant constraints arise from processes sensitive to initial–state charm, such as the charm structure functions of the EMC experiment [186], other fixed–target DIS datasets, and collider electroweak gauge boson production.

For the first approach, the CT14IC analysis [55] provides a recent and representative example. Here, the charm PDF is parametrized according to two theoretical scenarios. The first scenario assumes either the exact or the approximate BHPS model [483], which predicts a valence–like charm PDF at the input scale. In the case of the approximate solution of the model, we have

$$c(x, Q_0) = \frac{1}{2}Ax^2 \left[\frac{1}{3}(1-x)(1+10x+x^2) - 2x(1+x)\ln(1/x) \right], \quad (130)$$

while a more complicated, non–analytic expression is used for the exact BHPS solution. The other scenario explored in the CT14IC study is the SEA model, where the charm PDF is parametrized by a ‘sea–like’ function taken to be proportional to the light quark sea PDFs, namely

$$c(x, Q_0) = A \left(\bar{d}(x, Q_0) + \bar{u}(x, Q_0) \right). \quad (131)$$

In these three scenarios, the overall normalization A of the fitted charm PDF is a free parameter to be determined from the experimental data.

To illustrate the main results of this study, In Fig. 62 we plot the deviation of the χ^2 in the CT14IC fits, computed with respect to the best–fit value of the CT14 fit with perturbative charm, as a function of the charm momentum fraction $\langle x \rangle_{IC} = C(Q_0)$, where we have defined

$$C(Q^2) \equiv \int_0^1 dx x \left(c(c, Q^2) + \bar{c}(x, Q^2) \right). \quad (132)$$

Results are shown for the BHPS and SEA models, with the ‘1’ points labelling the preferred value of $\langle x \rangle_{IC}$, and those labelled with ‘2’ indicating the largest values of the charm momentum fraction allowed by the fit tolerance criteria. We observe that the BHPS model is preferred, leading to a best–fit value of $\langle x \rangle_{IC} \simeq 0.6\%$.

As mentioned above, a different approach to fitted charm is adopted by the NNPDF collaboration [56, 169]. In this case, the charm PDF is treated on an equal footing to the light quark PDFs, and therefore it is

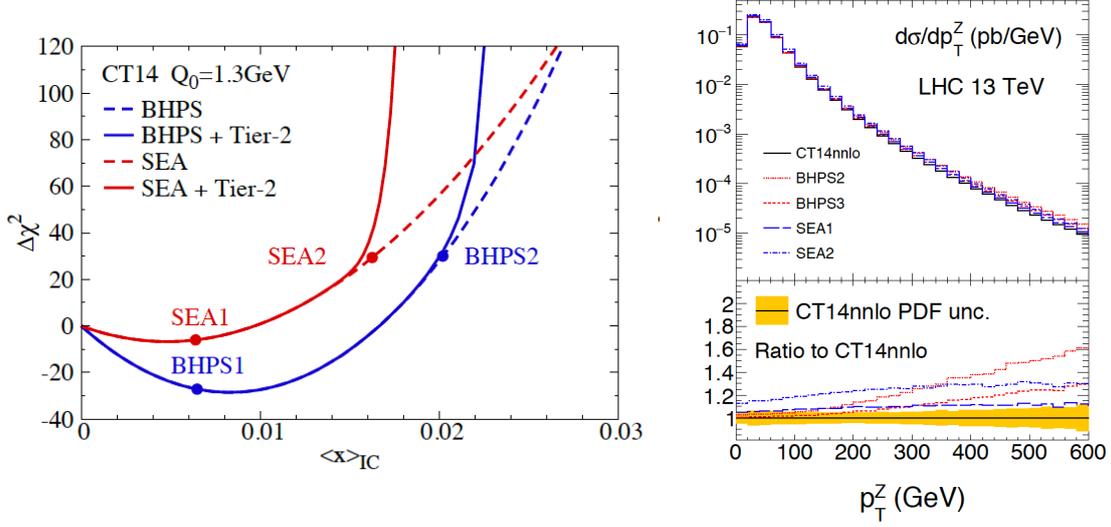


Figure 62: Left: the deviation of the χ^2 in the CT14IC fits, with respect to the best-fit value of the CT14 fit with perturbative charm, as a function of $\langle x \rangle_{IC}$. Results are shown for the BHPS and SEA models, with the “1” points labeling the preferred value of $\langle x \rangle_{IC}$, with those labelled with “2” indicate the largest values of the charm momentum fraction allowed by the fit tolerance criteria. Right: the transverse momentum distribution of Z bosons in the $pp \rightarrow Z + c$ process at 13 TeV, comparing the CT14 NNLO result with various of the CT14 IC models, as a function of p_T^Z .

parametrized with a 37-parameter artificial neural network with 2-5-3-1 architecture,

$$c^+(x, Q_0) = c(x, Q_0) + \bar{c}(x, Q_0) = x^{a_{c^+}} (1-x)^{b_{c^+}} \text{NN}_{c^+}(x), \quad (133)$$

where a_{c^+}, b_{c^+} are the corresponding preprocessing exponents, whose range is determined by an iterative procedure. The charm asymmetry, on the other hand, is assumed to vanish $c^-(x, Q_0) = 0$. The charm PDF is then determined at the input evolution scale $Q_0 = 1.64$ GeV from the experimental data, finding that the recent LHC high-precision electroweak gauge boson production measurements provide the best constraints [169]. DIS structure functions are treated with the FONLL general-mass VFN scheme, modified to account for initial-state massive contributions [163, 164].

One of the consequences of this model-independent approach is that it improves the stability of the fitted PDFs with respect to the value of the charm mass m_c , since its variations can be re-absorbed into the fitted charm boundary condition. To illustrate this point, in Fig. 64 we show the dependence of the quark-anti-quark PDF luminosity at the LHC 13 TeV in the NNPDF3IC fits with the value of the charm mass m_c used in the fit. We find that even for a relatively large variation of $\delta m_c = \pm 0.14$ GeV, the $q\bar{q}$ luminosity is very stable in most of the M_X range. A similar stability is found at the level of relevant LHC cross-sections such as the differential distributions for Higgs and Z production, as discussed in [56].

The amount of charm present inside the proton is most usefully quantified by its momentum fraction, defined in Eq. (132). In the case of perturbative charm, by construction we have $C(Q^2 < \mu_c^2) = 0$, while if there is a non-perturbative charm component in the proton in general we have $C(Q^2) \neq 0$ at all values of the scale Q^2 . In Fig. 63 we show the momentum fraction carried by charm quarks both at a low scale $Q = 1.51$ GeV and at a high scale $Q = M_Z$, comparing the results from NNPDF3.0, based on perturbative charm, with those from NNPDF3.1, based on fitted charm, with and without the inclusion of the EMC charm data, as well as with the BHPS and SEA scenarios of the CT14IC fits. The uncertainty bands of the CT14IC

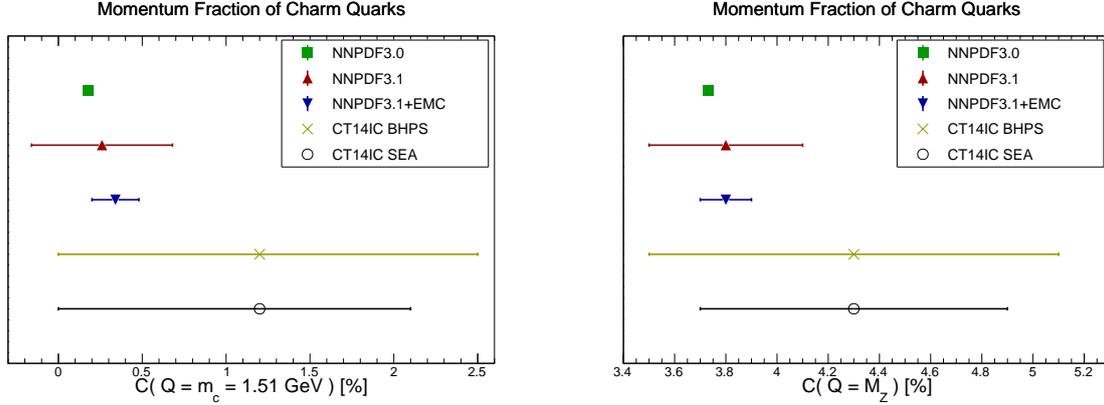


Figure 63: The momentum fraction carried by charm quarks, $C(Q)$ Eq. (132), both at a low scale $Q = 1.51$ GeV (left) and at a high scale $Q = M_Z$ (right plot). We compare NNPDF3.0 (perturbative charm) with NNPDF3.1 (based on fitted charm) with and without the inclusion of the EMC charm data, as well as with the BHPS and SEA scenarios of the CT14IC fits. See text for more details.

results indicate the range between no intrinsic charm and the maximum amount of IC allowed within the CT14 tolerance, with the central value corresponding to the best fit. Note that the standard CT14IC fits do not include the EMC data, with the rationale that the experimental systematics are not well documented. It was found in alternative fits that neither of the BHPS or SEA models can describe the EMC data well. This is also true for the fit with purely perturbative charm. Thus, within the CT framework, the EMC charm structure function data is found to provide no definitive discrimination between the purely perturbative and intrinsic charm models.

The comparisons of Fig. 63 highlight, first of all, that when the charm PDF is generated perturbatively its uncertainties are very small, but that this is not necessarily the case once charm is fitted. Reassuringly, once charm is fitted (NNPDF3.1), the results with perturbative charm (NNPDF3.0) are consistent within uncertainties. The NNPDF3.1 analysis also finds that adding the EMC charm data helps in reducing the PDF uncertainties on $\langle x \rangle_{IC}$ by around a factor 3, but that even without it one can achieve a quite competitive determination of $c(x, Q)$, due to the precision collider electroweak data. The CT14IC results are consistent within PDF errors with the NNPDF3.1IC, although the highest values within the CT14IC are disfavoured by the NNPDF fit with latest LHC data. The rapid growth of $C(Q)$ from $Q = 1.51$ GeV to $Q = M_Z$, driven by the perturbative component, is also apparent from this comparison.

In Fig. 64 we show a comparison of the fitted charm PDF at $Q = 1.65$ GeV between the NNPDF3IC set and the different model scenarios considered in the CT14IC analysis. We see that NNPDF3IC prefers a valence-like shape, along the lines of the BHPS model, though uncertainties are still large. The CT14IC BHPS results tend to have a maximum at slightly lower values of x ; note also that they develop a perturbative tail since the plot is performed at a value $Q > Q_0$. The CT14IC SEA models predict that the enhancement of the charm PDF is localised at medium and small x , while in the valence region the fitted charm agrees with the perturbative calculation.

In order to better disentangle the consequences of parametrizing the charm PDF in the NNPDF case, in Fig. 65 we show the comparison between the charm PDF $x c(x, Q^2)$ in NNPDF3.1 NNLO with fitted and with perturbative charm, both at $Q = 1.65$ GeV and at $Q = 100$ GeV, in the latter case normalized to the fitted charm results. One finds that both approaches are consistent within errors, with differences at the two-sigma level at most in the region around $x \simeq 0.05$, where the fitted charm PDF undershoots the perturbative

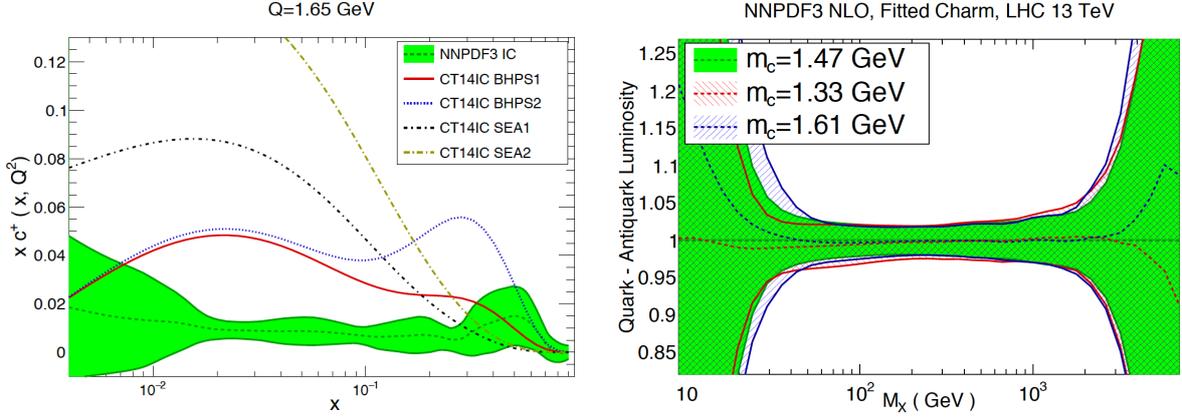


Figure 64: Left: comparison of the fitted charm PDF at $Q = 1.65$ GeV between the NNPDF3IC set and the different models of the CT14IC analysis. Right: the dependence of the quark–anti–quark PDF luminosity at the LHC 13 TeV in the NNPDF3IC fits with the value of the charm mass m_c used in the fit.

ansatz, both at low and at high scales. We emphasize that the latter behaviour (which is absent for instance in the CT14IC fits) is unrelated to the EMC charm structure function measurements, which are not included in the NNPDF3.1 baseline dataset and thus in the fits shown in Fig. 65

If the charm content of the nucleon is indeed different from the one generated purely from the perturbative evolution, there are a number of phenomenological consequences that could be studied at the LHC. To begin with, it would modify the kinematic distributions of the Z bosons in the $pp \rightarrow Z + c$ process [505–507], leading to an enhancement of the cross section which grows with the value of p_T^Z . To illustrate this point, in Fig. 62 we show the transverse momentum distribution of Z bosons in the $pp \rightarrow Z + c$ process at $\sqrt{s} = 13$ TeV, comparing the CT14 NNLO result with the CT14 IC models, as a function of p_T^Z . Depending on the specific model considered, enhancements of up to 50% are predicted, those this may be diluted by fragmentation contributions.

A closely related process is photon production in association with charm mesons [508, 509], which is however theoretically less clean as it is affected by the poorly understood parton–to–photon fragmentation component. Another important process where non–perturbative charm could play a role is open D meson production [510, 511], in particular at large p_T and at forward rapidities, which enhance the sensitivity to the large– x region. By exploiting the information from these various LHC processes using data from the present and future runs, we can hope to shed more light on the charm content of the proton in the future.

7. Electroweak corrections and the photon PDF

In this section we explore a topic that has recently received a lot of attention in PDF fits, namely the role of QED and weak corrections, and in particular photon-initiated processes. We first discuss the role of QED corrections and the photon PDF, before reviewing the case of pure weak corrections to the hard scattering matrix elements arising from virtual massive weak boson exchange.

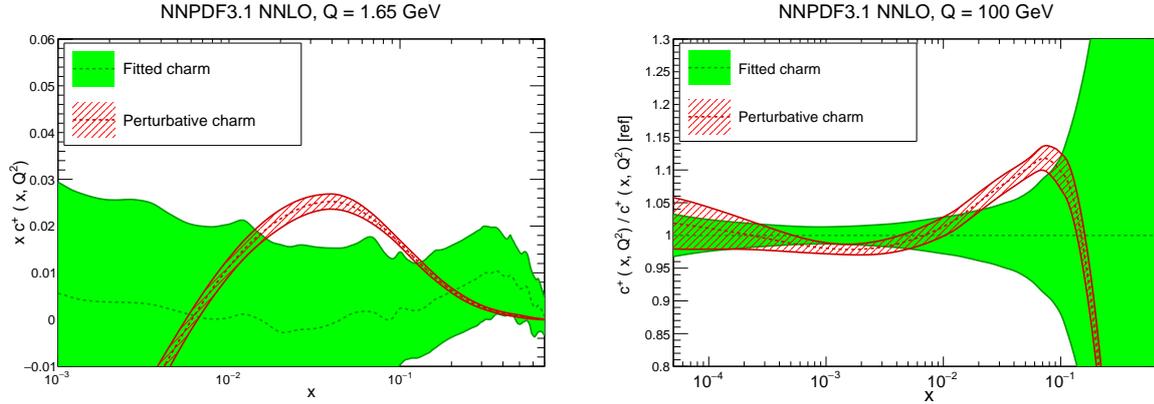


Figure 65: Comparison between the charm PDF $x c(x, Q^2)$ in the NNPDF3.1 NNLO analysis with fitted and with perturbative charm, both at $Q = 1.65$ GeV (left) and at $Q = 100$ GeV (right plot), in the latter case normalized to the fitted charm results. Note that these fits do not include the EMC charm structure function data.

7.1. Photon-induced processes

It has been over a decade since the calculation of the splitting functions at NNLO in α_s [122, 123] provided the necessary tools to be able to carry out NNLO PDF fits. Moreover, we have seen in Sect. 3 that for the majority of PDF-sensitive observables, the perturbative calculation is available at this NNLO order. Given that the data from the LHC are available at increasing precision, to below the percent level, NNLO PDF fits are essential to match this unprecedented precision and have naturally become the standard. However, a simple counting of powers of α_s indicates that

$$\alpha_s^2(M_Z) \sim \frac{1}{70}, \quad \alpha_{\text{EM}}(M_Z^2) \sim \frac{1}{130}. \quad (134)$$

That is, we may roughly expect the NNLO QCD and NLO EW corrections to be of the same order of magnitude. While such an argument clearly neglects the non-trivial differences in the structure of the QCD and EW corrections, this nonetheless serves as a warning that we must at least consider the impact of going to NLO EW if we are to claim percent-level precision to LHC cross sections.

A specific type of EW correction of particular relevance to PDF studies is the contribution from photon-initiated processes, such as those shown schematically in Fig. 66. As this involves a photon in the initial state, this requires the introduction of a PDF for the photon in the proton⁹. This is included in complete analogy to the QCD partons, and moreover as it involves a massless boson in the initial state, higher-order QED $q \rightarrow q\gamma$ and $\gamma \rightarrow q\bar{q}$ splitting will generate collinear singularities that must be absorbed into the corresponding PDFs. In other words, this will produce QED corrections to the DGLAP evolution of the PDFs. Another important type of EW corrections relevant for PDF fits, namely those associated to virtual massive weak boson exchange, are discussed in Sect. 7.2.

⁹For brevity, we will refer to this throughout as the photon PDF, but this should not be confused with the partonic content of the photon itself, which often receives a similar label, see e.g. [512].

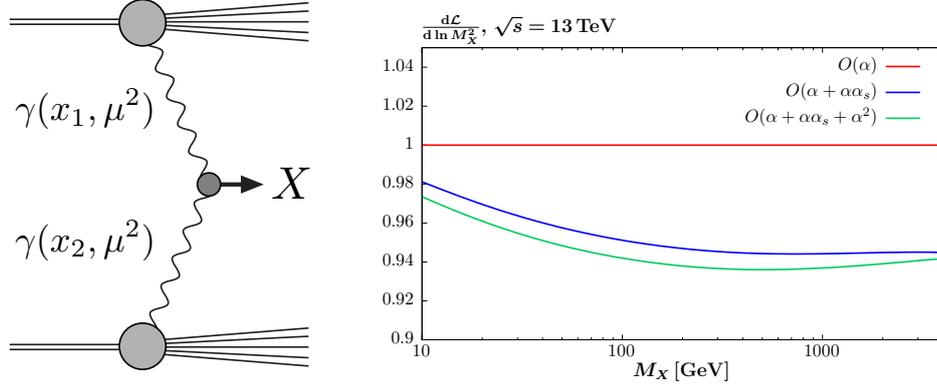


Figure 66: Left: schematic diagram for the photon–initiated production of a system X , and the corresponding photon PDFs. Right: the $\gamma\gamma$ luminosity as a function of the invariant mass, M_X , of the produced final state. The ratio of results with $O(\alpha_s)$ and $O(\alpha^2)$ to the leading $O(\alpha)$ DGLAP evolution shown. Calculated using the approach described in [476].

QED corrections to DGLAP evolution

The introduction of the photon PDF requires the following straightforward extension of the DGLAP evolution equations,

$$\begin{aligned}
 Q^2 \frac{\partial}{\partial Q^2} g(x, Q^2) &= \sum_{q, \bar{q}, g} P_{ga}(x, \alpha_s(Q^2)) \otimes f_a(x, Q^2) + P_{g\gamma}(x, \alpha_s(Q^2)) \otimes \gamma(x, Q^2), \\
 Q^2 \frac{\partial}{\partial Q^2} q(x, Q^2) &= \sum_{q, \bar{q}, g} P_{qa}(x, \alpha_s(Q^2)) \otimes f_a(x, Q^2) + P_{q\gamma}(x, \alpha_s(Q^2)) \otimes \gamma(x, Q^2), \\
 Q^2 \frac{\partial}{\partial Q^2} \gamma(x, Q^2) &= P_{\gamma\gamma} \otimes \gamma(x, Q^2) + \sum_{q, \bar{q}, g} P_{\gamma a}(x, \alpha_s(Q^2)) \otimes f_a(x, Q^2).
 \end{aligned} \tag{135}$$

The splitting functions can then be expanded in powers of both the QCD and QED coupling

$$P_{ij} = \sum_{m,n} \left(\frac{\alpha_S}{2\pi} \right)^m \left(\frac{\alpha}{2\pi} \right)^n P_{ij}^{(m,n)}. \tag{136}$$

The lowest order QED splitting function $P_{\gamma q}^{(0,1)}$ is due to the same type of Feynman diagram as in the LO QCD case for $P_{gq}^{(1,0)}$, with the gluon simply replaced by a photon, and similarly for P_{qq} and $P_{q\gamma}$. Thus these are trivially related by suitable adjustments of the colour factors and inclusion of the electric charges e_q of the quark, with

$$P_{qq}^{(0,1)} = \frac{e_q^2}{C_F} P_{qq}^{(1,0)}, \quad P_{q\gamma}^{(0,1)} = \frac{e_q^2}{T_R} P_{qg}^{(1,0)}, \quad P_{\gamma q}^{(0,1)} = \frac{e_q^2}{C_F} P_{gq}^{(1,0)}, \quad P_{\gamma\gamma}^{(0,1)} = -\frac{2}{3} \sum_f e_f^2 \delta(1-x), \tag{137}$$

where for the $P_{\gamma\gamma}$ case only the Abelian contribution is present and the sum is over all fermions in the loop, that is quarks and leptons¹⁰. The calculation of the $O(\alpha_s \alpha)$ ($m, n = 1$) terms, where the $P_{g\gamma}$ and $P_{\gamma g}$

¹⁰To be consistent, and in particular to preserve momentum fully, this requires the introduction of lepton PDFs in the proton. However as discussed in [513] the contribution from these is generally of limited phenomenological relevance, and can be safely neglected. Note in any case that including leptons in the running of the QED coupling $\alpha(Q)$ is still required.

splittings enter for the first time, is given in [514], while the $O(\alpha^2)$ ($m = 0, n = 2$) terms are given in [515]. Publicly available implementations of the DGLAP evolution including these QED corrections are provided by the APFEL [58] and QEDEVOL [516] packages.

The impact of the $O(\alpha\alpha_S)$ and $O(\alpha^2)$ corrections on the $\gamma\gamma$ luminosity is shown in Fig. 66. The $O(\alpha\alpha_S)$ corrections have a fairly small but clearly non-negligible impact on the luminosity, giving up to a $\sim 5\%$ negative correction. As we would expect, the $O(\alpha^2)$ corrections are significantly smaller, but can reach the percent level. The results of Fig. 66 have been computed using the MMHT framework [476], which is closely based on the LUXqed formalism described below.

As discussed in detail in [517], the $P_{\gamma\gamma}$ self-energy contribution to the DGLAP evolution of the photon PDF is intimately connected to the choice of renormalization scale for the initial-state photon coupling to the hard process. It is well known in QED that for on-shell external photons the coupling receives no renormalization, and is completely determined to be $\alpha(0)$. However, the $P_{\gamma\gamma}$ term breaks this simple picture, and we should instead use $\alpha(\mu_F)$ in the calculation. Physically, the photon substructure is being resolved by the introduction of a photon PDF and the contribution from $\gamma \rightarrow q\bar{q}$ splittings in the evolution, such that a purely on-shell renormalization scheme is no longer appropriate. This has been confirmed at NLO EW order in [518], where it is shown that using the on-shell scheme leads to uncanceled fermion-mass singularities in the hard cross section.

The photon PDF

The first attempts at describing the photon PDF can be divided into two distinct categories, either being phenomenological approaches that model the photon PDF, as in the MRST2004QED [49] and more recent CT14QED [50] sets, or the data-driven approach of the NNPDF2.3/3.0QED [53, 54] sets. The first attempt to include the photon in a PDF set was provided by MRST2004QED. Here, a simple model for the photon PDF at input scale Q_0 due to one-photon emission off the valence quarks was assumed. In other words, the quark valence distributions are frozen at Q_0 and the LO QED DGLAP evolution for the photon is integrated between the light quark mass m_q and Q_0 , so that

$$\gamma(x, Q_0^2) = \frac{\alpha}{2\pi} \left[\frac{4}{9} \log\left(\frac{Q_0^2}{m_u^2}\right) u(x, Q_0) + \frac{1}{9} \log\left(\frac{Q_0^2}{m_d^2}\right) d(x, Q_0) \right] \otimes \frac{1 + (1-x)^2}{x}. \quad (138)$$

The CT14QED set generalised this approach, allowing additional freedom in the normalization of the photon, which was fit to ZEUS data [468] on isolated photon production¹¹. Thus, within such approaches the photon PDF is completely *predicted* within the specific model, up to any freedom in the model parameters, such as the choice of quark masses for MRST2004QED or the overall normalization for CT14QED.

On the other hand, the NNPDF2.3QED [53] set (subsequently updated to NNPDF3.0QED [54]), freely parameterises the photon PDF at input. In other words, the photon is treated on exactly the same footing as the QCD partons. This is then extracted from a fit (or more precisely, a Bayesian reweighting) to DIS and LHC W, Z data; in the former case no explicit photon-initiated contribution was included, and so the constraint came purely from the effect on the PDF evolution. However, in general the contribution of photon-initiated process are small, leading to significant uncertainties on the extracted photon PDF.

More recently, there has been a great deal of progress in our understanding of the photon PDF. One crucial point that was missed in the above approaches is the long range nature of QED. That is, the proton is itself an electrically charged object which can coherently emit a photon, with the proton remaining intact

¹¹In fact, as we will discuss below, this has subsequently been supplemented with the elastic component to give the inclusive set CT14QEDinc.

afterwards. The possibility for such elastic photon emission is of course very well established. Elastic ep scattering is an extremely well measured process, providing for example the first measurement of the proton charge radius [73, 74] in the 1950s, with further precise measurements of this process [519] continuing to this day. Theoretically, the well known equivalent photon approximation (EPA) [520] provides a precise foundation for describing the elastic scattering process in terms of a flux of coherently emitted photons from the proton.

The connection of this fact to the photon PDF was discussed in [521] and more recently in [522, 523]. Following the equivalent photon approximation, it is straightforward to show that elastic photon emission leads to a contribution to the photon PDF at a scale $Q_0 \sim 1$ GeV given by

$$\gamma_{\text{el}}(x, Q_0^2) = \frac{1}{x} \frac{\alpha}{\pi} \int_{\frac{x^2 m_p^2}{1-x}}^{Q_0^2} \frac{dQ^2}{Q^2} \left[\left(1 - x - \frac{x^2 m_p^2}{Q^2} \right) F_E(Q^2) + \frac{x^2}{2} F_M(Q^2) \right], \quad (139)$$

where F_E and F_M are the elastic and magnetic form factors of proton, which are related to the electric and magnetic charge distributions in the proton. These are steeply falling functions of Q^2 that are probed very precisely in a range of elastic ep scattering experiments, see e.g. [519].

To demonstrate the connection of this elastic component to the inclusive photon PDF, if we for simplicity omit the small backreaction that the photon has on the quark and gluon PDFs via the evolution equations, then we can solve Eq. (135) to get [524]

$$\gamma(x, \mu^2) = \frac{1}{\alpha(\mu^2)} \left(\alpha(Q_0^2) \gamma(x, Q_0^2) + \int_{Q_0^2}^{\mu^2} \frac{dQ^2}{Q^2} \alpha(Q^2) \sum_{q, \bar{q}, g} P_{\gamma a}(x, \alpha_s(Q^2)) \otimes f_a(x, Q^2) \right), \quad (140)$$

$$\equiv \gamma_{\text{input}}(x, \mu^2) + \gamma_{\text{evol}}(x, \mu^2). \quad (141)$$

Thus the photon is given separately in terms of an input at low scale Q_0 and an evolution component due to the usual DGLAP $q \rightarrow q\gamma$ emission for $Q^2 > Q_0$. The latter is completely determined in terms of the quark and gluon PDFs, leaving the input photon at Q_0 , which is dominantly due to elastic emission. Thus this already provides quite a strong constraint on the photon PDF; as we will see below, the impact in comparison to the model-independent NNPDF approach can be dramatic.

However, even for relatively low photon virtualities, $Q^2 < Q_0^2$, the emission may also be inelastic, such that the proton no longer remains intact afterwards. In other words we have

$$\gamma(x, Q_0^2) = \gamma_{\text{el}}(x, Q_0^2) + \gamma_{\text{inel}}(x, Q_0^2), \quad (142)$$

In [522, 523] fairly simple phenomenological models for this inelastic component, given by suitable generalizations of (138), were taken, while the CT14QED set allowed an additionally free normalization to be fitted to ZEUS data on isolated photon production, as described above.

Given that the elastic component is directly determined from the form factors F_E and F_M , that are themselves measured from elastic ep scattering, it is natural to ask whether the inelastic component can be similarly determined. In other words, rather than relying on a phenomenological model, can γ_{inel} instead be calculated directly from the form factors for inelastic ep scattering, that is, from the proton structure functions? In the analysis of [47] it was shown that this is indeed the case, with the corresponding LUXqed

PDF set made publicly available. In particular, they find that the photon PDF can be expressed as¹²

$$\begin{aligned}
x\gamma(x, \mu^2) = & \frac{1}{2\pi\alpha(\mu^2)} \int_x^1 \frac{dz}{z} \left\{ \int_{\frac{x^2 m_p^2}{1-z}}^{\frac{\mu^2}{1-z}} \frac{dQ^2}{Q^2} \alpha^2(Q^2) \left[\left[z P_{\gamma q}(z) + \frac{2x^2 m_p^2}{Q^2} \right] F_2\left(\frac{x}{z}, Q^2\right) - z^2 F_L\left(\frac{x}{z}, Q^2\right) \right] \right. \\
& \left. - \alpha^2(\mu^2) z^2 F_2\left(\frac{x}{z}, \mu^2\right) \right\}. \tag{143}
\end{aligned}$$

Thus the photon PDF is in fact a completely derived quantity, which can be written purely in terms of the proton structure functions, which are known quite precisely from the experimental point of view. More recently, the detailed study of [48] has demonstrated how this expression may be derived in a process-independent way by using the operator definition of the photon PDF, as well as generalising this expression to the case of the polarized and transverse momentum dependent PDFs. This approach is also shown to provide a simple derivation of the known $O(\alpha_s)$ and $O(\alpha^2)$ splitting functions, $P_{\gamma i}$.

While the connection of Eq. (143) to the considerations above is not immediately clear, some similarity in the overall form with Eq. (140) is apparent. Indeed, by substituting for $F_{2,L}$ in terms of the quark and gluon PDFs, at high Q^2 this readily reduces to γ_{evol} in (140); indeed this is how the LUXqed photon PDF is calculated in this region. In addition, using the known expression for the elastic contributions to $F_{2,L}$ reproduces Eq (139) when combined with Eq. (140); this elastic contribution is also included. By using fits to the experimentally determined inelastic structure functions at low Q^2 , including in the resonance region, it is shown in [47] that the remaining inelastic contribution, and therefore the photon PDF in its entirety, is very precisely determined. In Fig. 67 we show an overview of the various contributions to the photon PDF $\gamma(x, Q^2)$ in the LUXqed approach as a function of x at $Q = 100$ GeV. We see that at small x it is dominated by the PDF contribution, while at large x the elastic contribution accounts for up to half of the size of $\gamma(x, Q)$.

It is worth emphasizing that the expression Eq. (143) does not rely on any explicit separation between an input and evolution component to the photon as in (140), and corresponds to the exact result for the photon within the quoted accuracy of [47, 48], valid to all orders in QED and QCD, and including non-perturbative corrections. Indeed, applying standard DGLAP above the starting scale Q_0 the power-like $\sim m_p^2/Q^2$ correction would be missed, while for $Q^2 > Q_0^2$ the contribution from the elastic component would be omitted and the inelastic resonance component, which contributes at higher x in this region, would not be correctly modelled.

However, from the point of view of a global PDF set it may be preferable to use Eq. (143) in a form that can be more directly implemented within the standard fitting framework. That is, applying this approach after suitable modification to calculate the input photon, which can then be included as part of the default input parton set for any future fits and studies. This strategy has been implemented to construct the recent NNPDF3.1luxQED global analysis [527], see also [476] for initial results in the context of the MMHT framework. An alternative iterative approach is proposed in [48].

To illustrate the differences and similarities between these various determinations of $\gamma(x, Q)$, in Fig. 67 we show the comparison of the photon PDFs from CT14qed_inc, MRST2004, NNPDF2.3/3.0 and LUXqed, normalized to the central value of the latter. It is clear from this comparison that the theoretical uncertainties associated with the LUXqed determination are much smaller than in any other of the previous approaches. Interestingly, the LUXqed PDF is in good agreement within uncertainties with the model-independent

¹²Following the publication of [47] it was discovered that this expression had been derived in the earlier papers of [525, 526], but without the correct limits on the Q^2 integral.

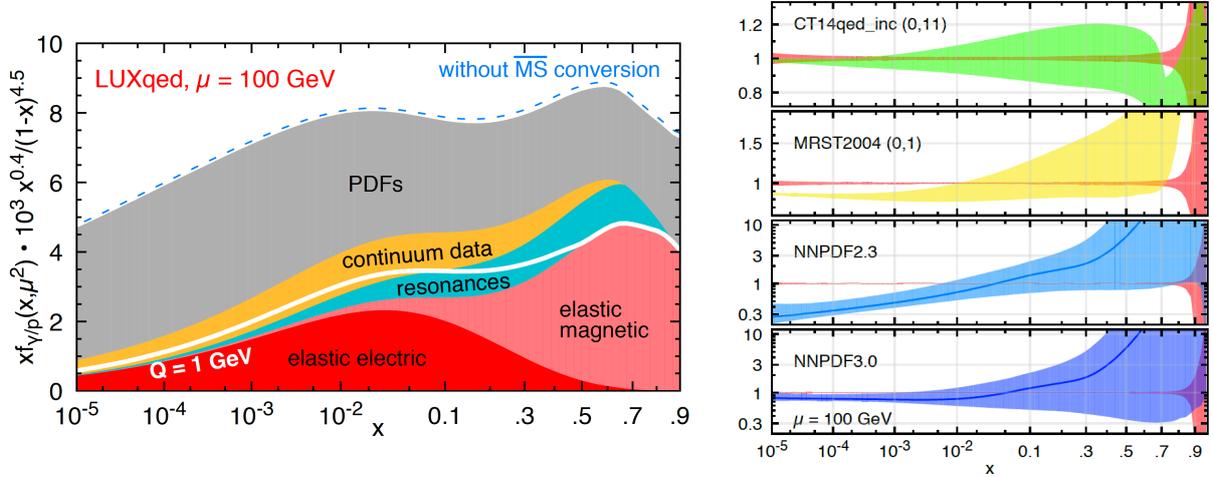


Figure 67: Left: overview of the various contributions to the photon PDF $\gamma(x, Q^2)$ in the LUXqed approach as a function of x at $Q = 100$ GeV. Right: comparison of the photon PDFs from CT14qed_inc, MRST2004, NNPDF2.3/3.0 and LUXqed, normalized to the central value of the latter.

NNPDF3.0QED fit in the entire range of x , though the PDF uncertainties are much larger in the latter case in particular at large x .

Phenomenology

In Fig. 68 (left) we show the NNPDF3.0QED $\gamma\gamma$ luminosity at $\sqrt{s} = 13$ TeV, including the 68% C.L. error bands. A very large PDF uncertainty is evident, in particular at higher system mass M_X . As discussed above, the input component in Eq. (141) is poorly determined within this approach, due to limited constraints placed by the available experimental data. It is therefore unsurprising that the PDF errors should be most significant at higher mass, as here the relative contribution from this input component is larger, due the reduced phase space for PDF evolution. In addition, the central value of the luminosity is seen to lie towards the upper end of the uncertainty band. As discussed in [28, 523], this exhibits a much gentler decrease with M_X in comparison to the QCD parton luminosities. However, also plotted is the LUXqed result, and the difference is dramatic. The central value lies towards the lower end of the NNPDF band at higher mass, with a PDF uncertainty that is smaller than the line width of the plot. We also show the result of [523], labelled HKR16, which demonstrates a similar trend. Thus, simply applying basic physical conditions on the form of the photon PDF, and including the dominant coherent input Eq. (139) gives a qualitatively similar result.

Taking a closer look, in Fig. 68 (right) we show the ratio of the HKR16, CT14QED and xFitter_HMDYep results to the LUXqed luminosity. The xFitter_HMDYep set is extracted in [52] by applying a similar agnostic methodology to NNPDF, but including the more constraining ATLAS high mass Drell–Yan data [528] in the fit; this therefore represents the most up to date set within such an approach. Again, the LUXqed uncertainty band is barely visible on the curve, varying from 1 – 2% over the considered mass interval. The CT14QEDinc prediction, which includes an elastic component, is consistent, but with larger $\sim 20 - 40\%$ uncertainties, due to the more limited constraints placed by the ZEUS isolated photon production data on the inelastic input¹³. The HKR prediction lies somewhat below the LUXqed result, outside of the quoted

¹³In addition, the ZEUS data are selected by requiring that at least one track associated with the proton side is reconstructed

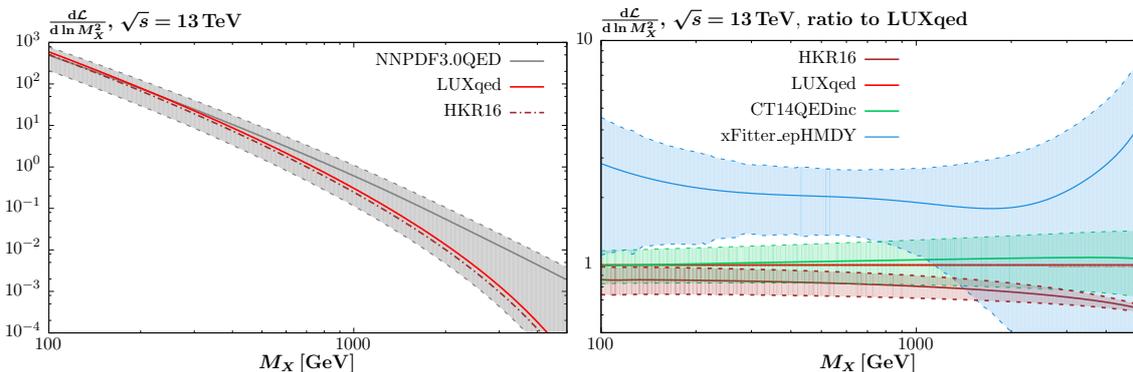


Figure 68: $\gamma\gamma$ luminosities at $\sqrt{s} = 13$ TeV. (Left) Absolute values for the HKR16, NNPDF3.0QED and LUXqed sets. (Right) Ratios of the CT14QED, HKR16 and xFitter_HMDYep sets to the LUXqed prediction. 68% uncertainty bands are shown, with the exception of the HKR16 set, where the error is due to model variation in the inelastic input (lower edge corresponds to elastic only).

model variation band, in particular at larger M_X . This is due in large part to the lack of any explicit resonant contribution in the inelastic input, which becomes more important at higher x and hence M_X . Finally, the xFitter_HMDYep set is seen to have a sizeable uncertainty band (albeit smaller than the NNPDF3.0 set [52]), and interestingly appears to lie somewhat above the LUXqed result. From this it is clear that any attempt to extract the photon PDF within such an approach will almost certainly not be competitive. More generally, we can see that the LUXqed set exhibits by far the smallest PDF uncertainties.

Prior to these most recent developments, a range of phenomenological studies pointed out similar trends in the NNPDFQED predictions for the photon-initiated contributions to lepton and W pair [28, 529, 530] and $t\bar{t}$ [377] production. At high system invariant mass these could be significant, and even dominant over the standard channels, with a large PDF uncertainty. From Fig. 68 the reason for this is clear, being driven by the large PDF uncertainty in the $\gamma\gamma$ luminosity at high mass, and the relatively gentle decrease with mass in the central value. However, from the discussion above we know that using the NNPDF set will dramatically overestimate the uncertainty on the photon-initiated contribution. In Fig. 69 we show the lepton pair production cross section at high mass, at the $\sqrt{s} = 13$ TeV LHC and a $\sqrt{s} = 100$ TeV FCC. We can see that indeed at the LHC, the NNPDF prediction for the photon-initiated contribution can even be larger than the standard Drell–Yan contribution. However, the up-to-date LUXqed prediction exhibits no such behaviour. The prediction is under good theoretical control, and gives a small, though not negligible, contribution.

Thus, by considering the physics that generates the photon PDF, and recognising the dominance of the elastic emission process, we already achieve a significant reduction in PDF uncertainty in comparison to the model-independent approach, even when accounting for the most sensitive data available in the latter case. Moreover, the additional information provided by Eq. (143) in combination with the high precision data on the inelastic (and elastic) proton structure functions give extremely tight constraints on the photon PDF, resulting in a $\sim 1\%$ level PDF uncertainty. It is worth emphasising that while consistency tests are of course to be encouraged, this is not the result of a particular theoretical model, to be treated on the same footing as older PDF sets. The LUXqed set is a fundamentally experimental determination of the photon PDF; it is

within the detector acceptance. While this will remove the elastic component entirely, which will produce no extra tracks, it is also possible for the proton dissociation products in the inelastic case to fall outside the acceptance and therefore not pass such a cut. Thus while CT14QED use these data to extract the total inelastic component, at least part of this will also be removed by this cut.

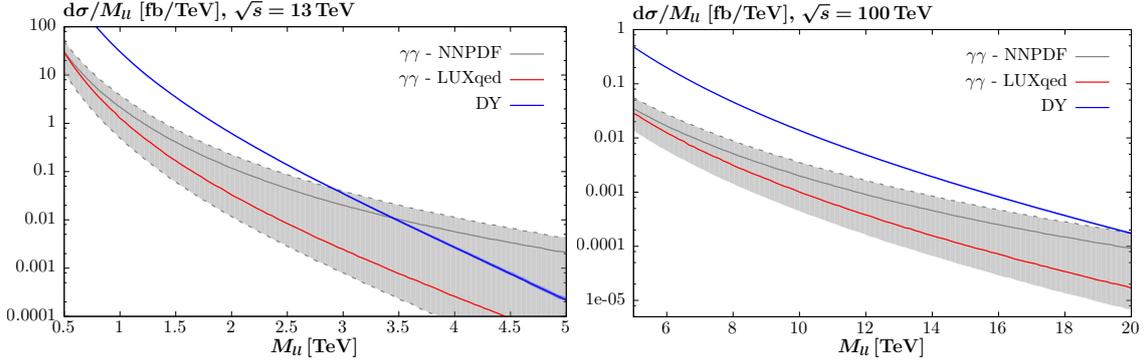


Figure 69: The differential lepton pair production cross sections at $\sqrt{s} = 13$ TeV and 100 TeV with respect to the invariant mass of the pair M_{ll} , for lepton $|\eta| < 2.5$ and $p_{\perp} > 20$ GeV. The photon-initiated contributions predicted using the LUXqed and NNPDF3.0QED sets, including the 68% C.L. uncertainty bands. The NLO Drell–Yan cross section, calculated with MCFM [306], is also shown.

simply that by doing this directly in terms of the measured structure functions the tightest constraints can be achieved. Such information must be included in any future photon PDF set, and we have therefore moved beyond the era of large photon PDF uncertainties. Indeed, the photon PDF, which used to be the poorest known of all the proton PDFs, now has the smallest uncertainty.

7.2. Electroweak corrections

In addition to the QED photon-initiated corrections discussed above, it can also be important to include other EW contributions, in particular those arising from virtual EW bosons, in a PDF fit. These corrections are most important at larger invariant masses of the produced system, $Q \gg M_W$, where virtual EW contributions receive logarithmic enhancements, see Ref. [531] for a review. In particular, the virtual exchange of soft or collinear weak bosons leads to Sudakov logarithms of the form $\alpha_W \ln^2 Q^2/M_W^2$, where $\alpha_W = \alpha/\sin^2 \theta_W$, which can lead to large (negative) corrections for large values of Q . Given that many of the LHC datasets that enter into the global PDF are sensitive to the TeV region, from high-mass Drell–Yan production to the large p_T tail of Z production and inclusive jets and dijets, the inclusion of such EW corrections is in general required to achieve the best possible description of experimental data in this region.

The state-of-the-art for EW corrections is NLO, that is $\mathcal{O}(\alpha_W)$ relative to the Born-level process, including in addition in some cases mixed terms of the form $\mathcal{O}(\alpha_W \alpha_s)$ and related terms. These corrections are available for most of the hadron collider processes that enter a typical global fit, including inclusive jet and dijet production [255], inclusive electroweak gauge boson production at high p_T [44, 339] and high invariant mass $M_{ll(\nu)}$ [296, 297] and differential top quark pair production [376, 377]. Most of these calculations are implemented in publicly available programs. For instance, EW corrections to inclusive gauge boson production are available in programs such as FEWZ [296, 297] and HORACE [302]. The latest version of MCFM [532] also includes the calculation of weak corrections to Drell–Yan, top quark pair, and dijet production at hadron colliders. Recently, there has also been progress in the automation of the calculation of these corrections, both in the framework of MadGraph5_aMC@NLO [533] and of Sherpa/OpenLoops [534].

In Fig. 70 we show two representative examples of NLO EW corrections for processes relevant for PDF determinations, computed with MCFM at $\sqrt{s} = 13$ TeV [532]. In the left plot, we show the percentage NLO EW correction for high-mass dilepton production as a function of M_{ll} . The ZGRAD calculation [301] is also shown. We see that these corrections are negligible for $M_{ll} \lesssim 500$ GeV, but that they can become

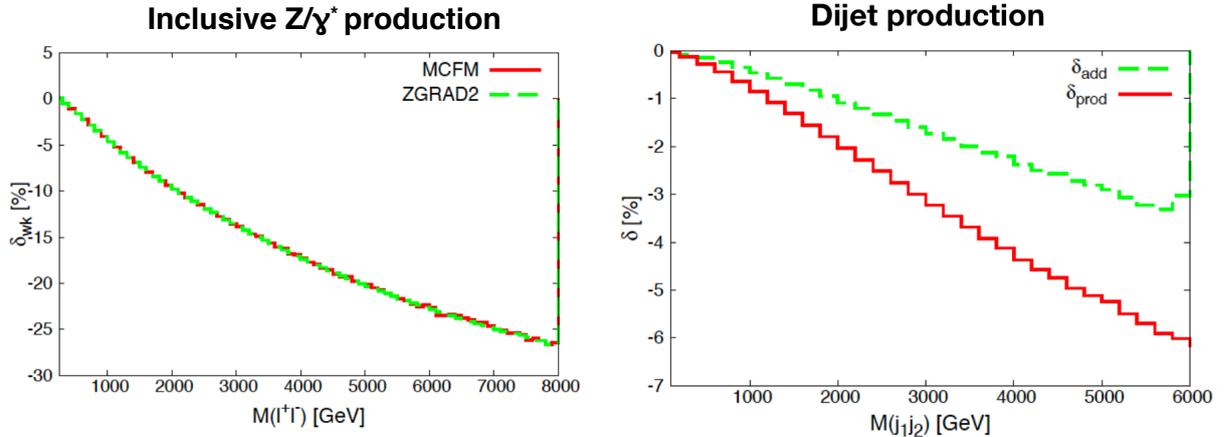


Figure 70: Two representative examples of NLO EW corrections for processes relevant for PDF determinations, computed with MCFM at $\sqrt{s} = 13$ TeV [532]. In the left plot, we show the percentage NLO EW correction for high-mass dilepton production as a function of M_{ll} , comparing also with the corresponding ZGRAD calculation. In the right plot, we show the same quantity, now for dijet production as a function of the invariant mass of the dijet M_{jj} . The two curves correspond to two possible ways to combine NLO QCD and EW corrections, known as additive (δ_{add}) and multiplicative (δ_{prod}).

significant as we increase M_{ll} , reaching $\delta_{\text{wk}} \sim -20\%$ at 5 TeV. In the right plot, we show the same quantity, now for dijet production as a function of the invariant mass of the dijet M_{jj} . The two curves correspond to two possible ways to combine NLO QCD and EW corrections, known as additive (δ_{add}) and multiplicative (δ_{prod}). Here the corrections are more moderate (since the Born is a pure QCD process) but they can still reach a few percent in the region accessible at the LHC. The results of Fig. 70 illustrate how a careful inclusion of NLO EW corrections is important for the description of the LHC data in the TeV region used for PDF determinations.

8. Implications for LHC phenomenology

In this section we present an overview of some of the most representative implications of PDFs and their uncertainties for LHC phenomenology. First we discuss the role of PDFs for the predictions of the Higgs boson production cross-sections at the LHC. Then we assess the role of PDF uncertainties in searches for new heavy resonances predicted by various BSM scenarios. Finally, we highlight the importance of PDFs for the precision measurements of SM parameters such as the W mass and the strong coupling constant.

8.1. Higgs production cross-sections

In the Standard Model, once the Higgs mass m_h is measured, all other parameters of the Higgs sector, such as the strength of its coupling to fermions and vector bosons and its branching ratios, are uniquely determined [10, 535, 536]. On the other hand, deviations of these Higgs couplings with respect to the SM predictions are generically expected in BSM scenarios [537]. Therefore, the precision measurements of the Higgs couplings represents a unique opportunity for BSM searches, since any deviation with respect to the tightly fixed SM predictions would represent a smoking gun for New Physics. Crucially, realising this program requires not only high precision experimental measurements of Higgs boson production and its

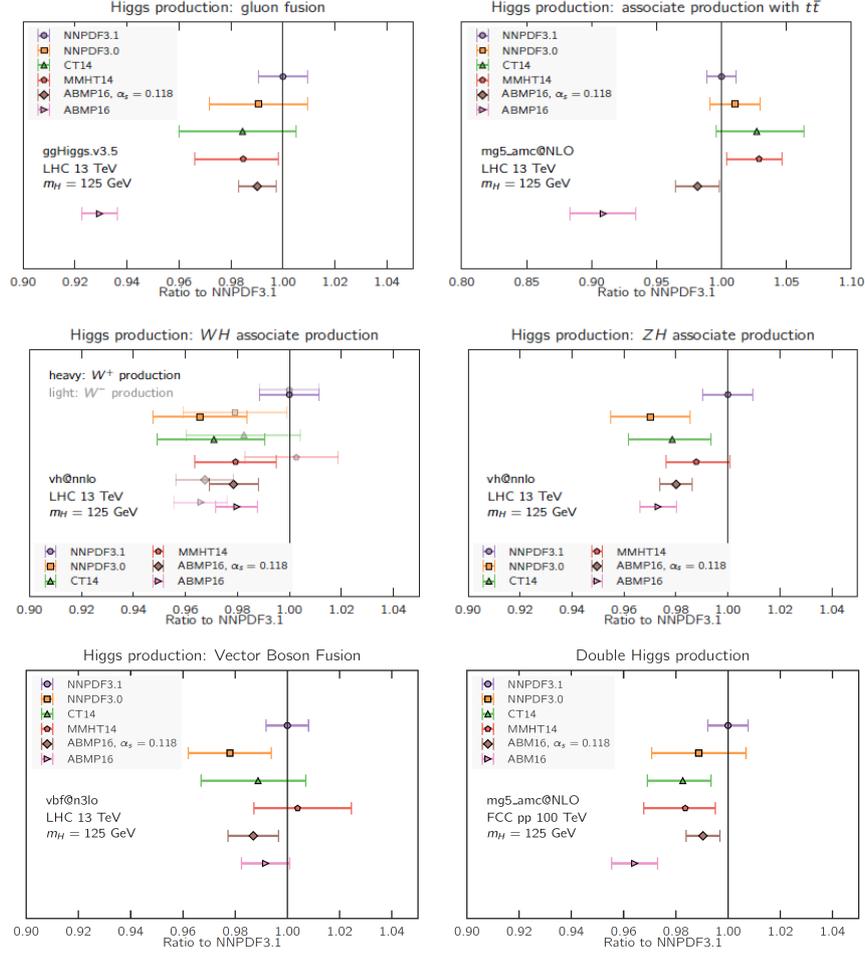


Figure 71: The PDF dependence of the most important Higgs production inclusive cross sections at the LHC 13 TeV. The results are normalized to the central value of the NNPDF3.1 prediction, and only PDF uncertainties are shown. In the case of ABMP16, we show results for the nominal α_s value as well as for the common $\alpha_s(m_Z) = 0.118$ value. See [169] for more details of the theoretical calculations.

decay in various channels, but also the calculation of the SM cross sections and decay rates with matching theoretical precision. Moreover, PDFs are one of the largest sources of theoretical uncertainty affecting the predictions for Higgs boson production [10].

We therefore present a comparison of inclusive Higgs production cross sections at 13 TeV with the latest releases of all PDF groups. The settings of this comparison, and the codes used for each process, are described in Ref. [169]. Specifically, we show the dominant Higgs boson production modes at hadron colliders: gluon fusion, associated production with a $t\bar{t}$ pair, VH associated production, and Higgs production in vector–boson fusion. In addition, we also show the results for double Higgs production in the dominant gluon–fusion channel, which should become accessible at the High Luminosity LHC (HL–LHC) [538–542]. Results are provided for NNPDF3.0 and 3.1, CT14, MMHT14 and for the ABMP16 NNLO sets for $\alpha_s(m_Z) = 0.118$, while in the latter case we also show the result corresponding to their best–fit value of $\alpha_s(m_Z) = 0.1149$. The theoretical settings for each cross section calculation are based on state–of–the–art

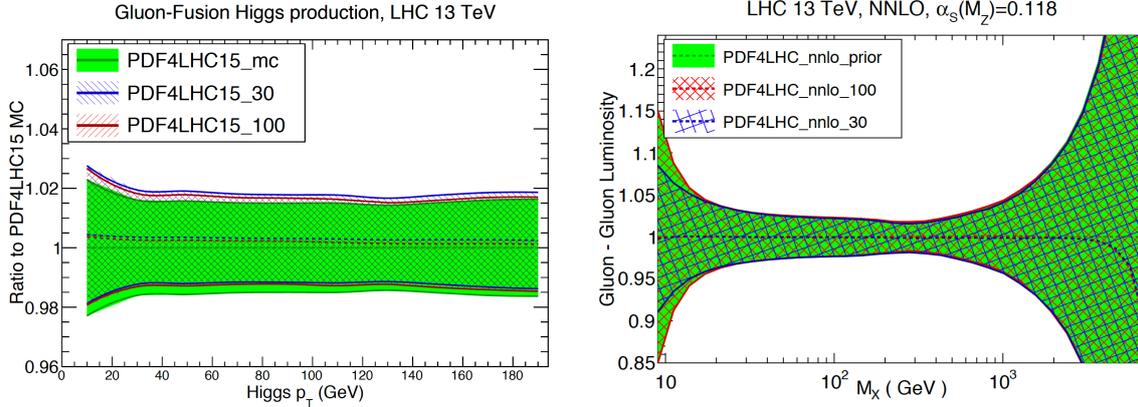


Figure 72: Left: the PDF uncertainties in the p_T^h distribution of Higgs bosons produced in the gluon–fusion mode at the LHC 13 TeV for $0 \leq p_T^h \leq 200$ GeV, computed using the PDF4LHC15 NNLO sets. Right: the gluon–gluon PDF luminosity with the same set now extending up to higher values of the invariant mass of the final state M_X .

matrix element calculations, for instance the gluon–fusion and VBF results are computed at N3LO using the `ggHiggs` [543] and `vbf@n3lo` [544] codes respectively. We only show the PDF uncertainties, while other sources of theoretical errors affecting these cross sections are listed in e.g. the latest Higgs Cross Section Working Group report [10]. We note that the uncertainty associated to the input value of $\alpha_s(m_Z)$ can be comparable to the PDF uncertainties in some channels.

There are a number of noteworthy features in the comparison of Fig. 71. First, it shows that in general there is good agreement between the three global fits, NNPDF3.1, CT14 and MMHT14 for all the Higgs boson production modes. The comparison between NNPDF3.1 and its predecessor NNPDF3.0 highlights good agreement for the gluon initiated channels, with a reduction of the PDF uncertainties in the former case, while for quark–initiated cross sections such as VH and VBF there is an upper shift by around one sigma. Another significant feature of this comparison is that the recent ABMP16 set is also in reasonable agreement with the other groups, provided that the same common value of the strong coupling constant $\alpha_s(m_Z) = 0.118$ is used. On the other hand, if their best–fit value $\alpha_s(m_Z) = 0.1149$ is used in the calculation, significant differences in the predicted cross sections arise for the gluon–initiated channels. In particular, the ABMP16 prediction is around 5–7% (10–12%) lower than the other predictions for the gluon–fusion ($t\bar{t}$ associated production) cross section.

It is also worth mentioning here that PDF uncertainties are relevant not only for the extraction of Higgs couplings from inclusive cross sections, but also for the differential measurements [545, 546] that will become available due to the large statistics that will be accumulated by the end of Run II as well as for the HL–LHC. To illustrate this point, in Fig. 72 we show the PDF uncertainties in the Higgs boson p_T^h distribution in the gluon–fusion mode at the LHC 13 TeV for $0 \leq p_T^h \leq 200$ GeV, computed using the PDF4LHC15 NNLO sets. In this case we find that PDF uncertainties are at around the $\sim 2\%$ level. However, these uncertainties will increase as the LHC becomes sensitive to higher p_T values: as shown in the right plot of Fig. 72, at high invariant masses (high p_T values) the PDF uncertainties in the gluon–gluon luminosity become rather larger, affecting the extraction of Higgs properties in that region.

8.2. PDF uncertainties and searches for new massive particles

Many BSM physics scenarios predict the existence of new heavy particles with masses around the TeV scale. For example, supersymmetry, composite Higgs, and extra dimensions, are all classes of models, among many other others, that motivate the search for new heavy resonances at the LHC in the high-mass tail of various kinematic distributions. Here, PDF uncertainties play an important role in setting robust exclusion limits based on available null results, and would become even more important in the case of eventual discovery. In particular, PDFs represent the dominant theoretical uncertainty for the production of new heavy particles in the TeV region, as such processes are sensitive to the large- x behaviour of quarks and gluons. As discussed in Sect. 6, PDF uncertainties are large in this region due to the limited amount of experimental constraints.

In order to quantify the size of the PDF uncertainties in the large invariant mass region, as well as the relative agreement between the PDF groups, it is useful to compare the PDF luminosities for $M_X \geq 1$ TeV. We will restrict ourselves to ABMP16, CT14, MMHT14 and NNPDF3.1, in all cases using $\alpha_s(m_Z) = 0.118$. Results are shown in Fig. 73 for $\sqrt{s} = 13$ TeV normalized to the central value of the MMHT14 calculation. From the comparison in Fig. 73, we find that PDF uncertainties are small, at the few-percent level, up to $M_X \simeq 5$ TeV for the quark-quark luminosities. This is due to the fact that \mathcal{L}_{qq} is dominated by the rather accurately known up and down quark valence PDFs, which are constrained by measurements of e.g. fixed-target DIS structure functions.

For the gluon-gluon luminosity, \mathcal{L}_{gg} , we find a rather large spread in the predictions between the different groups, with MMHT14 (ABMP16) leading to the largest (smallest) central values. For instance, at $M_X \sim 5$ TeV, which is close to the upper limit of the kinematic coverage of the LHC, the envelope of the PDF uncertainty bands spans $\sim 100\%$. Even for more moderate invariant masses the spread is quite large, with the values of \mathcal{L}_{gg} at $M_X \sim 2.5$ TeV varying between $\sim +10\%$ and -30% in comparison to the central MMHT14 result. It is thus clear that these uncertainties would represent one of the limiting factors for BSM characterization in the case of an eventual discovery.

For the quark-gluon luminosity, \mathcal{L}_{qg} , we observe a similar trend to the gluon-gluon case, but with reduced PDF uncertainties due to the contribution of the well-constrained large- x quark PDFs. Interestingly, in the above two cases the results from all sets do not necessarily agree within the quoted 68% C.L. uncertainty bands. As discussed further in Sect. 6.1, it will be informative to see how this comparison changes when data sensitive to the high x gluon, such as the $t\bar{t}$ differential distributions, currently only in the NNPDF3.1 set, are included in all the fits shown in this comparison.

PDF uncertainties, as well as the differences between groups, are also large for the quark-anti-quark PDF luminosity $\mathcal{L}_{q\bar{q}}$, also shown in Fig. 73. The reason for this behaviour is two-fold. On the one hand, the large- x anti-quarks are notoriously difficult to pin down, although recent high-precision measurements from the Tevatron and the LHC are improving the situation. On the other hand, various groups parameterize the quark sea content of the proton with rather different assumptions [416] (see Sect. 5), and this can have important implications for the quark-anti-quark luminosities. We find that the spread of the different results ranges between $+5\%$ and -30% for $M_X = 3$ TeV, with PDF uncertainties becoming $O(100\%)$ for higher invariant masses. Note here that the PDF uncertainties are the largest for NNPDF3.1, despite this being the global analysis which includes the largest amount of LHC electroweak data sensitive to anti-quarks. This highlights the fact that methodological differences in the flavour assumptions and parametrization of anti-quarks are one of the dominant factors in explaining the difference between the various groups for $\mathcal{L}_{q\bar{q}}$ at large invariant masses M_X .

In order to illustrate the phenomenological consequences of these large PDF uncertainties at high M_X , in Fig. 74 we show the PDF uncertainties for high-mass graviton production in the Randall-Sundrum

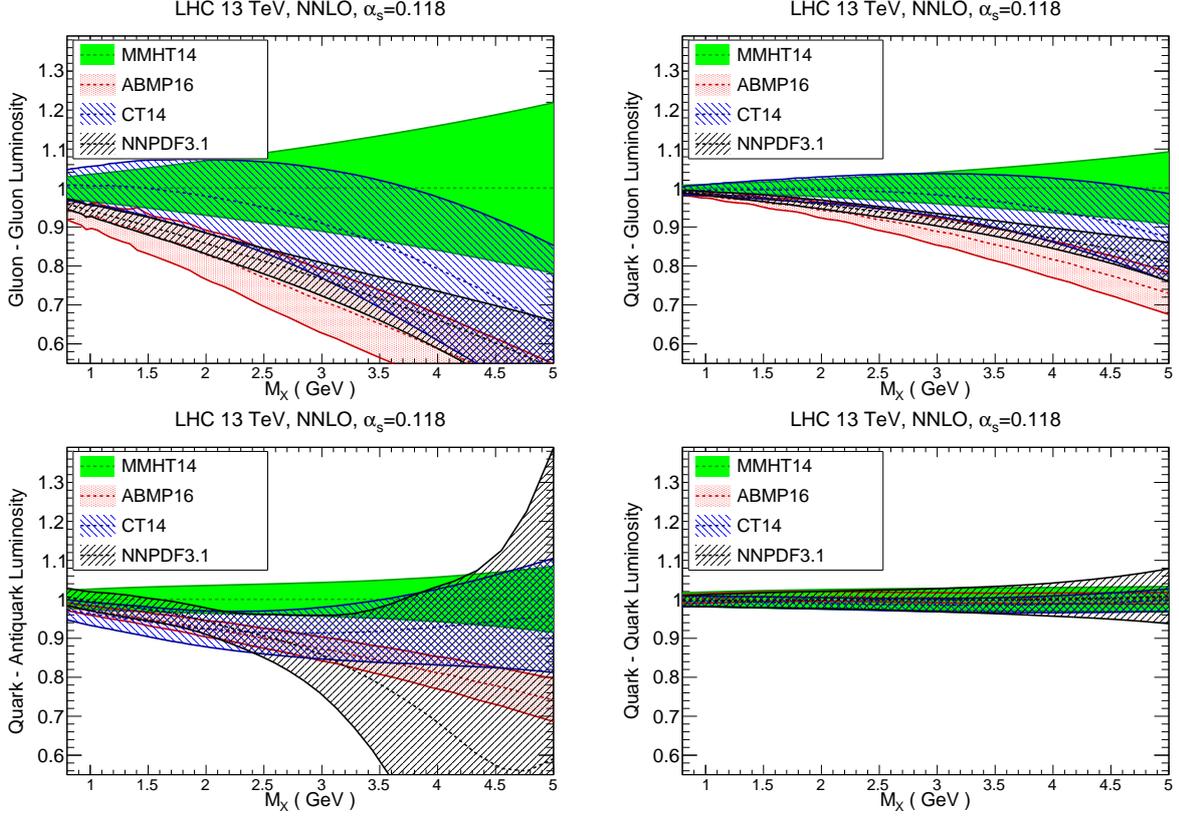


Figure 73: Comparison of PDF luminosities in the large invariant mass M_X region between MMHT14, ABMP16, CT14 and NNPDF3.1. From left to right and from top to bottom we show the results of the gluon–gluon, gluon–quark, quark–anti–quark and quark–quark luminosities, normalized to the central value of MMHT14. In this comparison, NNLO PDFs with $\alpha_s(m_Z) = 0.118$ sets are used.

scenario [547, 548] induced by gluon fusion at the LHC 8 TeV, computed at LO with MadGraph5 [549]. We compare the results of the NNPDF2.3 fit with those of the same fit including the constraints from top quark production cross sections [380]. We observe that PDF uncertainties become $\mathcal{O}(100\%)$ at large values of the graviton mass, consistent with the estimates from the gluon-gluon luminosity shown in Fig. 73. We also see how these PDF uncertainties can be reduced by the inclusion of top quark pair production total cross sections, highlighting the cross-talk between precision SM measurements and improving BSM searches. Further constraints on the large- x gluon can be obtained *e.g.* fitting top quark differential distributions, see Sect. 3.7.

In Fig. 74 we also show the K -factor for the NLO+NLL cross sections for the production of a squark-anti-squark pair $\tilde{q}\tilde{q}^*$ at the LHC 13 TeV. The calculation has been performed using various PDF sets, including PDF uncertainties, with each set normalized to the corresponding NLO result [11]. This production channel is dominated by the quark–anti–quark luminosity. We compare the predictions of NNPDF3.0, CTEQ6.6, and MSTW2008, all at NLO. Here we also see that PDF uncertainties become very large at high mass, in particular in the case of NNPDF3.0, reflecting the underlying behaviour of the quark-anti-quark luminosities.

Note that by construction the central values of the three predictions shown in Fig. 74 are rather close,

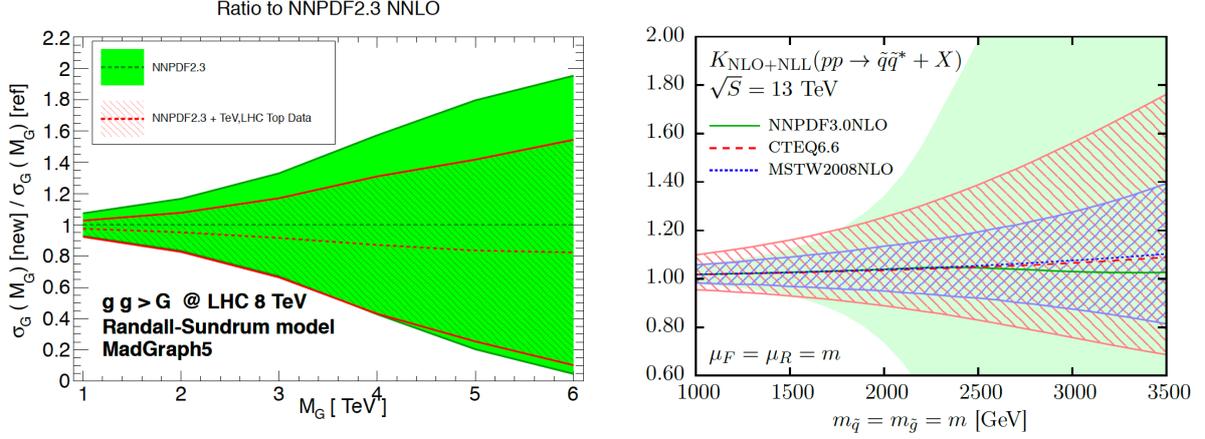


Figure 74: Left: the PDF uncertainties for high-mass graviton production in the Randall-Sundrum scenario induced by gluon fusion at the LHC 8 TeV, computed at LO with MadGraph5. We compare the results of the NNPDF2.3 fit with that of the same fit including the constraints from top quark production cross sections [380]. Right: the K -factor for the NLO+NLL cross section, including PDF uncertainties, for the production of a squark-anti-squark pair $\tilde{q}\tilde{q}^*$ at the LHC 13 TeV with various PDF sets, from [11].

since differences mostly cancel in this K -factor ratio. The usefulness of this comparison is thus related to estimating the PDF uncertainties in each case. The corresponding differences between central values can be read from the quark-antiquark PDF luminosity of Fig. 73, which is the relevant initial-state for the SUSY cross-sections shown in Fig. 74. From that comparison we see that differences in the central values of up to 40% are present at $M_X \simeq 4$ TeV, well within the LHC coverage, highlighting our poor understanding of the proton structure in the large- x region.

The results shown in Fig. 74 (together with those of Fig. 73) highlight that, in the case of an eventual discovery of novel high-mass particles at the LHC, it will be crucial to improve our knowledge of the large- x PDFs in order to be able to characterize the underlying BSM scenario. To achieve this it will be essential to exploit as much as possible high-precision collider data, mostly from the LHC, in fits, in order to pin down the large- x gluons and anti-quarks and thus reduce the PDF uncertainties associated to high-mass BSM particle production. For example, in Ref. [379] it was shown that by including the y_t and $y_{\bar{t}}$ differential distributions for top quark pair production in a global PDF fit, it is possible to reduce the PDF uncertainties that affect the high-mass tail of the $m_{\bar{t}}$ distribution by up to a factor of two. This distribution is widely used for searches, for instance of new resonances that couple strongly to the top quark [550]. In the future, it might also be feasible to provide indirect constraints on BSM models, for instance on the coefficients of the SM-EFT [551] higher-dimension operators, by including these in the global PDF fit, along the lines of early studies aiming to constrain colored sparticles from Tevatron jet production [552].

8.3. Precision measurements of SM parameters

The precision measurements of SM parameters, such as the mass of the W boson M_W or the running of the strong coupling constant $\alpha_S(Q)$, represent powerful ways of constraining BSM dynamics at the LHC. For instance, following the discovery of the Higgs boson, in the absence of new physics the SM is an over-constrained theory: one can use a set of input parameters, such as the Higgs mass m_h and the top quark mass m_t , in the context of the global electroweak precision fit [553], to predict other parameters, such as M_W . By comparing these indirect predictions of the W mass with direct experimental measurements, one can

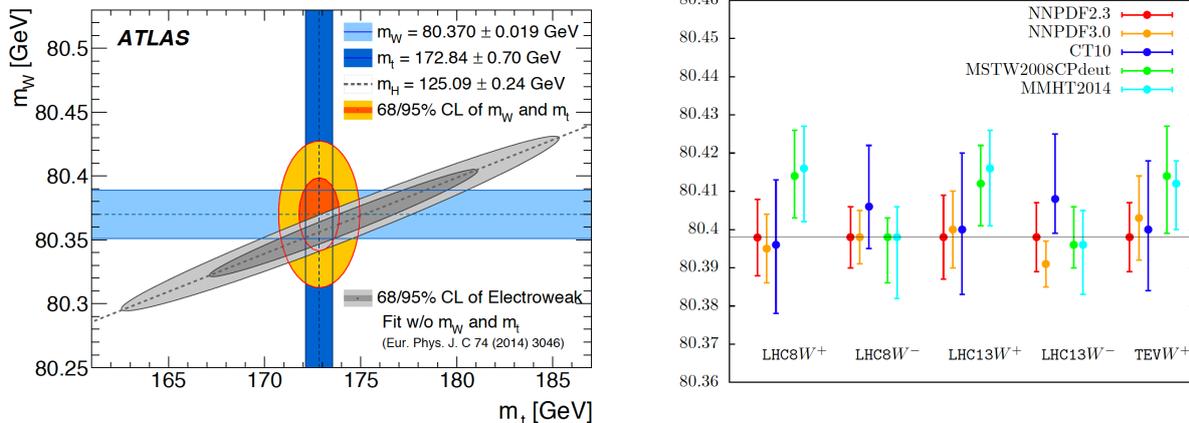


Figure 75: Left: comparison between the direct measurements of m_W , m_t , and m_h from ATLAS with the predictions from the global electroweak fit, from [15]. Right: estimate of the PDF uncertainties in the m_W determination using different PDF sets and collider scenarios, from [12]. This estimate has been obtained from template fits to the p_T^l distribution, imposing the constraint that $p_T^W \leq 15$ GeV.

provide a stress test of the SM, where any tension might indicate hints for BSM dynamics at scales higher than those that are currently directly accessible. Tests of this type have already resulted in the famous $g_\mu - 2$ anomaly, where a persistent 3 to 4-sigma discrepancy is found between the theoretical predictions of the muon anomalous magnetic moment [554] and the corresponding experimental measurements.

In order to make these comparison between indirect predictions and direct measurements as stringent as possible, it is important to improve the precision of the latter. For many SM parameters, PDF uncertainties represent one of the dominant theoretical uncertainties in their determination, providing another motivation for the need of improved PDFs. Focusing on the case of the W mass measurements, the role of PDF uncertainties has been quantified in detail in a number of studies, both from the phenomenological [12–14] and from the experimental point of view. In the latter case, the first direct measurement of M_W at the LHC has recently been presented by the ATLAS collaboration [15], yielding a total uncertainty of only 19 MeV, of which around half of it (9 MeV) is estimated to come from PDF uncertainties. In Fig. 75 we show a comparison between the direct measurements of m_W , m_t , and m_h from ATLAS with the predictions from the global electroweak fit, from [15]. We observe that there is good agreement between the direct measurements and the indirect predictions, providing a highly non-trivial validation test of the SM. Future measurements of m_W and m_t , as well as their combination with other experiments, should be able to reduce the uncertainties in this comparison.

In Fig. 75 we also show a phenomenological estimate of the PDF uncertainties associated to the m_W measurements using different PDF sets and collider scenarios, from [12]. This estimate has been obtained from template fits to the p_T^l distribution, imposing the additional constraint that $p_T^W \leq 15$ GeV. A number of NNLO PDF sets are used in this comparison, in order to achieve a robust estimate of the PDF uncertainties. In general one finds that there is good agreement within PDF uncertainties, although in some cases this agreement is only marginal, as in the case of NNPDF3.0 and CT10 at the LHC 13 TeV for the W^- case. From this study, one estimates that at the LHC 7 TeV PDF uncertainties using state-of-the-art sets are around 6 MeV, a number close to the one estimated in the ATLAS measurement [15].

Another SM parameter that could potentially provide indirect information on BSM dynamics is the

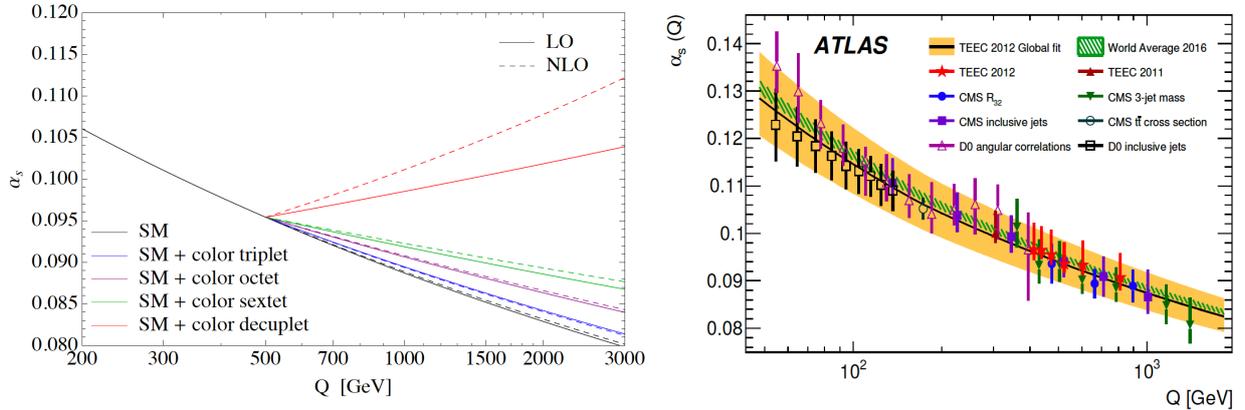


Figure 76: Left: the modifications in the running of $\alpha_s(Q)$ induced by a new heavy colored fermion of mass $m = 0.5$ TeV as compared to the SM expectation, for various representations of its color gauge group, from [21]. Right: comparison of recent direct determinations of $\alpha_s(Q)$ at the Tevatron and the LHC as a function of Q , together with the PDG 2016 world average and with the results of the ATLAS TEEC 2012 extraction [17].

QCD coupling $\alpha_s(Q)$, and specifically its running at the TeV scale. It is well known that the presence of new colored heavy degrees of freedom will modify the QCD β -function and therefore lead to a different running with Q as compared to the corresponding SM prediction. This fact forms for example the basis of the improved agreement at high scales between the strong, weak and electromagnetic couplings in the case of low-scale supersymmetry, which suggest the unification of the three forces at a GUT scale of around $\Lambda_{\text{GUT}} \sim 10^{16}$ GeV [22].

If new BSM heavy particles are present at the TeV scale, the difference induced in the QCD coupling running could be accessible at the LHC, see e.g. [21] and Fig. 76, where we show the modifications in the running of $\alpha_s(Q)$ induced by a new heavy colored fermion of mass $m = 0.5$ TeV for various representations of its color gauge group. With this motivation, as well as to compare with other measurements of α_s at lower energies, the ATLAS and CMS collaborations have presented a number of measurements of both $\alpha_s(m_Z)$ and of $\alpha_s(Q)$ for restricted Q ranges, mostly from jet production [16–18] but also from top quark pair production [20] (see also [19] for a review, and Fig. 76 for a graphical overview).

In these collider-based determinations of the strong coupling, PDF uncertainties, which are significant at the TeV scale (see Sect. 8.2), represent an important source of theoretical uncertainties. For instance, in the recent ATLAS determination of $\alpha_s(m_Z)$ from transverse energy–energy correlations (TEEC) at 8 TeV [17], the PDF uncertainty is $\delta_{\text{pdf}} = 0.0018$, almost a factor 2 larger than the experimental uncertainty of $\delta_{\text{exp}} = 0.0011$. While in this analysis PDF uncertainties are sub-dominant with respect to the scale uncertainties, $\delta_{\text{scale}} \simeq 0.006$, the latter were computed using NLO theory and can be reduced significantly by exploiting the NNLO calculation. Likewise, in the CMS measurement at 7 TeV based on a QCD analysis of the inclusive jet cross sections [16], one finds that the PDF uncertainties $\delta_{\text{pdf}} = 0.0028$ are larger than the experimental uncertainties $\delta_{\text{exp}} = 0.0019$, though still sub-dominant with respect to the large scale variations of the NLO calculation $\delta_{\text{scale}} = {}^{+0.0053}_{-0.0024}$. In both cases, it is clear that if one is able to reduce PDF uncertainties, and exploit the reduction of scale errors arising from the NNLO calculation, it should be possible to achieve a rather competitive determination of $\alpha_s(m_Z)$ and also perform stringent tests of its running in the TeV region.

As a related topic, we would like to mention that there are also proposals to measure the running of the

electroweak running couplings at the LHC [555] and use these to impose model-independent constraints on new particles with electroweak quantum numbers without any assumptions about their decay properties. Again, here PDFs are one of the dominant sources of theoretical uncertainty, for instance in the high-mass tail of W and Z production at the LHC due to the quark-anti-quark luminosity. Improving our knowledge of large- x anti-quarks will therefore be helpful in providing such indirect constraints of new heavy electroweak sectors.

9. The future of PDF determinations

In the final section of this Report we discuss three topics that could play an important role in shaping global analyses of PDFs in the coming years. First we discuss the problem of theoretical uncertainties in PDF fits, whose estimate is becoming increasingly necessary given the size of PDF uncertainties in recent global analysis. Second, we summarise recent progress in lattice QCD computations of PDFs, including the first efforts towards a determination of their x -dependence, and discuss how in the near future lattice inputs could contribute to global PDF fits. Third, we briefly review the status and plans for future high-energy colliders, such as the Large Hadron electron Collider (LHeC) or the Future Circular Collider (FCC), and the role that PDFs would play in these.

9.1. PDFs with theoretical uncertainties

The development of sophisticated methodologies for PDF fits, as well as the availability of a wealth of high-precision data, have reduced the PDF uncertainties in global analysis, arising mostly from experimental data and procedural choices, to the few percent-level in the most constrained regions. At this level of accuracy, various theoretical uncertainties become more and more important, representing a major limitation for present and (even more) for future studies. Therefore, a robustly account of these theoretical uncertainties is one of the main goals of PDF fitters for the near future.

In this section, we focus specifically on the role of the theoretical uncertainties due to missing higher orders (MHO) in the QCD coupling constant, namely those arising from the truncation of the asymptotic perturbative expansion. There have been a number studies recently on how to estimate MHO, although we are still far from a conclusive answer. In the following we first review progress on MHO of calculations of non-hadronic and hadronic processes and then several recent studies related to PDF determination.

We emphasise that these theoretical uncertainties from MHO should not be confused with the parametric theoretical uncertainties, that is, those arising from the choice of the values of input parameters such as α_s and m_c . These have been reviewed in Sect. 4.5, and the procedure for estimating the impact on the PDF fit and to propagating these parametric uncertainties into collider cross sections is well established.

9.1.1. MHO in matrix element calculations

The most frequently used and probably also the simplest method of estimating the MHO corrections is the variation of the renormalization and factorization scales in a given fixed-order calculation. In the case of the total inclusive cross sections in e^+e^- collisions or decay rates with a single hard scale Q , usually one varies the QCD renormalization scale μ_R within the interval $[Q/r, rQ]$. The induced changes in the physical observable, either from three-point evaluations or from a scan over the entire range, are taken as the *uncertainty* of the MHO, assuming then typically either a Gaussian (or two half-Gaussians) or a flat distribution. The conventional choice is $r = 2$, which is found to work well in most cases, but that underestimates the true higher order corrections in certain cases, especially if the fixed-order calculation is carried out at leading order.

At hadron colliders, there exists in addition the factorization scale μ_F , arising from the factorization of collider QCD divergences due to initial-state hadrons. The two scales μ_R and μ_F can be varied either simultaneously or independently within the above range, with the latter case usually further restricted to $1/r \leq \mu_F/\mu_R \leq r$. However, even for a single scale problem, there can still be different choices of the central or nominal scale, e.g., $Q/2$ or $2Q$, motivated by either QCD resummation or speed of convergence of the series [41], which leads to further ambiguities in the estimation of theoretical uncertainties from scale variations. There are also some studies on utilizing the so-called principal of maximum conformality to determine the QCD renormalization scale at different orders, which claims a much smaller MHO [556] than traditional scale variations.

Determining a suitable prescription for scale variations becomes more complicated when moving to differential observables, since here more hard scales, including those related to the kinematics, are involved. This therefore usually requires a dynamical choice for the central scale, which often also depends on the specific distribution considered. For example, in a recent study on the hadronic production of top quark pairs [41] it was shown that the preferable scale is half of the transverse mass of the top quark when studying the transverse momentum distribution of the top quark, and one fourth of the sum of transverse mass of top quark and anti-quark when studying rapidity distribution of the top quark. Starting from a given choice of the nominal scale, the scale variations can then be evaluated in a similar way to the inclusive case and serve as an estimate of MHO. One further complexity arises concerning the correlations of the MHO or scale variations in different regions of the distribution. Typically, they are assumed to be fully (anti-)correlated in the entire region which leads to very small theoretical uncertainties in the case of a normalized distribution. There have however been attempts to decorrelate these scale variations based on consideration of the kinematic dependence of the QCD corrections [557].

Alternative proposals for estimating MHO based on results at known orders exists, such as the so-called Cacciari-Houdeau (CH) approach [558]. The basic idea is to express the full perturbation series in terms of the expansion parameter $\alpha_s(Q)$ and assume that all the expansion coefficients follow the same uniform bounded probability density distribution, in the Bayesian sense. Bayesian inference can be used to calculate the probability density of the unknown higher-order coefficients given those known coefficients at lower orders. Thus the MHO, including its probability density distribution, which will be non-Gaussian in general, can be constructed. The original CH method was developed for the study of non-hadronic processes. Subsequently, in the modified CH ($\overline{\text{CH}}$) [559] approach, it was generalised to include hadronic processes. There the expansion parameter has been adjusted to $\alpha_s(Q)/\lambda$, with the parameter λ determined from a global survey of selected processes with known higher-order corrections. To be specific, the best value of λ computed from the predicted probability density of higher orders is required to match the distribution from a frequency count in the survey. For hadronic processes, the optimal value of λ is found to be about 0.6, meaning that the *true* perturbative expansion parameter is actually around $1.7\alpha_s(Q)$ rather than $\alpha_s(Q)$.

To give another example, the series acceleration method [560] can also be applied to approximate the full result for physical observables based on the available information from a finite number terms of the asymptotic series, e.g., using Levin-Weniger sequence transforms. In the ‘Passarino-David’ method of Ref. [560] a uniform distribution (in the Bayesian sense) for the theory prediction is assumed between the last known partial sum of the perturbation series and its approximated value from a Weniger δ -transform.

To illustrate how these various methods compare to each other, Fig. 77 shows the predictions for the production cross sections of the Higgs boson via gluon fusion at the LHC 8 TeV, calculated at LO, NLO, NNLO and approximate N³LO with a nominal scale of $\mu_R = \mu_F = m_H$ [561]. The MHO as estimated from different approaches are shown, including scale variations, CH, $\overline{\text{CH}}$ and the series acceleration method

of [560] at various perturbative orders. Note that different approaches may have different interpretations on the uncertainties. In the case of the $\overline{\text{CH}}$ method, the λ value has been adjusted to give almost equal expansion coefficients for the known orders [561]. The $\overline{\text{CH}}$ method predicts a larger uncertainty, while the scale variation uncertainty is larger than the CH result. Encouragingly, the difference in the MHOUs is observed to decrease with increasing perturbative order. Indeed, the different approaches turn out to give similar sizes for the MHOUs associated to the N^3LO calculation, with the exception of the original CH. However, for case of the Higgs boson production, none of the methods work very well, i.e., the N^3LO cross section is not captured by uncertainty bands from lower orders. In addition, we note that the series acceleration method also induces a shift in the central value of the prediction.

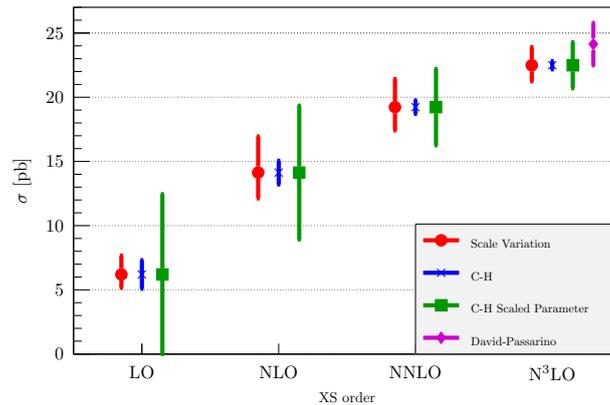


Figure 77: The cross section for Higgs boson production in gluon fusion calculated at increasing perturbative orders [561]. At each order the theoretical uncertainty is shown for using scale variation (red circles), the CH method (blue crosses), and the $\overline{\text{CH}}$ method (green squares); at N^3LO the Passarino–David uncertainty based on series acceleration method is also shown (purple diamonds).

9.1.2. MHOUs in PDF determination

Global determinations of PDFs are based on perturbative calculations of matrix elements and DGLAP splitting kernels, suitably combined to predict a variety of physical cross sections. In these perturbative calculations, in principle one should account for their associated MHOUs, which then propagates into the resulting PDFs via the fitting of the theoretical predictions to the experimental data. Therefore, sophisticated treatments of the MHOUs from different sources are required in order to study the impact on the PDFs, not unlike the treatment of the experimental systematic uncertainties. Crucially, the correlations between theoretical predictions of different experimental bins of one process and between different processes must be accounted for. Furthermore, when making any theoretical prediction, one should also take care of the correlations between the MHOUs of the PDFs and of those coming from the MHOUs of the process studied, since they may rely on the same perturbative expansion for the relevant matrix elements.

Due to the significant complexity of this problem, there is still no satisfactory solution, and the MHOUs have not been included in any of the public PDFs from global determinations. However, it is possible to restrict ourselves to a region where a single process is most likely dominant in the MHOUs, such as inclusive jet cross sections and the gluon PDF at large x . In such a case, it may still be possible to construct simple prescriptions such as the use of scale variations. To illustrate this, in Fig. 78 (Left) we show the impact of the choice of the QCD scales in calculations of the inclusive jet cross sections on the gluon PDFs at $Q = 85$ GeV for alternative CT10 NNLO fits [125]. Note that the theoretical predictions here are only at NLO,

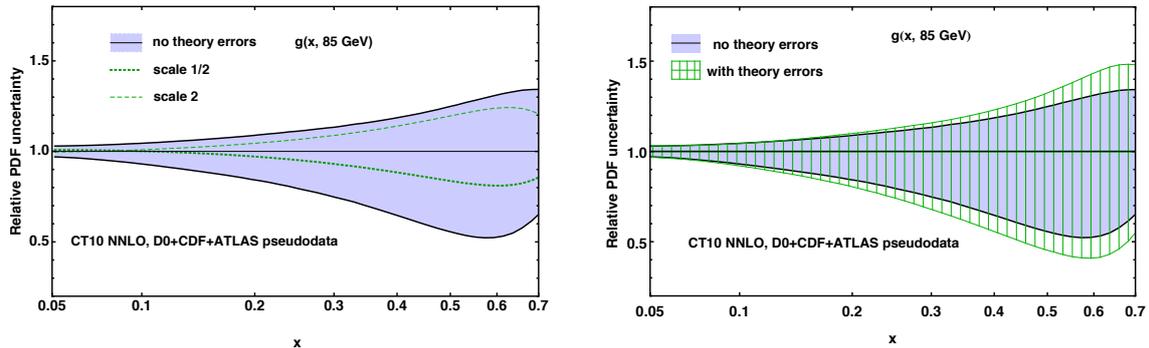


Figure 78: Left plot: dependence of the gluon PDF on the choice of QCD scales used in the calculation of inclusive jet cross sections in CT10 NNLO fits [125] by changing scales to half and twice of the nominal value. Right plot: impact of the theoretical uncertainties from the inclusive jet cross sections included in the CT10 NNLO fits [1] in the resulting gluon PDF uncertainties.

though the PDFs are determined at NNLO.

From Fig. 78 we can see that in the higher $x > 0.2$ region, the spread of the gluon PDFs by using scales of 0.5, 1 and 2 times the central scale (in this case the jet p_T) is not negligible in comparison to the nominal PDF uncertainties. This illustrates the potential significance of the MHOU in PDF determination. In a related study, also based on alternative CT10 NNLO fits, scale variations of the NLO inclusive jet cross sections are further decomposed into several correlated systematics described by five nuisance parameters [1]. By treating those systematics in a similar way to the experimental correlated systematic errors, it is possible to include the MHOU in the standard PDF uncertainty on the same footing as the experimental systematic uncertainties. As shown in Fig. 78 (Right) the inclusion of theory errors from the jet cross sections in the CT10 NNLO fit results in an increase of the gluon PDF uncertainty at large x , consistent with Fig. 78 (Left).

Another possibility to provide a rough estimate of the MHOU consists in checking the convergence of the fitted PDFs with increasing orders. Fig. 79 shows the comparison of the nominal PDF uncertainties with the difference of the central PDFs fitted at NLO and NNLO for gluon and total singlet PDFs at $Q = 100$ GeV in the NNPDF3.0 fits [30]. This difference between the PDF central values at NLO and NNLO provides a conservative upper bound of the MHOU associated to the NNLO PDFs. From this comparison, we see that there are regions where the shifts of NLO to NNLO PDFs are comparable to or even larger than the conventional PDF uncertainties. Based on similar approaches Ref. [561] found that the MHOU of the NNLO PDFs have a negligible impact on the Higgs production cross section through gluon fusion, but on the other hand it could be relevant for top quark pair production.

9.2. Lattice QCD calculations of the proton structure

As discussed in Sect. 2.2, PDFs arise from non-perturbative QCD dynamics. Therefore, since we are currently not able to analytically solve strongly-coupled non-Abelian gauge theories, it is very challenging to compute PDFs from first principles. Perhaps the only possibility, beyond model calculations, consists in exploiting recent progress in lattice QCD [562]. This method is based on discretizing the QCD Lagrangian in a finite-volume Euclidean lattice, which naturally introduces an ultraviolet cutoff, and then computing non-perturbative QCD quantities directly on this lattice, before taking the continuum limit. Such lattice QCD calculations require minimal external input, in particular only the hadronic mass scale Λ_{QCD} and the values of the quark masses, or alternatively, the physical pion and kaon masses.

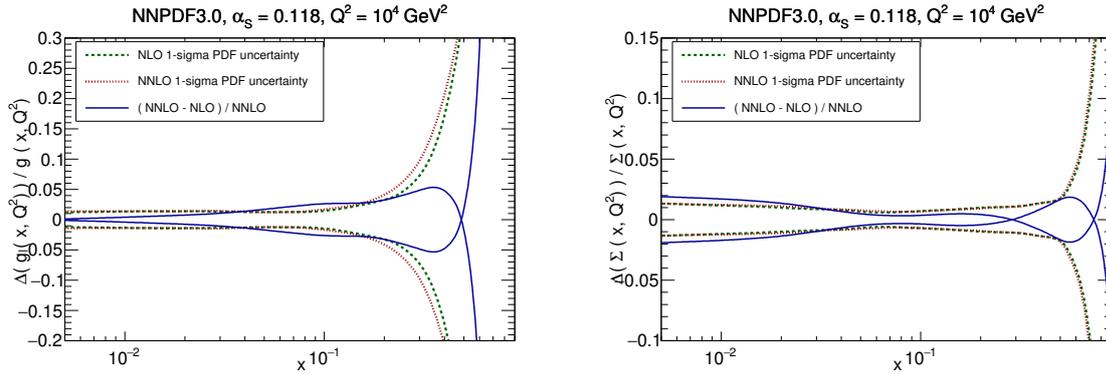


Figure 79: Left plot: comparisons between the nominal PDF uncertainties with the difference of the central PDFs determined at NLO and NNLO for the gluon PDF at $Q = 100$ GeV in the NNPDF3.0 fit. Right plot: same for the total quark singlet, $\Sigma(x, Q^2)$.

Here we briefly review some of these recent developments in lattice QCD calculations of PDFs. For a more detailed overview of this progress, together with the study of their interplay with state-of-the-art global analysis, see Ref. [563], a White Paper that has been produced as the outcome of the dedicated workshop “*Parton Distributions and Lattice QCD calculation in the LHC era*”,¹⁴ which took place in Oxford in March 2017. This workshop brought together experts from the two communities (lattice QCD and global PDF fitting) to explore the synergies and complementary aspects between the two approaches. The discussion and results shown in this section represent an brief excerpt of the material contained in that White Paper (see also [564]).

Given that PDFs have a formal definition in terms of the nucleon matrix elements of certain non-local operators (see the discussion in Sect. 2.2), it is in principle possible to compute PDFs using lattice QCD. From the practical point of view, however, given the extremely CPU intensive nature of these calculations, most lattice QCD results on PDFs have been limited for a long time to the first two moments of non-singlet flavour combinations for large (unphysical) quark masses. These restrictions have been overcome in the recent years, with several groups providing now results for PDF moments with physical pion and kaon masses.

Moreover, the development of novel strategies to overcome the limitations in computing the first few moments [565–567] has allowed the determination of more challenging quantities, such as gluon and flavour-singlet matrix elements as well as higher PDF moments. More recently, both conceptual and numerical breakthroughs in lattice QCD computations have allowed even greater progress, providing the first attempts to evaluate PDFs and related quantities directly in Bjorken- x space [568–571]. These developments have pushed lattice QCD calculations to the point where, for the first time, it is possible to provide information on the PDF shape of specific flavour combinations, both for quarks and for antiquarks, and meaningful comparison with global fits can start to be made.

Recent progress in lattice QCD calculations of PDFs and related quantities has also been partly driven by a greatly improved control on the systematic uncertainties that enter the calculation of relatively simple quantities such as nucleon matrix elements, which correspond to low moments of the PDFs. These include, among others, using physical pion masses, reducing the excited-state contamination, and using large lattices to remove finite-size effects. Moreover, to make contact with the physical world and experimental

¹⁴<http://www.physics.ox.ac.uk/confs/PDFlattice2017>

data, the numerical results are extrapolated to the continuum and infinite-volume limits. In addition, the past decade has seen significant progress in the development of efficient algorithms for the generation of ensembles of gauge field configurations, which represent the QCD vacuum, and tools for extracting the relevant information from lattice QCD correlation functions.

In order to be reliably used in phenomenological applications, lattice QCD calculations must demonstrate control over all relevant sources of systematic uncertainty introduced by the discretisation of QCD on the lattice, such as those discussed above. For a coherent assessment of the present state of lattice QCD calculations of various quantities, the degree to which each systematic uncertainty has been controlled in a given calculation should be an important consideration. The quality of individual lattice calculations can be quantitatively assessed based on criteria such as those from the FLAG analysis of flavour physics on the lattice [572].

The traditional approach for lattice QCD calculations of PDFs has been to determine the matrix elements of local twist-two operators, where twist is the dimension minus the spin, that can be related to the Mellin moments of PDFs. In principle, given a sufficient number of Mellin moments, PDFs can be reconstructed from the inverse Mellin transform. In practice, however, the calculation is limited to the lowest three moments, because power-divergent mixing occurs between twist-two operators on the lattice. Three moments are unfortunately insufficient to reconstruct the full x -dependence of the PDFs without significant model dependence [573]. The lowest three moments do provide, however, useful information, both as benchmarks of lattice QCD calculations and as constraints in global extractions of PDFs. For instance, provided systematic uncertainties are kept under control, one can envisage adding these lattice QCD calculations of PDF moments as an additional theoretical constraint to the global fit, on the same footing as the momentum and valence sum rules.

Now we briefly summarize the state-of-the-art of lattice QCD calculations of the first moment of unpolarized PDFs, which are those for which systematic uncertainties are under better control. See Ref. [574] for a more exhaustive set of comparisons, including those of lattice QCD calculations with global fits. The observables that are discussed here are defined as follows:

1. The second moment of the flavour triplet combination, $T_3 = u^+ - d^+$,

$$\langle x \rangle_{u^+ - d^+}(\mu^2) \Big|_{\mu^2 = Q^2} \equiv \int_0^1 dx x \{u(x, Q^2) + \bar{u}(x, Q^2) - d(x, Q^2) - \bar{d}(x, Q^2)\}. \quad (144)$$

2. The second moment of the individual quark $q^+ = q + \bar{q}$ PDFs,

$$\langle x \rangle_{q^+ = u^+, d^+, s^+, c^+}(\mu^2) \Big|_{\mu^2 = Q^2} \equiv \int_0^1 dx x \{q(x, Q^2) + \bar{q}(x, Q^2)\}. \quad (145)$$

3. The second moment of the gluon PDF,

$$\langle x \rangle_g(\mu^2) \Big|_{\mu^2 = Q^2} \equiv \int_0^1 dx x g(x, Q^2). \quad (146)$$

In Table 2 we show a selection of recent results for the moments defined in Eqns. (144-146). As can be seen, for $\langle x \rangle_{u^+ - d^+}$ the lattice QCD uncertainties vary between 5% and 15%, with the quoted results not agreeing among themselves within errors. For the first moment of the gluon, $\langle x \rangle_g$, the uncertainties are around 10%, and for the individual total quark combinations they vary between 10% and 20%. Therefore, while current determinations of the first moments are unlikely to provide constraints on global PDF fits (where uncertainties are the few-percent level), future calculations with improved systematic errors might be able

Mom.	Collab.	Ref.	N_f	Value
$\langle x \rangle_{u^+ - d^+}$	LHPC 14	[575]	2+1	0.140(21)
	ETMC 17	[576]	2	0.194(9)(11)
	RQCD 14	[577]	2	0.217(9)
$\langle x \rangle_{u^+}$	ETMC 17	[576]	2	0.453(57)(48)
$\langle x \rangle_{d^+}$	ETMC 17	[576]	2	0.259(57)(47)
$\langle x \rangle_{s^+}$	ETMC 17	[576]	2	0.092(41)(0)
$\langle x \rangle_g$	ETMC 17	[576]	2	0.267(22)(27)

Table 2: Selected recent lattice QCD calculations of the first moments of unpolarized PDFs, defined in Eqns. (144-146), evaluated at $\mu^2 = 4 \text{ GeV}^2$. See Ref. [574] for more details.

to make a difference. On the other hand, existing calculations can already be used to provide meaningful constraints on polarized PDFs, where uncertainties are rather larger than in the unpolarized case due to the scarcer dataset.

To quantify this point, in Fig. 80 we show the comparison between lattice QCD calculation of different moments of unpolarized PDFs at $Q = 2 \text{ GeV}$ with the corresponding results from global PDF fits, both in absolute scale (left) and normalized to the lattice result (right plot). We find that there is agreement within uncertainties for most cases, except for the gluon, but that lattice QCD calculations are still far from being competitive with the global PDF fit results. However it is also clear that future lattice calculations with reduced systematic uncertainties could be able to produce useful constraints in the global fit framework.

While the lowest moments of the PDFs provide crucial benchmarks to assess the reliability of lattice QCD calculations of the nucleon structure, as well as potentially useful information for global PDF fits, they do not allow the complete x -dependence of the PDFs to be reconstructed. In particular, the calculation of PDF moments is mostly insensitive to the small- x region. To bypass these limitations, recently a number of approaches have been developed, aiming to determine the x -dependence of PDFs directly from lattice QCD. One of the most important approaches goes under the name of ‘quasi-PDFs’, first formulated by Ji in Refs. [578, 579]. For simplicity, we focus in the following on the flavour nonsinglet case so that we can neglect the mixing with gluons.

In this approach, the unpolarized quark quasi-PDF $\tilde{q}(x, \Lambda, p_z)$ is defined as a momentum-dependent nonlocal static matrix element:

$$\tilde{q}(x, \Lambda, p_z) \equiv \int \frac{dz}{4\pi} e^{-ixzp_z} \frac{1}{2} \sum_{s=1}^2 \langle p, s | \bar{\psi}(z) \gamma_z e^{ig \int_0^z A_z(z') dz'} \psi(0) | p, s \rangle, \quad (147)$$

where Λ is an UV cut-off scale, typically chosen to be the inverse lattice spacing $1/a$. As we see from Eq. (147), these quasi-PDFs are defined for nucleon states at finite momentum (as p_z is finite, the momentum fraction x can be larger than unity). Therefore, in order to make contact with the standard collinear PDFs and thus with phenomenology, they must be related to the corresponding light-front PDF, for which the nucleon momentum is taken to infinity.

In the large-momentum effective field theory (LaMET) approach, the quasi-PDF $\tilde{q}(x, \Lambda, p_z)$ can be

related to the p_z -independent light-front PDF $q(x, Q^2)$ through the following relation [578, 579]

$$\tilde{q}(x, \Lambda, p_z) = \int_{-1}^1 \frac{dy}{|y|} Z\left(\frac{x}{y}, \frac{\mu}{p_z}, \frac{\Lambda}{p_z}\right)_{\mu^2=Q^2} q(y, Q^2) + O\left(\frac{\Lambda_{\text{QCD}}^2}{p_z^2}, \frac{m^2}{p_z^2}\right), \quad (148)$$

where μ is the renormalisation scale; Z is a matching kernel; and m is the nucleon mass. Here the $O(m^2/p_z^2)$ terms are target-mass corrections and the $O(\Lambda_{\text{QCD}}^2/p_z^2)$ terms are higher twist effects, both of which are suppressed at large nucleon momentum. A complementary approach to the LaMET methods instead views the quasi-PDF as a ‘lattice cross section’ from which the light-front PDF can be factorized [580, 581]. An alternative construction, using so-called ‘pseudo-PDFs’, has been explored in [582, 583].

Preliminary results from lattice calculations of quasi-PDFs have been rather encouraging [568–571], although a number of important limitations still need to be overcome. To illustrate this progress in lattice calculations of x -space PDFs, in Fig. 80 we also display the comparison between lattice QCD and global PDF fit results now for the unpolarized isovector quark distributions, $x(u - d)$ (left) and $x(\bar{u} - \bar{d})$ (right plot) at $Q = 2$ GeV as a function of Bjorken- x . See the discussions in the White Paper [563] for more details. From this comparison, we see that the lattice QCD results are still far from the global fits (and thus also from the experimental data), but given recent progress we expect sizable improvements in the coming years. In addition, the same methodology can be also applied to other nucleon matrix elements, including the polarized PDFs.

Despite these remarkable developments, there still remain a number of important challenges that must be overcome before one can achieve a complete determination of the x -dependence of PDFs directly from lattice QCD that is competitive with global PDF fits. In particular, excellent control over the various sources of systematic uncertainties that affect the calculation must be reached. Some of these are common to the calculations of PDF moments, as discussed above, but there are also a number of additional systematic errors specific to quasi-PDFs, such as those associated with the finite nucleon momentum of the lattice calculations and with the renormalisation of quasi-PDFs. Given the recent fast progress, it is quite possible that this may be achieved and that in the future lattice QCD calculations of x -space PDFs may be used to constrain global analyses.

9.3. Parton distributions at future high-energy colliders

Now we turn to discuss the role that PDFs would have in some of the recent proposals for future colliders involving hadrons in the initial state. There are three main families of potential future high-energy colliders currently under active discussion. Electron-positron colliders, such as the ILC [584], CLIC [585], TLEP/FCC-ee [586], or CEPC [587], offer the potential for ultra-high precision measurements of the Higgs, electroweak and top quark sectors. On the other hand, hadron colliders with energy much greater than the LHC would allow the continued exploration of the high-energy frontier and significantly extend the coverage of searches for new BSM particles, including Dark Matter candidates, while at the same time providing unprecedented opportunities for the study of the Higgs sector, for example the Higgs self-interactions. With this motivation, there is ongoing work towards a circular collider hosted at the CERN site which would accelerate protons up to the extreme energies of $\sqrt{s} = 100$ TeV [28, 588], dubbed FCC-hh, while a similar machine is under study by the Chinese HEP community [587].

Another avenue for future high-energy collisions would be new machines based on electron-proton collisions, exploiting the successful strategy adopted at HERA. One of the open proposals is the Large Hadron electron Collider (LHeC) [29], where the $E_p = 7$ TeV proton beam from the LHC would collide with an electron/positron beam with $E_e = 60$ GeV coming from a new LinAc, and that would be able to

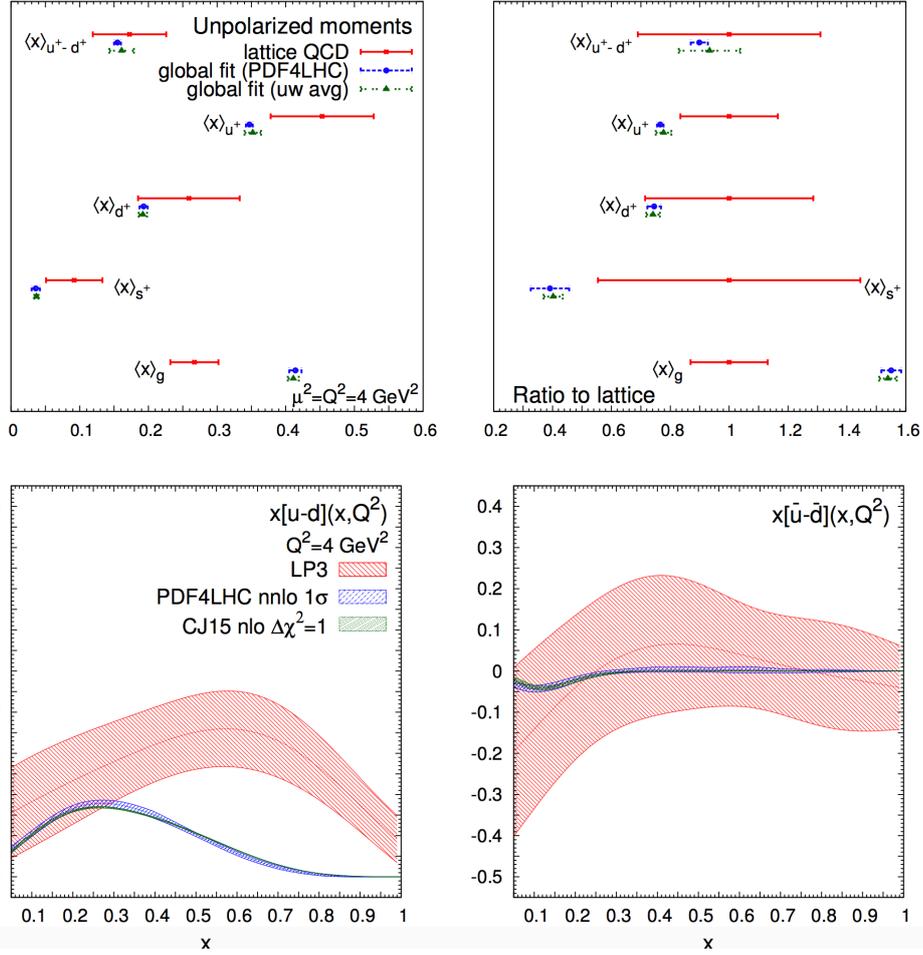


Figure 80: Upper plots: comparison between lattice QCD calculation of different moments of unpolarized PDFs at $Q = 2 \text{ GeV}$ with the corresponding results from global PDF fits, both in absolute scale (left) and normalized to the lattice result (right plot). Bottom plots: the same comparison now for the unpolarized isovector quark distributions, $x(u - d)$ (left) and $x(\bar{u} - \bar{d})$ (right plot) at $Q = 2 \text{ GeV}$. See the discussion in the White Paper [563] for more details.

reach the region down to $x_{\min} \simeq 2 \cdot 10^{-6}$ at $Q^2 = 2 \text{ GeV}^2$. A more extreme incarnation of the same idea corresponds to colliding these $E_e = 60 \text{ GeV}$ electrons with the $E_p = 50 \text{ TeV}$ beam of the FCC–hh. The resulting collider, dubbed FCC–eh, would be able to reach down to $x_{\min} \simeq 2 \cdot 10^{-7}$ at $Q^2 = 2 \text{ GeV}^2$. These two machines would therefore extend the PDFs coverage at small x by more than two orders of magnitude in comparison to HERA.

The Electron Ion Collider (EIC) [589], which might start construction soon either at the BNL or the JLAB site, falls under the same category. The EIC would offer the possibility to polarize both leptons and protons and to accelerate heavy nuclei, although its \sqrt{s} would be smaller than that of HERA. It would therefore not be able to provide new information on the unpolarized proton PDFs, though it would be a unique machine to explore the spin and 3D content of the proton, as well as map in detail the nuclear modifications of the free–proton PDFs.

In this section, we review the role that PDFs would play first at a future higher energy lepton–proton collider such as the LHeC/FCC–eh, and then at a future hadron collider with a centre–of–mass energy of

$$\sqrt{s} = 100 \text{ TeV.}$$

9.3.1. PDFs at high-energy lepton-hadron colliders

As mentioned above, one of the possibilities for a future high-energy collider now under active discussion would be to exploit the LHC/FCC proton beam and collide it with a high energy lepton beam, which would be delivered by a new LinAc to be built at the CERN site. In the case of using the LHC beams, the LHeC [29, 590, 591] would then represent a scaled-up version of HERA, and as such would offer immense opportunities for improved determinations of the proton structure down to very low x and up to high Q^2 , as well as providing a wealth of information on nuclear PDFs in a kinematic region where they are currently essentially unconstrained. Several options are now being considered, with some preference now for synchronous operation during the final years of the HL-LHC upgrade, since then the LHeC program can be extended to include measurements of the Higgs sector.

In Fig. 81 we show the kinematic coverage in (x, Q^2) of several existing and future DIS experiments, including the EIC, the LHeC, and the FCC-eh. We observe that by starting from the fixed-target experiments and then moving to HERA, the LHeC and finally the FCC-eh, as the centre of mass energy increases, the kinematic reach extends both towards higher Q^2 and smaller x values. At the FCC-eh in particular, it should be able to cover the region down to $x \simeq 10^{-7}$ without leaving the perturbative region $Q \gtrsim 1 \text{ GeV}$. It is important to emphasise that the same coverage would be achieved for nuclear PDFs, extending the coverage by four or five orders of magnitude in x as compared to existing measurements.

One of the most important aspects of the LHeC/FCC-eh scientific case is the capability to probe the proton/nuclear PDFs with an unprecedented precision, not only by means of inclusive structure functions but also with measurements of the strange, charm, and bottom structure functions, that provide a direct handle on the heavy flavour PDFs. The LHeC/FCC-eh would also allow a measurement of the strong coupling constant $\alpha_s(M_Z)$ with per-mille experimental uncertainties, for instance using jet production [231], and high-precision measurements of the electroweak sector parameters. The challenge in this respect will be to provide theoretical predictions with matching uncertainties, which specially for jet production will be difficult. To illustrate these possibilities, in Fig. 81 we also show the results of an xFitter PDF feasibility study that compares the impact on the gluon PDF of adding either LHeC or FCC-eh (or both) pseudo-data in addition to the HERA inclusive structure function dataset. The reduction of the PDF uncertainties down to very small- x values reflects the extended kinematic reach of these future high energy lepton proton colliders. A similar reduction of the PDF uncertainty is expected for the quark PDFs.

Another important aspect of the interplay between PDFs and the LHeC/FCC-eh is related to the small- x resummation framework [170, 171, 592–594]. This framework is based on extending the collinear DGLAP formalism to account for the all-order resummation of terms of the form $\alpha_s^k \ln^m(1/x)$ in the splitting functions and the partonic coefficient functions, as implemented in the BFKL equation [595–597]. Using this approach, the reliability of theoretical predictions for DIS structure functions and collider cross sections can be extended down to much smaller values of x in comparison to the calculations based on the collinear DGLAP framework. Recently, a version of the NNPDF3.1 global analysis, called NNPDF3.1sx, based on NLO+NLL x and NNLO+NLL x theory, has been presented [172], which suggest that accounting for small- x effects leads to an improved description of the inclusive HERA structure function data (see also the discussion in [598, 599]). One may expect such effects to become even more relevant for higher-energy lepton-proton colliders (see Fig. 81).

With the motivation of providing a first estimate of the relevance of small- x resummation for the LHeC/FCC-eh, in Fig. 82 we show predictions for the F_2 and F_L structure functions using the NNPDF3.1sx NNLO and NNLO+NLL x fits at $Q^2 = 5 \text{ GeV}^2$ for the kinematics of the LHeC and the FCC-eh. For these calculations, we have used APFEL to produce NNLO(+NLL x) predictions, each using as input the cor-

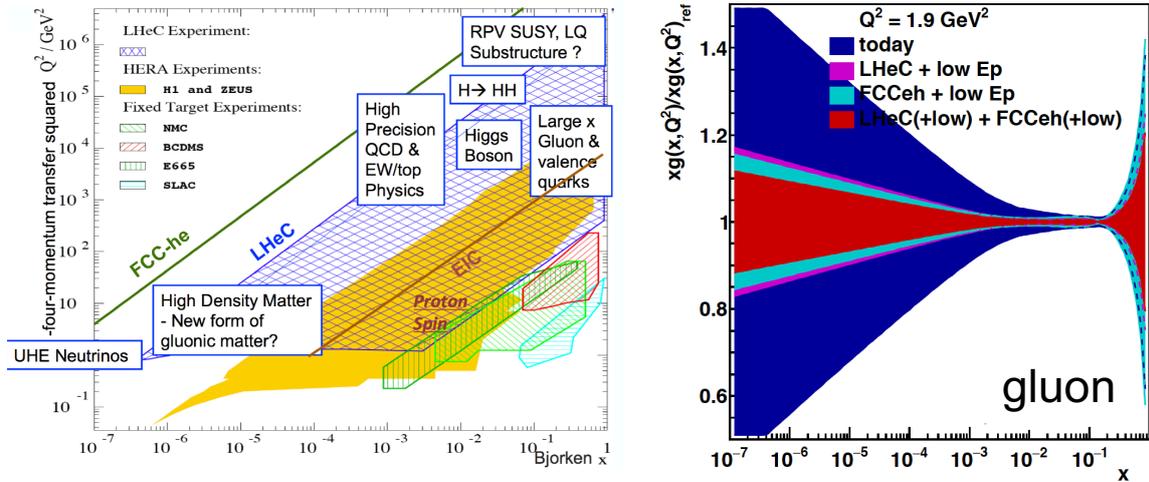


Figure 81: Left: kinematic coverage in the (x, Q^2) of several existing and proposed DIS experiments. Starting from the fixed-target experiments and then moving to HERA, the LHeC and finally the FCC-eh, as the centre of mass energy increases, the kinematic reach extends both towards higher Q^2 and smaller x values. Right: results of an xFitter PDF feasibility study that compares the impact on the gluon PDF of adding either LHeC or FCC-eh (or both) pseudo-data in addition to the HERA inclusive structure function dataset.

responding NNPDF3.1sx fits, for the most updated version of the simulated LHeC/FCC-eh pseudo-data kinematics. In the case of F_2 , we also show the expected total experimental uncertainties based on the simulated pseudo-data, assuming the NNLO+NLL x curve as central prediction. The total uncertainties of the simulated pseudo-data are at the few percent level, and therefore they are rather smaller than the PDF uncertainties in the entire kinematic range. Moreover, the inset in the F_2 plot provides a magnified view of the region $x > 3 \times 10^{-5}$ which is also covered by the HERA measurements.

From the comparisons in Fig. 82 we see how the FCC-eh would allow the small- x region to be probed by about an order of magnitude deeper than the LHeC (which in turn extends HERA by about the same amount). The differences between NNLO and NNLO+NLL x theory are moderate for F_2 , in particular taking into account the large PDF uncertainties. This implies in turn that refitting the pseudo-data is required to first reduce the PDF errors and then discriminate between the two theoretical scenarios. This has been demonstrated explicitly in the study of [172], which highlights how inclusive F_2^p measurements would represent a powerful probe of small- x dynamics. We also see that differences are more marked for F_L , with the central values differing by several times the PDF uncertainty in most of the accessible kinematic range. This illustrates the unique sensitivity of F_L^p measurements to probe small- x QCD dynamics.

In this respect, we emphasize that both structure functions F_2^p and F_L^p provide complementary information on small- x QCD: the former, which can be measured with smaller uncertainties, via scaling violations, and the latter, affected by larger uncertainties, directly via the different overall normalization. We also note that small- x predictions based on non-linear effects [600, 601] (‘saturation’) have typically the opposite trend to small- x resummation, suppressing the structure functions in comparison to the NNLO fixed-order calculation. Therefore, the measurements in Fig. 82 and related ones would open a unique window to the novel dynamical regime of QCD at very small x .

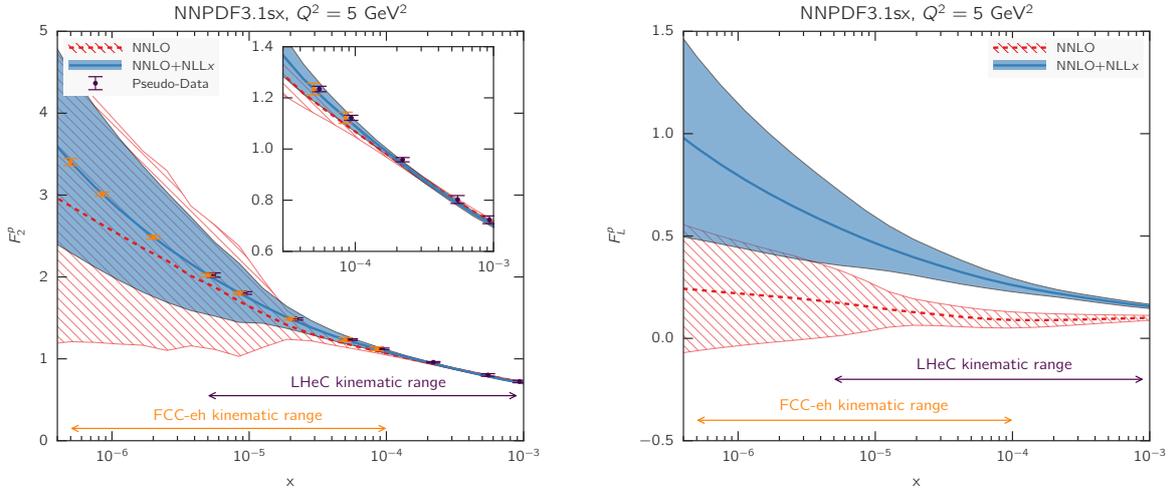


Figure 82: Theoretical predictions for the $F_2(x, Q^2)$ and $F_L(x, Q^2)$ structure functions evaluated with the NNPDF3.1sx NNLO and NNLO+NLL x sets. These structure functions are showed at $Q^2 = 5 \text{ GeV}^2$ for the range in x that would be accessible at the LHeC and FCC-eh. In the case of F_2 , we show in addition the expected total experimental uncertainties with the NNLO+NLL x values as central prediction. The inset in the F_2 plot provides a magnified view of the region $x > 3 \times 10^{-5}$ which is also covered by the HERA measurements.

9.3.2. PDFs at a 100 TeV hadron collider

Next we discuss parton distributions at the FCC-hh, a proposal for a future hadron collider with a centre of mass energy of $\sqrt{s} = 100 \text{ TeV}$. In order to illustrate the extended kinematic coverage that would be achieved at a 100 TeV proton-proton collider in comparison to the one at the LHC, in Fig. 83, taken from [602], we compare the (x, M_X) coverage at 100 TeV and 14 TeV, where the dotted lines indicate the regions of constant rapidity at the FCC-hh. In addition, we also indicate the relevant M_X regions for some representative processes, from low masses (Drell-Yan, low p_T jets), to electroweak scale processes (Higgs, W , Z , top) and new high-mass particles (squarks, Z'). A significant increase in the kinematic coverage relative to the LHC is clear. A particularly interesting aspect is that at the FCC-hh even electroweak-scale processes such as W , Z , or h production become sensitive to the small- x region.

As discussed in the FCC-hh Yellow Report [28], there are two main aspects of PDF phenomenology relevant for a 100 TeV collider. On the one hand, as for the LHC, at $\sqrt{s} = 100 \text{ TeV}$ parton distributions are one of the dominant theoretical uncertainties for several cross sections. In particular, electroweak scale cross sections, such as W or h production, become sensitive to the small- x region where PDF uncertainties are currently large.

To illustrate this point, in Fig. 84 we show the comparison of cross sections for different representative processes at the FCC with $\sqrt{s} = 100 \text{ TeV}$, between the NNPDF3.0 predictions and those of the NNPDF3.0+LHCb sets [27]. Specifically, the NNPDF3.0+LHCb sets used here include D meson production data from LHCb at $\sqrt{s} = 5, 7$ and 13 TeV , as discussed in Sect. 3.8. The acceptance cuts are different in each process, see [603] for more details. In the left plot we show the results for direct photon production, off-peak Drell-Yan cross sections, and inclusive weak boson production. In the right plot we show the fiducial cross sections for $c\bar{c}$ and $b\bar{b}$ production. In all cases, and in particular for heavy quark pair production, we can observe the reduction of PDF uncertainties that is derived once the NNPDF3.0+LHCb sets are used, highlighting the sensitivity to the small- x region.

Kinematics of a 100 TeV FCC

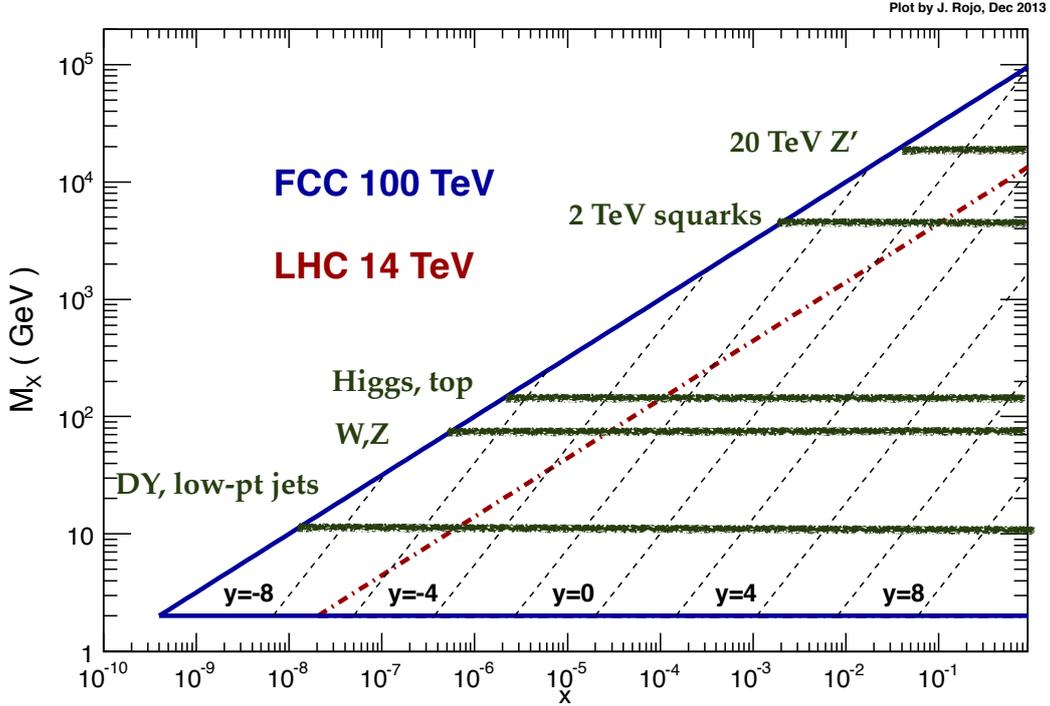


Figure 83: Kinematic coverage in the (x, M_X) plane of a $\sqrt{s} = 100$ TeV hadron collider (solid blue line), compared with the corresponding coverage of the LHC at $\sqrt{s} = 14$ TeV (dot-dashed red line), from [602].

The other aspect of PDFs at the FCC–hh is the onset of new phenomenon that are absent at the lower energies of the LHC. These include the possibility of treating the top quark as a massless parton [604, 605], the need for resummation of “collinear” weak gauge boson radiation and the consequent introduction of electroweak PDFs [606–608], as well as the increased role of photon–induced processes [28, 51, 513]. Moreover, just as in the case of the LHeC/FCC–eh, the role of small– x resummation is expected to become more important at the FCC–hh than at the LHC, given the sensitivity of even standard candles such as W , Z and Higgs production to the small– x region.

Here we provide two representative illustrations of these new PDF–related phenomena at the FCC–hh. First of all, as mentioned above, at 100 TeV the electroweak gauge bosons become effectively massless, and thus it is possible to construct electroweak PDFs with the corresponding evolution equations. In Fig. 85 we show the PDF of the W^+ boson normalized to that of the gluon, as a function of x for different scales: $q = 10^4$ GeV, 10^6 GeV, 10^8 GeV, computed using the framework of [606]. We observe that the dependence of the W PDF with the energy q is rather mild. For most of the range of x , the W PDF is at most a few percent of the gluon PDF, while for $x \geq 0.1$ it becomes larger, up to 40% of the gluon PDF. This does not necessarily mean that the effects of the W PDF will be phenomenologically relevant: this can be assessed only at the cross section level, comparing calculations with massive gauge bosons and those where these are treated as massless (and thus resummed into the electroweak PDFs). Interestingly, the chiral nature of the weak interactions implies that once electroweak effects are accounted for in the DGLAP evolution, the proton can become spontaneously polarized, with implications for BSM searches [609].

Following a similar line of thought, at 100 TeV it is tempting to treat the top quark as a massless

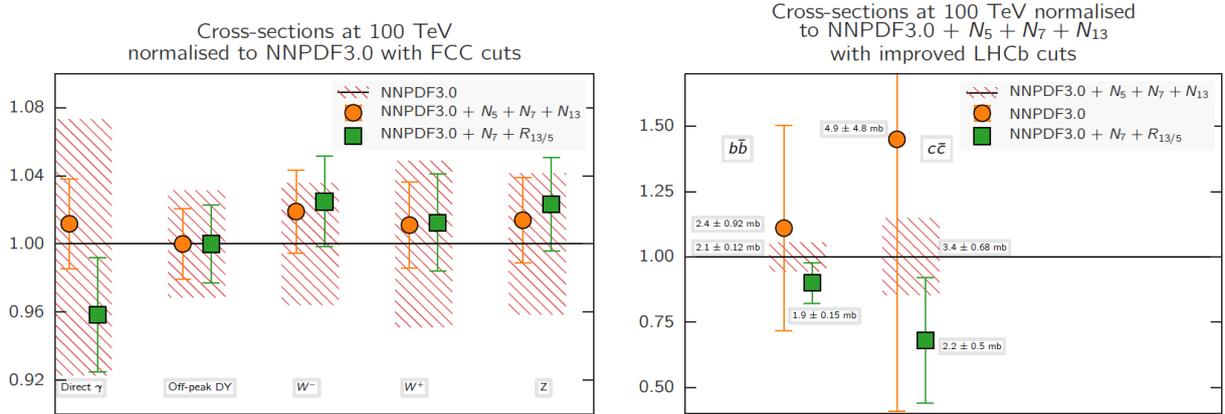


Figure 84: Comparison of cross sections for different representative processes at the FCC with $\sqrt{s} = 100$ TeV, between the NNPDF3.0 predictions and those of the NNPDF3.0+LHCb sets. Note that the acceptance cuts are different in each process. In the left plot we show the results for direct photon production, off–peak Drell–Yan cross sections, and inclusive weak boson production. In the right plot we show the fiducial cross sections for $c\bar{c}$ and $b\bar{b}$ production. Note that the range of the y axis is different in the two plots, reflecting the larger uncertainties in heavy quark production.

parton, much in the same way as at the LHC the bottom quark is treated as massless in most calculations. In Fig. 85 we show the cross section of inclusive heavy scalar (“Higgs”) production by $t\bar{t}$ associated production, comparing the results of the $n_f = 5$ scheme ($gg \rightarrow H^0 t\bar{t}$), the $n_f = 6$ scheme ($t\bar{t} \rightarrow H^0$), and of their interpolation by means of the ACOT general–mass scheme [604] (see also Sect. 2.5). The comparison is performed as a function of the heavy Higgs boson mass m_{H^0} . We find that the $n_f = 6$ calculation, where the top quark is treated as massless and resummed into a top quark PDF $t(x, Q)$, is rather far from the matched calculation up to at least $m_{H^0} = 10$ TeV. This suggest that the massless top approximation is not suitable even for the extreme FCC energies. On the other hand, the concept of top PDF is still useful in order to improve fixed order calculations, using general–mass schemes such as ACOT or FONLL, but this seems to be inadvisable without accounting for the top quark mass effects.

10. Conclusions

The wealth of experimental data collected by the LHC experiments so far represents only a small fraction of the complete dataset that it will provide over the coming two decades of planned operations. This will allow the SM to be studied and stress–tested to an even greater, completely unprecedented, level. Together with recent progress in theoretical calculations, we are therefore entering the era of the LHC precision physics, which will aim to compare data and theory at the percent level or even less. Given the null results of BSM searches so far, a systematic high–precision analysis of the SM predictions and the LHC data might be one of most promising approaches to look for BSM dynamics at the LHC, which may only be manifest as subtle differences with the SM predictions. The detailed mapping of the quark and gluon structure of the proton represents then a crucial component of this LHC precision physics program.

In this Report we have presented an overview of the most important recent developments in PDF determinations, with emphasis on their implications for LHC phenomenology. After a brief review of the theoretical foundations of the global QCD analysis framework, we have reviewed recent progress both from

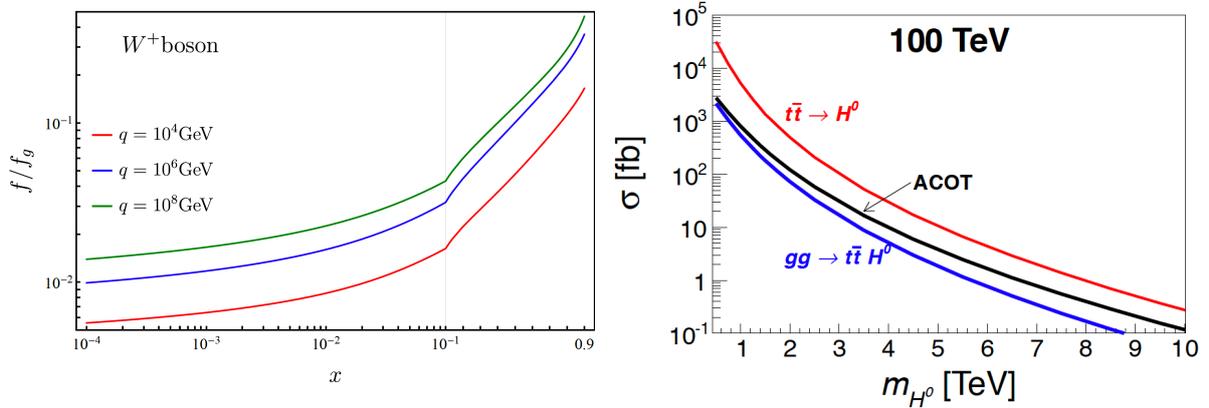


Figure 85: Left: the PDF of the W^+ boson normalized to that of the gluon, as a function of x for different scales: $q = 10^4$ GeV, 10^6 GeV, 10^8 GeV, from Ref. [606]. Right: the cross section of inclusive Higgs production by $t\bar{t}$ associated production, comparing the results of the $n_f = 5$ scheme ($gg \rightarrow ht\bar{t}$), the $n_f = 6$ scheme ($t\bar{t} \rightarrow h$), and of their interpolation by means of the ACOT general-mass scheme, from [604].

the theoretical and the experimental point of view for those hard-scattering cross sections used in PDF fits; we have compared the similarities and differences between the methodologies used for the various PDF fitting collaborations; and then presented the state-of-the-art fits from each group and assessed what we can learn about the internal structure of the nucleons from various points of view. We have then discussed the role of QED corrections in PDF fits, in particular concerning photon-initiated processes, and presented some of the most representative examples of the applications of PDFs for LHC phenomenology, from the measurement of the Higgs boson couplings to the determination of the mass of the W boson.

In the last section of this Report we have discussed some topics that will very likely play a crucial role in the near future of PDF determinations. One of these topics is the role of theoretical uncertainties, in particular those arising from missing higher-order terms in the perturbative expansion. Given the size of PDF uncertainties in the current generation of global analyses, it is likely that these theory errors will increasingly become comparable (if not larger) than the nominal PDF uncertainties, and thus finding a statistically sound method to account for these is of utmost importance. Another topic that might affect the PDF fitting paradigm is that of the interplay with lattice QCD calculations, where recent progress both in computing Mellin moments of various flavour combinations as well in direct x -space calculation of PDFs suggest that in the future the constraints from non-perturbative lattice calculations might be able to provide information for global PDF fits. Finally, we have summarized the importance of PDFs for future higher-energy colliders whose physics case is being discussed at present, such as a Large Hadron electron Collider or a new proton-proton collider with a centre-of-mass energy of up to 100 TeV.

We hope that this Report has managed to convey to the reader that the topic of PDF determinations is an interesting and lively one, with implications from the understanding of the non-perturbative dynamics of the strong interactions to searches for new BSM physics and ultra-high energy astrophysics. Moreover, PDF fits represent a unique internal consistency check of the SM and of the collinear QCD factorization framework. The latter appears to be in remarkably good shape, given that we are now able to simultaneously describe a few tens of individual experiments spanning a wide kinematical domain, some of them with extremely small uncertainties at the per-mile level.

PDF fits thrive at the crossroads of advanced data analysis, state-of-the-art perturbative calculations,

and modern robust statistical methodology, and thus provide guidance for other similar global analyses efforts, such as fits of the Wilson coefficients within the SM Effective Field Theory framework. Moreover, progress in unpolarized PDF fits is also one of the main drivers of recent improvements of other related aspects of the proton structure, from polarized PDFs to nuclear PDFs, as well as to other aspects of non-perturbative QCD such as light and heavy hadron fragmentation.

As we enter in the LHC precision era, ever-improving PDF determinations will keep providing a unique contribution to this exciting exploration of the high-energy frontier. We fully expect that the requirements of this LHC precision program will further drive improvements in global PDF determinations, leading to an ever more detailed picture of the inner life of the proton.

Acknowledgments

We are grateful to Giulia Zanderighi for giving us the opportunity to write this Report. We are also indebted to many colleagues for extensive discussions and fruitful collaborations of the topic of PDF determinations in the last years. A necessarily incomplete list of colleagues that we would like to thank for this includes Alberto Accardi, Sergei Alekhin, Richard D. Ball, Josh Bendavid, Valerio Bertone, Marco Bonvini, Ulla Blumenschein, Stefano Carrazza, Mandy Cooper-Sarkar, David D’Enterria, Sayipjamal Dulat, Stephen Farry, Stefano Forte, Rikkert Frederix, Stefano Frixione, Marco Guzzi, Nathan P. Hartland, Tie-Jiun Hou, Joey Huston, Zahari Kassabov, Valery Khoze, Jan Kretzschmar, Rhorry Gauld, Sasha Glazov, Joey Huston, Katerina Lipka, Jon Pumplin, Michelangelo Mangano, Alan Martin, Simone Marzani, Ronan McNulty, Alexander Mitov, Pavel Nadolsky, Emanuele R. Nocera, Luca Perrozzi, Ringaile Placakyte, Klaus Rabbertz, Voica Radescu, Misha Ryskin, Albert de Roeck, Luca Rottoli, Robert Thorne, Gavin Salam, Carl Schmidt, Peter Skands, Daniel Stump, Maria Ubiali, and C. P. Yuan, among many other colleagues that have helped pushing the boundaries of global PDF fits to a level that would have seemed impossible just a few years ago.

J. R. is supported by the European Research Council (ERC) Starting Grant “PDF4BSM” and by the Dutch Organization for Scientific Research (NWO). The work of J. G. is sponsored by Shanghai Pujiang Program. L. H. L thanks the Science and Technology Facilities Council (STFC) for support via grant award ST/L000377/1.

- [1] J. Rojo, et al., The PDF4LHC report on PDFs and LHC data: Results from Run I and preparation for Run II, *J. Phys. G*42 (2015) 103103. [arXiv:1507.00556](#), [doi:10.1088/0954-3899/42/10/103103](#).
- [2] J. Butterworth, et al., PDF4LHC recommendations for LHC Run II, *J. Phys. G*43 (2016) 023001. [arXiv:1510.03865](#), [doi:10.1088/0954-3899/43/2/023001](#).
- [3] R. D. Ball, S. Carrazza, L. Del Debbio, S. Forte, J. Gao, et al., Parton Distribution Benchmarking with LHC Data, *JHEP* 1304 (2013) 125. [arXiv:1211.5142](#), [doi:10.1007/JHEP04\(2013\)125](#).
- [4] S. Alekhin, et al., The PDF4LHC Working Group Interim Report. [arXiv:1101.0536](#).
- [5] S. Forte, G. Watt, Progress in the Determination of the Partonic Structure of the Proton, *Ann.Rev.Nucl.Part.Sci.* 63 (2013) 291. [arXiv:1301.6754](#), [doi:10.1146/annurev-nucl-102212-170607](#).
- [6] S. Forte, Parton distributions at the dawn of the LHC, *Acta Phys.Polon.* B41 (2010) 2859. [arXiv:1011.5247](#).
- [7] E. Perez, E. Rizvi, The Quark and Gluon Structure of the Proton, *Rep.Prog.Phys.* 76 (2013) 046201. [arXiv:1208.1178](#), [doi:10.1088/0034-4885/76/4/046201](#).
- [8] A. De Roeck, R. S. Thorne, Structure Functions, *Prog.Part.Nucl.Phys.* 66 (2011) 727. [arXiv:1103.0555](#), [doi:10.1016/j.pnpnp.2011.06.001](#).
- [9] A. Accardi, et al., A Critical Appraisal and Evaluation of Modern PDFs, *Eur. Phys. J. C*76 (8) (2016) 471. [arXiv:1603.08906](#), [doi:10.1140/epjc/s10052-016-4285-4](#).

- [10] D. de Florian, et al., Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector. [arXiv:1610.07922](#).
- [11] W. Beenakker, C. Borschensky, M. Krmer, A. Kulesza, E. Laenen, S. Marzani, J. Rojo, NLO+NLL squark and gluino production cross-sections with threshold-improved parton distributions, *Eur. Phys. J. C* 76 (2) (2016) 53. [arXiv:1510.00375](#), [doi:10.1140/epjc/s10052-016-3892-4](#).
- [12] G. Bozzi, L. Citelli, A. Vicini, Parton density function uncertainties on the W boson mass measurement from the lepton transverse momentum distribution, *Phys. Rev. D* 91 (11) (2015) 113005. [arXiv:1501.05587](#), [doi:10.1103/PhysRevD.91.113005](#).
- [13] G. Bozzi, L. Citelli, M. Vesterinen, A. Vicini, Prospects for improving the LHC W boson mass measurement with forward muons, *Eur. Phys. J. C* 75 (12) (2015) 601. [arXiv:1508.06954](#), [doi:10.1140/epjc/s10052-015-3810-1](#).
- [14] G. Bozzi, J. Rojo, A. Vicini, The Impact of PDF uncertainties on the measurement of the W boson mass at the Tevatron and the LHC, *Phys.Rev. D* 83 (2011) 113008. [arXiv:1104.2056](#), [doi:10.1103/PhysRevD.83.113008](#).
- [15] M. Aaboud, et al., Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. [arXiv:1701.07240](#).
- [16] V. Khachatryan, et al., Constraints on parton distribution functions and extraction of the strong coupling constant from the inclusive jet cross section in pp collisions at $\sqrt{s} = 7$ TeV, *Eur. Phys. J. C* 75 (6) (2015) 288. [arXiv:1410.6765](#), [doi:10.1140/epjc/s10052-015-3499-1](#).
- [17] M. Aaboud, et al., Determination of the strong coupling constant α_s from transverse energy-energy correlations in multijet events at $\sqrt{s} = 8$ TeV using the ATLAS detector. [arXiv:1707.02562](#).
- [18] S. Chatrchyan, et al., Measurement of the ratio of the inclusive 3-jet cross section to the inclusive 2-jet cross section in pp collisions at $\sqrt{s} = 7$ TeV and first determination of the strong coupling constant in the TeV range, *Eur. Phys. J. C* 73 (10) (2013) 2604. [arXiv:1304.7498](#), [doi:10.1140/epjc/s10052-013-2604-6](#).
- [19] J. Rojo, Constraints on parton distributions and the strong coupling from LHC jet data, *Int. J. Mod. Phys. A* 30 (2015) 1546005. [arXiv:1410.7728](#), [doi:10.1142/S0217751X15460057](#).
- [20] S. Chatrchyan, et al., Determination of the top-quark pole mass and strong coupling constant from the t t-bar production cross section in pp collisions at $\sqrt{s} = 7$ TeV, *Phys.Lett. B* 728 (2014) 496. [arXiv:1307.1907](#), [doi:10.1016/j.physletb.2014.08.040](#), [10.1016/j.physletb.2013.12.009](#).
- [21] D. Becciolini, M. Gillioz, M. Nardecchia, F. Sannino, M. Spannowsky, Constraining new colored matter from the ratio of 3 to 2 jets cross sections at the LHC, *Phys. Rev. D* 91 (1) (2015) 015010, [Addendum: *Phys. Rev.D*92,no.7,079905(2015)]. [arXiv:1403.7411](#), [doi:10.1103/PhysRevD.91.015010](#), [10.1103/PhysRevD.92.079905](#).
- [22] S. Dimopoulos, S. Raby, F. Wilczek, Supersymmetry and the Scale of Unification, *Phys. Rev. D* 24 (1981) 1681–1683. [doi:10.1103/PhysRevD.24.1681](#).
- [23] A. Cooper-Sarkar, P. Mertsch, S. Sarkar, The high energy neutrino cross-section in the Standard Model and its uncertainty, *JHEP* 08 (2011) 042. [arXiv:1106.3723](#), [doi:10.1007/JHEP08\(2011\)042](#).
- [24] R. Gauld, J. Rojo, L. Rottoli, S. Sarkar, J. Talbert, The prompt atmospheric neutrino flux in the light of LHCb, *JHEP* 02 (2016) 130. [arXiv:1511.06346](#), [doi:10.1007/JHEP02\(2016\)130](#).
- [25] M. V. Garzelli, S. Moch, O. Zenaiev, A. Cooper-Sarkar, A. Geiser, K. Lipka, R. Placakyte, G. Sigl, Prompt neutrino fluxes in the atmosphere with PROSA parton distribution functions, *JHEP* 05 (2017) 004. [arXiv:1611.03815](#), [doi:10.1007/JHEP05\(2017\)004](#).
- [26] O. Zenaiev, et al., Impact of heavy-flavour production cross sections measured by the LHCb experiment on parton distribution functions at low x, *Eur. Phys. J. C* 75 (8) (2015) 396. [arXiv:1503.04581](#), [doi:10.1140/epjc/s10052-015-3618-z](#).
- [27] R. Gauld, J. Rojo, Precision determination of the small-x gluon from charm production at LHCb, *Phys. Rev. Lett.* 118 (7) (2017) 072001. [arXiv:1610.09373](#), [doi:10.1103/PhysRevLett.118.072001](#).
- [28] M. L. Mangano, et al., Physics at a 100 TeV pp collider: Standard Model processes, CERN Yellow Report (3) (2017) 1–254. [arXiv:1607.01831](#), [doi:10.23731/CYRM-2017-003.1](#).
- [29] J. Abelleira Fernandez, et al., A Large Hadron Electron Collider at CERN: Report on the Physics and Design Concepts for Machine and Detector, *J.Phys. G* 39 (2012) 075001. [arXiv:1206.2913](#), [doi:10.1088/0954-3899/39/7/075001](#).
- [30] R. D. Ball, et al., Parton distributions for the LHC Run II, *JHEP* 04 (2015) 040. [arXiv:1410.8849](#), [doi:10.1007/JHEP04\(2015\)040](#).
- [31] S. Dulat, T.-J. Hou, J. Gao, M. Guzzi, J. Huston, P. Nadolsky, J. Pumplin, C. Schmidt, D. Stump, C. P. Yuan, New parton distribution functions from a global analysis of quantum chromodynamics, *Phys. Rev. D* 93 (3) (2016) 033006. [arXiv:1506.07443](#), [doi:10.1103/PhysRevD.93.033006](#).
- [32] L. A. Harland-Lang, A. D. Martin, P. Motylinski, R. S. Thorne, Parton distributions in the LHC era: MMHT 2014 PDFs, *Eur. Phys. J. C* 75 (5) (2015) 204. [arXiv:1412.3989](#), [doi:10.1140/epjc/s10052-015-3397-6](#).
- [33] S. Alekhin, J. Blümlein, S. Moch, R. Placakyte, Parton distribution functions, α_s , and heavy-quark masses for LHC Run II,

- Phys. Rev. D96 (1) (2017) 014011. [arXiv:1701.05838](#), [doi:10.1103/PhysRevD.96.014011](#).
- [34] H. Abramowicz, et al., Combination of measurements of inclusive deep inelastic $e^{\pm}p$ scattering cross sections and QCD analysis of HERA data, Eur. Phys. J. C75 (12) (2015) 580. [arXiv:1506.06042](#), [doi:10.1140/epjc/s10052-015-3710-4](#).
- [35] P. Jimenez-Delgado, E. Reya, Delineating parton distributions and the strong coupling, Phys.Rev. D89 (7) (2014) 074049. [arXiv:1403.1852](#), [doi:10.1103/PhysRevD.89.074049](#).
- [36] A. Accardi, L. T. Brady, W. Melnitchouk, J. F. Owens, N. Sato, Constraints on large- x parton distributions from new weak boson production and deep-inelastic scattering data, Phys. Rev. D93 (11) (2016) 114017. [arXiv:1602.03154](#), [doi:10.1103/PhysRevD.93.114017](#).
- [37] G. Watt, R. S. Thorne, Study of Monte Carlo approach to experimental uncertainty propagation with MSTW 2008 PDFs, JHEP 1208 (2012) 052. [arXiv:1205.4024](#), [doi:10.1007/JHEP08\(2012\)052](#).
- [38] J. Gao, P. Nadolsky, A meta-analysis of parton distribution functions, JHEP 1407 (2014) 035. [arXiv:1401.0013](#), [doi:10.1007/JHEP07\(2014\)035](#).
- [39] S. Carrazza, S. Forte, Z. Kassabov, J. I. Latorre, J. Rojo, An Unbiased Hessian Representation for Monte Carlo PDFs, Eur. Phys. J. C75 (8) (2015) 369. [arXiv:1505.06736](#), [doi:10.1140/epjc/s10052-015-3590-7](#).
- [40] S. Carrazza, J. I. Latorre, J. Rojo, G. Watt, A compression algorithm for the combination of PDF sets, Eur. Phys. J. C75 (2015) 474. [arXiv:1504.06469](#), [doi:10.1140/epjc/s10052-015-3703-3](#).
- [41] M. Czakon, D. Heymes, A. Mitov, Dynamical scales for multi-TeV top-pair production at the LHC, JHEP 04 (2017) 071. [arXiv:1606.03350](#), [doi:10.1007/JHEP04\(2017\)071](#).
- [42] J. Currie, E. W. N. Glover, J. Pires, Next-to-Next-to Leading Order QCD Predictions for Single Jet Inclusive Production at the LHC, Phys. Rev. Lett. 118 (7) (2017) 072002. [arXiv:1611.01460](#), [doi:10.1103/PhysRevLett.118.072002](#).
- [43] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, J. Pires, Precise predictions for dijet production at the LHC, Phys. Rev. Lett. 119 (15) (2017) 152001. [arXiv:1705.10271](#), [doi:10.1103/PhysRevLett.119.152001](#).
- [44] J. M. Campbell, R. K. Ellis, C. Williams, Direct Photon Production at Next-to-Next-to-Leading Order, Phys. Rev. Lett. 118 (22) (2017) 222001. [arXiv:1612.04333](#), [doi:10.1103/PhysRevLett.118.222001](#).
- [45] R. Boughezal, C. Focke, X. Liu, F. Petriello, W -boson production in association with a jet at next-to-next-to-leading order in perturbative QCD, Phys. Rev. Lett. 115 (6) (2015) 062002. [arXiv:1504.02131](#), [doi:10.1103/PhysRevLett.115.062002](#).
- [46] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, T. A. Morgan, The NNLO QCD corrections to Z boson production at large transverse momentum, JHEP 07 (2016) 133. [arXiv:1605.04295](#), [doi:10.1007/JHEP07\(2016\)133](#).
- [47] A. Manohar, P. Nason, G. P. Salam, G. Zanderighi, How bright is the proton? A precise determination of the photon parton distribution function, Phys. Rev. Lett. 117 (24) (2016) 242002. [arXiv:1607.04266](#), [doi:10.1103/PhysRevLett.117.242002](#).
- [48] A. V. Manohar, P. Nason, G. P. Salam, G. Zanderighi, The Photon Content of the Proton. [arXiv:1708.01256](#).
- [49] A. D. Martin, R. G. Roberts, W. J. Stirling, R. S. Thorne, Parton distributions incorporating QED contributions, Eur. Phys. J. C39 (2005) 155. [arXiv:hep-ph/0411040](#), [doi:10.1140/epjc/s2004-02088-7](#).
- [50] C. Schmidt, J. Pumplin, D. Stump, C. P. Yuan, CT14QED parton distribution functions from isolated photon production in deep inelastic scattering, Phys. Rev. D93 (11) (2016) 114015. [arXiv:1509.02905](#), [doi:10.1103/PhysRevD.93.114015](#).
- [51] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, Photon-initiated processes at high mass, Phys. Rev. D94 (7) (2016) 074008. [arXiv:1607.04635](#), [doi:10.1103/PhysRevD.94.074008](#).
- [52] F. Giuli, et al., The photon PDF from high-mass DrellYan data at the LHC, Eur. Phys. J. C77 (6) (2017) 400. [arXiv:1701.08553](#), [doi:10.1140/epjc/s10052-017-4931-5](#).
- [53] R. D. Ball, et al., Parton distributions with QED corrections, Nucl.Phys. B877 (2013) 290–320. [arXiv:1308.0598](#), [doi:10.1016/j.nuclphysb.2013.10.010](#).
- [54] V. Bertone, S. Carrazza, Combining NNPDF3.0 and NNPDF2.3QED through the APFEL evolution code, PoS DIS2016 (2016) 031. [arXiv:1606.07130](#).
- [55] T.-J. Hou, S. Dulat, J. Gao, M. Guzzi, J. Huston, P. Nadolsky, C. Schmidt, J. Winter, K. Xie, C. P. Yuan, CT14 Intrinsic Charm Parton Distribution Functions from CTEQ-TEA Global Analysis. [arXiv:1707.00657](#).
- [56] R. D. Ball, V. Bertone, M. Bonvini, S. Carrazza, S. Forte, A. Guffanti, N. P. Hartland, J. Rojo, L. Rottoli, A Determination of the Charm Content of the Proton, Eur. Phys. J. C76 (11) (2016) 647. [arXiv:1605.06515](#), [doi:10.1140/epjc/s10052-016-4469-y](#).
- [57] S. Alekhin, et al., HERAFitter, Eur. Phys. J. C75 (7) (2015) 304. [arXiv:1410.4412](#), [doi:10.1140/epjc/s10052-015-3480-z](#).
- [58] V. Bertone, S. Carrazza, J. Rojo, APFEL: A PDF Evolution Library with QED corrections, Comput.Phys.Commun. 185 (2014) 1647. [arXiv:1310.1394](#), [doi:10.1016/j.cpc.2014.03.007](#).
- [59] G. P. Salam, J. Rojo, A Higher Order Perturbative Parton Evolution Toolkit (HOPPET), Comput. Phys. Commun. 180 (2009)

- 120–156. [arXiv:0804.3755](#), [doi:10.1016/j.cpc.2008.08.010](#).
- [60] M. Botje, QCDNUM: Fast QCD Evolution and Convolution, *Comput.Phys.Commun.* 182 (2011) 490–532. [arXiv:1005.1481](#), [doi:10.1016/j.cpc.2010.10.020](#).
- [61] 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](#), [doi:10.1140/epjc/s10052-010-1255-0](#).
- [62] V. Bertone, S. Carrazza, N. P. Hartland, APFELgrid: a high performance tool for parton density determinations, *Comput. Phys. Commun.* 212 (2017) 205–209. [arXiv:1605.02070](#), [doi:10.1016/j.cpc.2016.10.006](#).
- [63] V. Bertone, R. Frederix, S. Frixione, J. Rojo, M. Sutton, aMCfast: automation of fast NLO computations for PDF fits, *JHEP* 1408 (2014) 166. [arXiv:1406.7693](#), [doi:10.1007/JHEP08\(2014\)166](#).
- [64] M. Wobisch, D. Britzger, T. Kluge, K. Rabbertz, F. Stober, Theory-Data Comparisons for Jet Measurements in Hadron-Induced Processes. [arXiv:1109.1310](#).
- [65] M. Czakon, D. Heymes, A. Mitov, fastNLO tables for NNLO top-quark pair differential distributions. [arXiv:1704.08551](#).
- [66] E. R. Nocera, R. D. Ball, S. Forte, G. Ridolfi, J. Rojo, A first unbiased global determination of polarized PDFs and their uncertainties, *Nucl.Phys. B* 887 (2014) 276. [arXiv:1406.5539](#), [doi:10.1016/j.nuclphysb.2014.08.008](#).
- [67] D. de Florian, R. Sassot, M. Stratmann, W. Vogelsang, Extraction of Spin-Dependent Parton Densities and Their Uncertainties, *Phys. Rev. D* 80 (2009) 034030. [arXiv:0904.3821](#), [doi:10.1103/PhysRevD.80.034030](#).
- [68] N. Sato, W. Melnitchouk, S. E. Kuhn, J. J. Ethier, A. Accardi, Iterative Monte Carlo analysis of spin-dependent parton distributions, *Phys. Rev. D* 93 (7) (2016) 074005. [arXiv:1601.07782](#), [doi:10.1103/PhysRevD.93.074005](#).
- [69] D. de Florian, R. Sassot, P. Zurita, M. Stratmann, Global Analysis of Nuclear Parton Distributions, *Phys. Rev. D* 85 (2012) 074028. [arXiv:1112.6324](#), [doi:10.1103/PhysRevD.85.074028](#).
- [70] K. J. Eskola, P. Paakkinen, H. Paukkunen, C. A. Salgado, EPPS16: Nuclear parton distributions with LHC data, *Eur. Phys. J. C* 77 (3) (2017) 163. [arXiv:1612.05741](#), [doi:10.1140/epjc/s10052-017-4725-9](#).
- [71] A. Kusina, F. Lyonnet, D. B. Clark, E. Godat, T. Jezo, K. Kovarik, F. I. Olness, I. Schienbein, J. Y. Yu, Vector boson production in pPb and PbPb collisions at the LHC and its impact on nCTEQ15 PDFs, *Eur. Phys. J. C* 77 (7) (2017) 488. [arXiv:1610.02925](#), [doi:10.1140/epjc/s10052-017-5036-x](#).
- [72] R. Angeles-Martinez, et al., Transverse Momentum Dependent (TMD) parton distribution functions: status and prospects, *Acta Phys. Polon. B* 46 (12) (2015) 2501–2534. [arXiv:1507.05267](#), [doi:10.5506/APhysPolB.46.2501](#).
- [73] R. Hofstadter, R. W. McAllister, Electron Scattering From the Proton, *Phys. Rev.* 98 (1955) 217–218. [doi:10.1103/PhysRev.98.217](#).
- [74] R. W. McAllister, R. Hofstadter, Elastic Scattering of 188-MeV Electrons From the Proton and the α Particle, *Phys. Rev.* 102 (1956) 851–856. [doi:10.1103/PhysRev.102.851](#).
- [75] G. Zweig, [An SU\(3\) model for strong interaction symmetry and its breaking. Version 2](#), in: D. Lichtenberg, S. P. Rosen (Eds.), *DEVELOPMENTS IN THE QUARK THEORY OF HADRONS. VOL. 1. 1964 - 1978, 1964*, pp. 22–101. URL <https://inspirehep.net/record/4674/files/cern-th-412.pdf>
- [76] M. Gell-Mann, A Schematic Model of Baryons and Mesons, *Phys. Lett.* 8 (1964) 214–215. [doi:10.1016/S0031-9163\(64\)92001-3](#).
- [77] E. D. Bloom, et al., High-Energy Inelastic e p Scattering at 6-Degrees and 10-Degrees, *Phys. Rev. Lett.* 23 (1969) 930–934. [doi:10.1103/PhysRevLett.23.930](#).
- [78] M. Breidenbach, J. I. Friedman, H. W. Kendall, E. D. Bloom, D. H. Coward, H. C. DeStaebler, J. Drees, L. W. Mo, R. E. Taylor, Observed Behavior of Highly Inelastic electron-Proton Scattering, *Phys. Rev. Lett.* 23 (1969) 935–939. [doi:10.1103/PhysRevLett.23.935](#).
- [79] J. D. Bjorken, Asymptotic Sum Rules at Infinite Momentum, *Phys. Rev.* 179 (1969) 1547–1553. [doi:10.1103/PhysRev.179.1547](#).
- [80] W. Albrecht, F. W. Brasse, H. Dorner, W. Flauger, K. H. Frank, J. Gayler, H. Hultschig, J. May, E. Ganssauge, Inelastic electron-proton scattering at fixed four-momentum transfer of 0.773 and 1.935 (gev/c)-squared, *Nucl. Phys. B* 13 (1969) 1–8. [doi:10.1016/0550-3213\(69\)90359-9](#).
- [81] R. P. Feynman, Very high-energy collisions of hadrons, *Phys. Rev. Lett.* 23 (1969) 1415–1417. [doi:10.1103/PhysRevLett.23.1415](#).
- [82] J. D. Bjorken, E. A. Paschos, Inelastic Electron Proton and gamma Proton Scattering, and the Structure of the Nucleon, *Phys. Rev.* 185 (1969) 1975–1982. [doi:10.1103/PhysRev.185.1975](#).
- [83] C. G. Callan, Jr., D. J. Gross, High-energy electroproduction and the constitution of the electric current, *Phys. Rev. Lett.* 22 (1969) 156–159. [doi:10.1103/PhysRevLett.22.156](#).
- [84] G. Miller, et al., Inelastic electron-Proton Scattering at Large Momentum Transfers, *Phys. Rev. D* 5 (1972) 528. [doi:10.1103/PhysRevD.5.528](#).
- [85] D. J. Gross, F. Wilczek, Ultraviolet Behavior of Nonabelian Gauge Theories, *Phys. Rev. Lett.* 30 (1973) 1343–1346. [doi:10.1103/PhysRevLett.30.1343](#).

- [86] H. D. Politzer, Reliable Perturbative Results for Strong Interactions?, Phys. Rev. Lett. 30 (1973) 1346–1349. doi:10.1103/PhysRevLett.30.1346.
- [87] J. Kuti, V. F. Weisskopf, Inelastic lepton - nucleon scattering and lepton pair production in the relativistic quark parton model, Phys. Rev. D4 (1971) 3418–3439. doi:10.1103/PhysRevD.4.3418.
- [88] R. McElhaney, S. F. Tuan, Some consequences of a modified Kuti Weisskopf quark parton model, Phys. Rev. D8 (1973) 2267–2272. doi:10.1103/PhysRevD.8.2267.
- [89] V. D. Barger, R. J. N. Phillips, Quark - parton model relations in deep inelastic lepton scattering, Nucl. Phys. B73 (1974) 269–294. doi:10.1016/0550-3213(74)90020-0.
- [90] I. Hinchliffe, C. H. Llewellyn Smith, Detailed Treatment of Scaling Violations in Asymptotically Free Gauge Theories, Nucl. Phys. B128 (1977) 93. doi:10.1016/0550-3213(77)90302-9.
- [91] D. J. Fox, et al., Test of Scale Invariance in High-Energy Muon Scattering, Phys. Rev. Lett. 33 (1974) 1504. doi:10.1103/PhysRevLett.33.1504.
- [92] G. Altarelli, G. Parisi, Asymptotic Freedom in Parton Language, Nucl. Phys. B126 (1977) 298–318. doi:10.1016/0550-3213(77)90384-4.
- [93] V. N. Gribov, L. N. Lipatov, Deep inelastic ep scattering in perturbation theory, Sov. J. Nucl. Phys. 15 (1972) 438–450.
- [94] Y. L. Dokshitzer, Calculation of the structure functions for deep inelastic scattering and e^+e^- annihilation by perturbation theory in quantum chromodynamics. (in russian), Sov. Phys. JETP 46 (1977) 641–653.
- [95] L. N. Lipatov, The parton model and perturbation theory, Sov. J. Nucl. Phys. 20 (1975) 94–102, [Yad. Fiz.20,181(1974)].
- [96] M. Gluck, E. Hoffmann, E. Reya, Scaling Violations and the Gluon Distribution of the Nucleon, Zeit. Phys. C13 (1982) 119. doi:10.1007/BF01547675.
- [97] L. Baulieu, C. Kounnas, A Direct Method for Computing QCD Predictions for Deep Inelastic Structure Functions, Nucl. Phys. B155 (1979) 429–446. doi:10.1016/0550-3213(79)90279-7.
- [98] E. Eichten, I. Hinchliffe, K. D. Lane, C. Quigg, Super Collider Physics, Rev. Mod. Phys. 56 (1984) 579–707. doi:10.1103/RevModPhys.56.579.
- [99] H. Abramowicz, et al., Neutrino and anti-neutrinos Charged Current Inclusive Scattering in Iron in the Energy Range 20-GeV < Neutrino Energy < 300-GeV, Z. Phys. C17 (1983) 283. doi:10.1007/BF01571895.
- [100] D. W. Duke, J. F. Owens, q^2 Dependent Parametrizations of Parton Distribution Functions, Phys. Rev. D30 (1984) 49–54. doi:10.1103/PhysRevD.30.49.
- [101] A. D. Martin, R. G. Roberts, W. J. Stirling, Structure Function Analysis and psi, Jet, W, Z Production: Pinning Down the Gluon, Phys. Rev. D37 (1988) 1161. doi:10.1103/PhysRevD.37.1161.
- [102] M. Diemoz, F. Ferroni, E. Longo, G. Martinelli, Parton Densities from Deep Inelastic Scattering to Hadronic Processes at Super Collider Energies, Z. Phys. C39 (1988) 21. doi:10.1007/BF01560387.
- [103] P. Aurenche, R. Baier, M. Fontannaz, J. F. Owens, M. Werlen, The Gluon Contents of the Nucleon Probed with Real and Virtual Photons, Phys. Rev. D39 (1989) 3275. doi:10.1103/PhysRevD.39.3275.
- [104] P. N. Harriman, A. D. Martin, W. J. Stirling, R. G. Roberts, Parton Distributions Extracted From Data on Deep Inelastic Lepton Scattering, Prompt Photon Production and the Drell-Yan Process, Phys. Rev. D42 (1990) 798–810. doi:10.1103/PhysRevD.42.798.
- [105] J. G. Morfin, W.-K. Tung, Parton distributions from a global QCD analysis of deep inelastic scattering and lepton pair production, Z. Phys. C52 (1991) 13–30. doi:10.1007/BF01412323.
- [106] M. Gluck, E. Reya, A. Vogt, Radiatively generated parton distributions for high-energy collisions, Z. Phys. C48 (1990) 471–482. doi:10.1007/BF01572029.
- [107] A. D. Martin, W. J. Stirling, R. G. Roberts, Parton distributions of the proton, Phys. Rev. D50 (1994) 6734–6752. arXiv:hep-ph/9406315, doi:10.1103/PhysRevD.50.6734.
- [108] H. L. Lai, J. Botts, J. Huston, J. G. Morfin, J. F. Owens, J.-w. Qiu, W. K. Tung, H. Weerts, Global QCD analysis and the CTEQ parton distributions, Phys. Rev. D51 (1995) 4763–4782. arXiv:hep-ph/9410404, doi:10.1103/PhysRevD.51.4763.
- [109] E. L. Berger, F. Halzen, C. S. Kim, S. Willenbrock, WEAK BOSON PRODUCTION AT TEVATRON ENERGIES, Phys. Rev. D40 (1989) 83, [Erratum: Phys. Rev.D40,3789(1989)]. doi:10.1103/PhysRevD.40.83, 10.1103/PhysRevD.40.3789.
- [110] F. Abe, et al., Inclusive jet cross section in $\bar{p}p$ collisions at $\sqrt{s} = 1.8$ TeV, Phys. Rev. Lett. 77 (1996) 438–443. arXiv:hep-ex/9601008, doi:10.1103/PhysRevLett.77.438.
- [111] J. Huston, E. Kovacs, S. Kuhlmann, H. L. Lai, J. F. Owens, D. E. Soper, W. K. Tung, Large transverse momentum jet production and the gluon distribution inside the proton, Phys. Rev. Lett. 77 (1996) 444–447. arXiv:hep-ph/9511386, doi:10.1103/PhysRevLett.77.444.
- [112] S. Alekhin, Extraction of parton distributions and α_s from DIS data within the Bayesian treatment of systematic errors, Eur. Phys. J. C10 (1999) 395–403. arXiv:hep-ph/9611213, doi:10.1007/s100520050763.

- [113] M. Botje, A QCD analysis of HERA and fixed target structure function data, Eur. Phys. J. C14 (2000) 285–297. [arXiv:hep-ph/9912439](#), [doi:10.1007/s100520000358](#).
- [114] V. Barone, C. Pascaud, F. Zomer, A new global analysis of deep inelastic scattering data, Eur. Phys. J. C12 (2000) 243–262. [arXiv:hep-ph/9907512](#), [doi:10.1007/s100529900198](#).
- [115] W. T. Giele, S. A. Keller, D. A. Kosower, Parton distribution function uncertainties. [arXiv:hep-ph/0104052](#).
- [116] J. Pumplin, et al., Uncertainties of predictions from parton distribution functions. 2. The Hessian method, Phys. Rev. D65 (2001) 014013. [arXiv:hep-ph/0101032](#), [doi:10.1103/PhysRevD.65.014013](#).
- [117] J. Pumplin, et al., New generation of parton distributions with uncertainties from global QCD analysis, JHEP 07 (2002) 012. [arXiv:hep-ph/0201195](#).
- [118] A. D. Martin, R. G. Roberts, W. J. Stirling, R. S. Thorne, Uncertainties of predictions from parton distributions. I: Experimental errors. ((T)), Eur. Phys. J. C28 (2003) 455–473. [arXiv:hep-ph/0211080](#), [doi:10.1140/epjc/s2003-01196-2](#).
- [119] M. A. G. Aivazis, J. C. Collins, F. I. Olness, W.-K. Tung, Leptoproduction of heavy quarks. 2. a unified qcd formulation of charged and neutral current processes from fixed target to collider energies, Phys. Rev. D50 (1994) 3102–3118. [arXiv:hep-ph/9312319](#).
- [120] R. S. Thorne, R. G. Roberts, An Ordered analysis of heavy flavor production in deep inelastic scattering, Phys. Rev. D57 (1998) 6871–6898. [arXiv:hep-ph/9709442](#), [doi:10.1103/PhysRevD.57.6871](#).
- [121] J. C. Collins, Hard scattering factorization with heavy quarks: A General treatment, Phys. Rev. D58 (1998) 094002. [arXiv:hep-ph/9806259](#), [doi:10.1103/PhysRevD.58.094002](#).
- [122] S. Moch, J. A. M. Vermaseren, A. Vogt, The three-loop splitting functions in qcd: The non-singlet case, Nucl. Phys. B688 (2004) 101–134. [arXiv:hep-ph/0403192](#).
- [123] A. Vogt, S. Moch, J. A. M. Vermaseren, The Three-loop splitting functions in QCD: The Singlet case, Nucl. Phys. B691 (2004) 129–181. [arXiv:hep-ph/0404111](#), [doi:10.1016/j.nuclphysb.2004.04.024](#).
- [124] A. D. Martin, W. J. Stirling, R. S. Thorne, G. Watt, Parton distributions for the LHC, Eur. Phys. J. C63 (2009) 189. [arXiv:0901.0002](#), [doi:10.1140/epjc/s10052-009-1072-5](#).
- [125] J. Gao, et al., CT10 next-to-next-to-leading order global analysis of QCD, Phys.Rev. D89 (2014) 033009. [arXiv:1302.6246](#), [doi:10.1103/PhysRevD.89.033009](#).
- [126] S. Alekhin, J. Blümlein, S. Klein, S. Moch, The 3-, 4-, and 5-flavor NNLO Parton from Deep-Inelastic- Scattering Data and at Hadron Colliders, Phys. Rev. D81 (2010) 014032. [arXiv:0908.2766](#), [doi:10.1103/PhysRevD.81.014032](#).
- [127] S. I. Alekhin, Global fit to the charged leptons DIS data: α_s , parton distributions, and high twists, Phys. Rev. D63 (2001) 094022. [arXiv:hep-ph/0011002](#), [doi:10.1103/PhysRevD.63.094022](#).
- [128] S. Alekhin, Parton distribution functions from the precise NNLO QCD fit, JETP Lett. 82 (2005) 628–631. [arXiv:hep-ph/0508248](#).
- [129] F. Aaron, et al., Combined Measurement and QCD Analysis of the Inclusive e^+p Scattering Cross Sections at HERA, JHEP 1001 (2010) 109. [arXiv:0911.0884](#), [doi:10.1007/JHEP01\(2010\)109](#).
- [130] P. Jimenez-Delgado, E. Reya, Dynamical NNLO parton distributions, Phys.Rev. D79 (2009) 074023. [arXiv:0810.4274](#), [doi:10.1103/PhysRevD.79.074023](#).
- [131] S. Forte, L. Garrido, J. I. Latorre, A. Piccione, Neural network parametrization of deep-inelastic structure functions, JHEP 05 (2002) 062. [arXiv:hep-ph/0204232](#).
- [132] R. D. Ball, et al., A first unbiased global NLO determination of parton distributions and their uncertainties, Nucl. Phys. B838 (2010) 136. [arXiv:1002.4407](#), [doi:10.1016/j.nuclphysb.2010.05.008](#).
- [133] R. D. Ball, et al., Impact of Heavy Quark Masses on Parton Distributions and LHC Phenomenology, Nucl. Phys. B849 (2011) 296. [arXiv:1101.1300](#).
- [134] R. D. Ball, et al., Unbiased global determination of parton distributions and their uncertainties at NNLO and at LO, Nucl.Phys. B855 (2012) 153. [arXiv:1107.2652](#).
- [135] R. D. Ball, et al., Parton distributions with LHC data, Nucl.Phys. B867 (2013) 244. [arXiv:1207.1303](#), [doi:10.1016/j.nuclphysb.2012.10.003](#).
- [136] S. Alekhin, J. Blümlein, S. Moch, The ABM parton distributions tuned to LHC data, Phys.Rev. D89 (2014) 054028. [arXiv:1310.3059](#), [doi:10.1103/PhysRevD.89.054028](#).
- [137] M. Botje, et al., The PDF4LHC Working Group Interim Recommendations. [arXiv:1101.0538](#).
- [138] J. C. Collins, D. E. Soper, The Theorems of Perturbative QCD, Ann. Rev. Nucl. Part. Sci. 37 (1987) 383–409.
- [139] J. C. Collins, D. E. Soper, G. F. Sterman, Factorization of Hard Processes in QCD, Adv. Ser. Direct. High Energy Phys. 5 (1989) 1–91. [arXiv:hep-ph/0409313](#).
- [140] J. M. Campbell, J. W. Huston, W. J. Stirling, Hard interactions of quarks and gluons: A primer for LHC physics, Rept. Prog. Phys. 70 (2007) 89. [arXiv:hep-ph/0611148](#), [doi:10.1088/0034-4885/70/1/R02](#).
- [141] R. Hamberg, W. L. van Neerven, T. Matsuura, A complete calculation of the order α_s^2 correction to the Drell-Yan K factor, Nucl. Phys. B359 (1991) 343–405, [Erratum: Nucl. Phys.B644,403(2002)]. [doi:10.1016/S0550-3213\(02\)00814-3](#),

- 10.1016/0550-3213(91)90064-5.
- [142] M. Czakon, P. Fiedler, A. Mitov, The total top quark pair production cross-section at hadron colliders through $O(\alpha_s^4)$, Phys.Rev.Lett. 110 (2013) 252004. [arXiv:1303.6254](#), [doi:10.1103/PhysRevLett.110.252004](#).
 - [143] C. Anastasiou, C. Duhr, F. Dulat, F. Herzog, B. Mistlberger, Higgs Boson Gluon-Fusion Production in QCD at Three Loops, Phys. Rev. Lett. 114 (21) (2015) 212001. [arXiv:1503.06056](#), [doi:10.1103/PhysRevLett.114.212001](#).
 - [144] J. Ablinger, A. Behring, J. Blümlein, A. De Freitas, A. von Manteuffel, C. Schneider, The three-loop splitting functions $P_{gg}^{(2)}$ and $P_{gg}^{(2,N_F)}$, Nucl. Phys. B922 (2017) 1–40. [arXiv:1705.01508](#), [doi:10.1016/j.nuclphysb.2017.06.004](#).
 - [145] S. Moch, B. Ruijl, T. Ueda, J. A. M. Vermaseren, A. Vogt, Four-Loop Non-Singlet Splitting Functions in the Planar Limit and Beyond. [arXiv:1707.08315](#).
 - [146] A. Vogt, Efficient evolution of unpolarized and polarized parton distributions with qcd-pegasus, Comput. Phys. Commun. 170 (2005) 65–92. [arXiv:hep-ph/0408244](#).
 - [147] M. R. Whalley, D. Bourilkov, R. C. Group, The Les Houches accord PDFs (LHAPDF) and LHAGLUE, in: HERA and the LHC: A Workshop on the implications of HERA for LHC physics. Proceedings, Part B, 2005, pp. 575–581. [arXiv:hep-ph/0508110](#).
 - [148] M. Dittmar, et al., Parton Distributions. [arXiv:0901.2504](#).
 - [149] S. Forte, E. Laenen, P. Nason, J. Rojo, Heavy quarks in deep-inelastic scattering, Nucl. Phys. B834 (2010) 116–162. [arXiv:1001.2312](#), [doi:10.1016/j.nuclphysb.2010.03.014](#).
 - [150] M. Kramer, I. F. I. Olness, D. E. Soper, Treatment of heavy quarks in deeply inelastic scattering, Phys. Rev. D62 (2000) 096007. [arXiv:hep-ph/0003035](#), [doi:10.1103/PhysRevD.62.096007](#).
 - [151] R. S. Thorne, A variable-flavour number scheme for nnlo, Phys. Rev. D73 (2006) 054019. [arXiv:hep-ph/0601245](#).
 - [152] J. Rojo, et al., Chapter 22 in: J. R. Andersen et al., "The SM and NLO multileg working group: Summary report", [arXiv:1003.1241](#) (2010).
 - [153] J. Butterworth, et al., Les Houches 2013: Physics at TeV Colliders: Standard Model Working Group Report. [arXiv:1405.1067](#).
 - [154] J. C. Collins, F. Wilczek, A. Zee, Low-Energy Manifestations of Heavy Particles: Application to the Neutral Current, Phys. Rev. D18 (1978) 242. [doi:10.1103/PhysRevD.18.242](#).
 - [155] E. Laenen, S. Riemersma, J. Smith, W. L. van Neerven, Complete $O(\alpha_s)$ corrections to heavy flavor structure functions in electroproduction, Nucl. Phys. B392 (1993) 162–228. [doi:10.1016/0550-3213\(93\)90201-Y](#).
 - [156] E. Laenen, S. Riemersma, J. Smith, W. L. van Neerven, $O(\alpha_s)$ corrections to heavy flavor inclusive distributions in electroproduction, Nucl. Phys. B392 (1993) 229–250. [doi:10.1016/0550-3213\(93\)90202-Z](#).
 - [157] E. L. Berger, J. Gao, C. S. Li, Z. L. Liu, H. X. Zhu, Charm-Quark Production in Deep-Inelastic Neutrino Scattering at Next-to-Next-to-Leading Order in QCD, Phys. Rev. Lett. 116 (21) (2016) 212002. [arXiv:1601.05430](#), [doi:10.1103/PhysRevLett.116.212002](#).
 - [158] J. Gao, Massive charged-current coefficient functions in deep-inelastic scattering at NNLO and impact on strange-quark distributions [arXiv:1710.04258](#).
 - [159] J. Ablinger, J. Blmlein, A. De Freitas, A. Hasselhuhn, C. Schneider, F. Wibrock, Three Loop Massive Operator Matrix Elements and Asymptotic Wilson Coefficients with Two Different Masses, Nucl. Phys. B921 (2017) 585–688. [arXiv:1705.07030](#), [doi:10.1016/j.nuclphysb.2017.05.017](#).
 - [160] M. Cacciari, M. Greco, P. Nason, The p(T) spectrum in heavy-flavour hadroproduction, JHEP 05 (1998) 007. [arXiv:hep-ph/9803400](#).
 - [161] P. M. Nadolsky, W.-K. Tung, Improved Formulation of Global QCD Analysis with Zero-mass Matrix Elements, Phys. Rev. D79 (2009) 113014. [arXiv:0903.2667](#), [doi:10.1103/PhysRevD.79.113014](#).
 - [162] A. Chuvakin, J. Smith, W. L. van Neerven, Comparison between variable flavor number schemes for charm quark electroproduction, Phys. Rev. D61 (2000) 096004. [arXiv:hep-ph/9910250](#), [doi:10.1103/PhysRevD.61.096004](#).
 - [163] R. D. Ball, V. Bertone, M. Bonvini, S. Forte, P. Groth Merrild, J. Rojo, L. Rottoli, Intrinsic charm in a matched general-mass scheme, Phys. Lett. B754 (2016) 49–58. [arXiv:1510.00009](#), [doi:10.1016/j.physletb.2015.12.077](#).
 - [164] R. D. Ball, M. Bonvini, L. Rottoli, Charm in Deep-Inelastic Scattering, JHEP 11 (2015) 122. [arXiv:1510.02491](#), [doi:10.1007/JHEP11\(2015\)122](#).
 - [165] R. Thorne, The effect on PDFs and $\alpha_S(M_Z^2)$ due to changes in flavour scheme and higher twist contributions, Eur.Phys.J. C74 (7) (2014) 2958. [arXiv:1402.3536](#), [doi:10.1140/epjc/s10052-014-2958-4](#).
 - [166] R. S. Thorne, A Variable-flavor number scheme for NNLO, Phys.Rev. D73 (2006) 054019. [arXiv:hep-ph/0601245](#), [doi:10.1103/PhysRevD.73.054019](#).
 - [167] M. Guzzi, P. M. Nadolsky, H.-L. Lai, C.-P. Yuan, General-Mass Treatment for Deep Inelastic Scattering at Two-Loop Accuracy, Phys.Rev. D86 (2012) 053005. [arXiv:1108.5112](#), [doi:10.1103/PhysRevD.86.053005](#).
 - [168] S. A. Malik, G. Watt, Ratios of W and Z cross sections at large boson p_T as a constraint on PDFs and background to new physics, JHEP 1402 (2014) 025. [arXiv:1304.2424](#), [doi:10.1007/JHEP02\(2014\)025](#).

- [169] R. D. Ball, et al., Parton distributions from high-precision collider data. [arXiv:1706.00428](#).
- [170] C. D. White, R. S. Thorne, A Global Fit to Scattering Data with NLL BFKL Resummations, *Phys. Rev. D* 75 (2007) 034005. [arXiv:hep-ph/0611204](#), [doi:10.1103/PhysRevD.75.034005](#).
- [171] M. Bonvini, S. Marzani, C. Muselli, Towards parton distribution functions with small- x resummation: HELL 2.0. [arXiv:1708.07510](#).
- [172] R. D. Ball, V. Bertone, M. Bonvini, S. Marzani, J. Rojo, L. Rottoli, Parton distributions with small- x resummation: evidence for BFKL dynamics in HERA data [arXiv:1710.05935](#).
- [173] M. Bonvini, S. Marzani, J. Rojo, L. Rottoli, M. Ubiali, R. D. Ball, V. Bertone, S. Carrazza, N. P. Hartland, Parton distributions with threshold resummation, *JHEP* 09 (2015) 191. [arXiv:1507.01006](#), [doi:10.1007/JHEP09\(2015\)191](#).
- [174] G. Altarelli, G. Martinelli, Transverse Momentum of Jets in Electroproduction from Quantum Chromodynamics, *Phys. Lett.* 76B (1978) 89–94. [doi:10.1016/0370-2693\(78\)90109-0](#).
- [175] V. Bertone, et al., A determination of $m_c(m_c)$ from HERA data using a matched heavy-flavor scheme, *JHEP* 08 (2016) 050. [arXiv:1605.01946](#), [doi:10.1007/JHEP08\(2016\)050](#).
- [176] J. Gao, M. Guzzi, P. M. Nadolsky, Charm quark mass dependence in a global QCD analysis, *Eur.Phys.J. C* 73 (2013) 2541. [arXiv:1304.3494](#), [doi:10.1140/epjc/s10052-013-2541-4](#).
- [177] S. Alekhin, J. Blümlein, K. Daum, K. Lipka, S. Moch, Precise charm-quark mass from deep-inelastic scattering, *Phys.Lett.* B720 (2013) 172. [arXiv:1212.2355](#), [doi:10.1016/j.physletb.2013.02.010](#).
- [178] A. C. Benvenuti, et al., A High Statistics Measurement of the Proton Structure Functions $F_2(x, Q^2)$ and R from Deep Inelastic Muon Scattering at High Q^2 , *Phys. Lett.* B223 (1989) 485. [doi:10.1016/0370-2693\(89\)91637-7](#).
- [179] A. C. Benvenuti, et al., A High Statistics Measurement of the Deuteron Structure Functions $F_2(x, Q^2)$ and R from Deep Inelastic Muon Scattering at High Q^2 , *Phys. Lett.* B237 (1990) 592. [doi:10.1016/0370-2693\(90\)91231-Y](#).
- [180] M. Arneodo, et al., Accurate measurement of F_2^d/F_2^p and $R_d - R_p$, *Nucl. Phys.* B487 (1997) 3–26. [arXiv:hep-ex/9611022](#), [doi:10.1016/S0550-3213\(96\)00673-6](#).
- [181] M. Arneodo, et al., Measurement of the proton and deuteron structure functions, F_2^p and F_2^d , and of the ratio σ_L/σ_T , *Nucl. Phys.* B483 (1997) 3–43. [arXiv:hep-ph/9610231](#), [doi:10.1016/S0550-3213\(96\)00538-X](#).
- [182] L. W. Whitlow, E. M. Riordan, S. Dasu, S. Rock, A. Bodek, Precise measurements of the proton and deuteron structure functions from a global analysis of the SLAC deep inelastic electron scattering cross-sections, *Phys. Lett.* B282 (1992) 475–482. [doi:10.1016/0370-2693\(92\)90672-Q](#).
- [183] S. Tkachenko, et al., Measurement of the structure function of the nearly free neutron using spectator tagging in inelastic $^2\text{H}(e, e'p)X$ scattering with CLAS, *Phys. Rev.* C89 (2014) 045206, [Addendum: *Phys. Rev.C*90,059901(2014)]. [arXiv:1402.2477](#), [doi:10.1103/PhysRevC.90.059901](#), [doi:10.1103/PhysRevC.89.045206](#).
- [184] A. Accardi, W. Melnitchouk, J. Owens, M. Christy, C. Keppel, et al., Uncertainties in determining parton distributions at large x , *Phys.Rev.* D84 (2011) 014008. [arXiv:1102.3686](#), [doi:10.1103/PhysRevD.84.014008](#).
- [185] A. Accardi, M. Christy, C. Keppel, P. Monaghan, W. Melnitchouk, et al., New parton distributions from large- x and low- Q^2 data, *Phys.Rev.* D81 (2010) 034016. [arXiv:0911.2254](#), [doi:10.1103/PhysRevD.81.034016](#).
- [186] J. J. Aubert, et al., Production of charmed particles in 250-GeV μ^+ - iron interactions, *Nucl. Phys.* B213 (1983) 31–64. [doi:10.1016/0550-3213\(83\)90174-8](#).
- [187] W. G. Seligman, et al., Improved determination of α_s from neutrino nucleon scattering, *Phys. Rev. Lett.* 79 (1997) 1213–1216. [arXiv:hep-ex/9701017](#), [doi:10.1103/PhysRevLett.79.1213](#).
- [188] U.-K. Yang, et al., Measurements of F_2 and $xF_3^y - xF_3^{\bar{y}}$ from CCFR ν_μ -Fe and $\bar{\nu}_\mu$ -Fe data in a physics model independent way, *Phys. Rev. Lett.* 86 (2001) 2742–2745. [arXiv:hep-ex/0009041](#), [doi:10.1103/PhysRevLett.86.2742](#).
- [189] G. Onengut, et al., Measurement of nucleon structure functions in neutrino scattering, *Phys. Lett.* B632 (2006) 65–75. [doi:10.1016/j.physletb.2005.10.062](#).
- [190] G. P. Zeller, et al., A Precise determination of electroweak parameters in neutrino nucleon scattering, *Phys. Rev. Lett.* 88 (2002) 091802, [Erratum: *Phys. Rev. Lett.*90,239902(2003)]. [arXiv:hep-ex/0110059](#), [doi:10.1103/PhysRevLett.88.091802](#).
- [191] M. Goncharov, et al., Precise measurement of dimuon production cross-sections in ν_μ Fe and $\bar{\nu}_\mu$ Fe deep inelastic scattering at the Tevatron, *Phys. Rev.* D64 (2001) 112006. [arXiv:hep-ex/0102049](#), [doi:10.1103/PhysRevD.64.112006](#).
- [192] D. A. Mason, Measurement of the strange - antistrange asymmetry at NLO in QCD from NuTeV dimuon data. FERMILAB-THESIS-2006-01.
- [193] A. Kayis-Topaksu, et al., Leading order analysis of neutrino induced dimuon events in the CHORUS experiment, *Nucl. Phys.* B798 (2008) 1–16. [arXiv:0804.1869](#), [doi:10.1016/j.nuclphysb.2008.02.013](#).
- [194] O. Samoylov, et al., A Precision Measurement of Charm Dimuon Production in Neutrino Interactions from the NOMAD Experiment, *Nucl.Phys.* B876 (2013) 339. [arXiv:1308.4750](#), [doi:10.1016/j.nuclphysb.2013.08.021](#).
- [195] H. Abramowicz, et al., Combination and QCD Analysis of Charm Production Cross Section Measurements in Deep-Inelastic ep Scattering at HERA, *Eur.Phys.J. C* 73 (2013) 2311. [arXiv:1211.1182](#), [doi:10.1140/epjc/s10052-013-2311-3](#).

- [196] F. D. Aaron, et al., Measurement of the Charm and Beauty Structure Functions using the H1 Vertex Detector at HERA, Eur. Phys. J. C65 (2010) 89–109. [arXiv:0907.2643](#), [doi:10.1140/epjc/s10052-009-1190-0](#).
- [197] H. Abramowicz, et al., Measurement of beauty and charm production in deep inelastic scattering at HERA and measurement of the beauty-quark mass, JHEP 09 (2014) 127. [arXiv:1405.6915](#), [doi:10.1007/JHEP09\(2014\)127](#).
- [198] F. Aaron, et al., Inclusive Deep Inelastic Scattering at High Q^2 with Longitudinally Polarised Lepton Beams at HERA, JHEP 1209 (2012) 061. [arXiv:1206.7007](#), [doi:10.1007/JHEP09\(2012\)061](#).
- [199] F. Aaron, et al., Measurement of the Inclusive e^+p Scattering Cross Section at High Inelasticity y and of the Structure Function F_L , Eur.Phys.J. C71 (2011) 1579. [arXiv:1012.4355](#), [doi:10.1140/epjc/s10052-011-1579-4](#).
- [200] H. Abramowicz, et al., Measurement of high- Q^2 neutral current deep inelastic e^+p scattering cross sections with a longitudinally polarized positron beam at HERA, Phys. Rev. D87 (5) (2013) 052014. [arXiv:1208.6138](#), [doi:10.1103/PhysRevD.87.052014](#).
- [201] H. Abramowicz, et al., Measurement of high- Q^2 charged current deep inelastic scattering cross sections with a longitudinally polarised positron beam at HERA, Eur.Phys.J. C70 (2010) 945–963. [arXiv:1008.3493](#), [doi:10.1140/epjc/s10052-010-1498-9](#).
- [202] J. Rojo, Progress in the NNPDF global analysis and the impact of the legacy HERA combination, in: Proceedings, 2015 European Physical Society Conference on High Energy Physics (EPS-HEP 2015), 2015. [arXiv:1508.07731](#).
- [203] L. A. Harland-Lang, A. D. Martin, P. Motylinski, R. S. Thorne, The impact of the final HERA combined data on PDFs obtained from a global fit, Eur. Phys. J. C76 (4) (2016) 186. [arXiv:1601.03413](#), [doi:10.1140/epjc/s10052-016-4020-1](#).
- [204] T.-J. Hou, S. Dulat, J. Gao, M. Guzzi, J. Huston, P. Nadolsky, J. Pumplin, C. Schmidt, D. Stump, C. P. Yuan, CTEQ-TEA parton distribution functions and HERA Run I and II combined data, Phys. Rev. D95 (3) (2017) 034003. [arXiv:1609.07968](#), [doi:10.1103/PhysRevD.95.034003](#).
- [205] V. Andreev, et al., Measurement of inclusive ep cross sections at high Q^2 at $\sqrt{s} = 225$ and 252 GeV and of the longitudinal proton structure function F_L at HERA, Eur.Phys.J. C74 (2014) 2814. [arXiv:1312.4821](#), [doi:10.1140/epjc/s10052-014-2814-6](#).
- [206] H. Abramowicz, et al., Deep inelastic cross-section measurements at large y with the ZEUS detector at HERA, Phys. Rev. D90 (7) (2014) 072002. [arXiv:1404.6376](#), [doi:10.1103/PhysRevD.90.072002](#).
- [207] S. Moch, J. A. M. Vermaseren, A. Vogt, The longitudinal structure function at the third order, Phys. Lett. B606 (2005) 123. [arXiv:hep-ph/0411112](#), [doi:10.1016/j.physletb.2004.11.063](#).
- [208] J. A. M. Vermaseren, A. Vogt, S. Moch, The third-order QCD corrections to deep-inelastic scattering by photon exchange, Nucl. Phys. B724 (2005) 3. [arXiv:hep-ph/0504242](#), [doi:10.1016/j.nuclphysb.2005.06.020](#).
- [209] J. Ablinger, A. Behring, J. Blümlein, A. De Freitas, A. von Manteuffel, C. Schneider, The 3-loop pure singlet heavy flavor contributions to the structure function $F_2(x, Q^2)$ and the anomalous dimension, Nucl. Phys. B890 (2014) 48–151. [arXiv:1409.1135](#), [doi:10.1016/j.nuclphysb.2014.10.008](#).
- [210] J. Ablinger, A. Behring, J. Blümlein, A. De Freitas, A. Hasselhuhn, A. von Manteuffel, C. G. Raab, M. Round, C. Schneider, F. Wibrock, 3-Loop Corrections to the Heavy Flavor Wilson Coefficients in Deep-Inelastic Scattering, PoS EPS-HEP2015 (2015) 504. [arXiv:1602.00583](#).
- [211] S. Kretzer, I. Schienbein, Heavy quark initiated contributions to deep inelastic structure functions, Phys. Rev. D58 (1998) 094035. [arXiv:hep-ph/9805233](#), [doi:10.1103/PhysRevD.58.094035](#).
- [212] V. Bertone, APFEL++: A new PDF evolution library in C++, in: 25th International Workshop on Deep Inelastic Scattering and Related Topics (DIS 2017) Birmingham, UK, April 3-7, 2017, 2017. [arXiv:1708.00911](#).
- [213] S. Alekhin, S. Moch, Heavy-quark deep-inelastic scattering with a running mass, Phys. Lett. B699 (2011) 345–353. [arXiv:1011.5790](#), [doi:10.1016/j.physletb.2011.04.026](#).
- [214] S. Alekhin, J. Blümlein, S. Moch, Strange sea determination from collider data. [arXiv:1708.01067](#).
- [215] S. Alekhin, J. Blümlein, L. Caminadac, K. Lipka, K. Lohwasser, S. Moch, R. Petti, R. Placakyte, Determination of Strange Sea Quark Distributions from Fixed-target and Collider Data, Phys. Rev. D91 (9) (2015) 094002. [arXiv:1404.6469](#), [doi:10.1103/PhysRevD.91.094002](#).
- [216] R. D. Ball, et al., Precision determination of electroweak parameters and the strange content of the proton from neutrino deep-inelastic scattering, Nucl. Phys. B823 (2009) 195–233. [arXiv:0906.1958](#), [doi:10.1016/j.nuclphysb.2009.08.003](#).
- [217] H. L. Lai, et al., The Strange Parton Distribution of the Nucleon: Global Analysis and Applications, JHEP 04 (2007) 089. [arXiv:hep-ph/0702268](#).
- [218] F. Olness, et al., Neutrino dimuon production and the strangeness asymmetry of the nucleon, Eur. Phys. J. C40 (2005) 145–156. [arXiv:hep-ph/0312323](#), [doi:10.1140/epjc/s2004-02099-4](#).
- [219] G. P. Salam, Towards Jetography, Eur. Phys. J. C67 (2010) 637–686. [arXiv:0906.1833](#), [doi:10.1140/epjc/s10052-010-1314-6](#).

- [220] M. Cacciari, G. P. Salam, G. Soyez, The Anti-k(t) jet clustering algorithm, JHEP 0804 (2008) 063. [arXiv:0802.1189](#), [doi:10.1088/1126-6708/2008/04/063](#).
- [221] S. Catani, Y. L. Dokshitzer, M. H. Seymour, B. R. Webber, Longitudinally invariant K_t clustering algorithms for hadron hadron collisions, Nucl. Phys. B406 (1993) 187–224. [doi:10.1016/0550-3213\(93\)90166-M](#).
- [222] S. D. Ellis, D. E. Soper, Successive combination jet algorithm for hadron collisions, Phys. Rev. D48 (1993) 3160–3166. [arXiv:hep-ph/9305266](#), [doi:10.1103/PhysRevD.48.3160](#).
- [223] Y. L. Dokshitzer, G. D. Leder, S. Moretti, B. R. Webber, Better jet clustering algorithms, JHEP 08 (1997) 001. [arXiv:hep-ph/9707323](#), [doi:10.1088/1126-6708/1997/08/001](#).
- [224] G. P. Salam, G. Soyez, A practical Seedless Infrared-Safe Cone jet algorithm, JHEP 05 (2007) 086. [arXiv:0704.0292](#).
- [225] G. C. Blazey, et al., Run II jet physics, in: QCD and weak boson physics in Run II. Proceedings, Batavia, USA, March 4-6, June 3-4, November 4-6, 1999, 2000, pp. 47–77. [arXiv:hep-ex/0005012](#).
- [226] M. Cacciari, J. Rojo, G. P. Salam, G. Soyez, Quantifying the performance of jet definitions for kinematic reconstruction at the LHC, JHEP 12 (2008) 032. [arXiv:0810.1304](#), [doi:10.1088/1126-6708/2008/12/032](#).
- [227] V. Khachatryan, et al., Measurement and QCD analysis of double-differential inclusive jet cross-sections in pp collisions at $\sqrt{s} = 8$ TeV and ratios to 2.76 and 7 TeV, JHEP 03 (2017) 156. [arXiv:1609.05331](#), [doi:10.1007/JHEP03\(2017\)156](#).
- [228] S. Alioli, K. Hamilton, P. Nason, C. Oleari, E. Re, Jet pair production in POWHEG, JHEP 1104 (2011) 081. [arXiv:1012.3380](#), [doi:10.1007/JHEP04\(2011\)081](#).
- [229] J. Currie, E. W. N. Glover, A. Gehrmann-De Ridder, T. Gehrmann, A. Huss, J. Pires, Single jet inclusive production for the individual jet p_T scale choice at the LHC, in: 23rd Cracow Epiphany Conference on Particle Theory Meets the First Data from LHC Run 2 Cracow, Poland, January 9-12, 2017, 2017. [arXiv:1704.00923](#).
- [230] G. Aad, et al., Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, Phys. Rev. D86 (2012) 014022. [arXiv:1112.6297](#).
- [231] J. Currie, T. Gehrmann, J. Niehues, Precise QCD predictions for the production of dijet final states in deep inelastic scattering, Phys. Rev. Lett. 117 (4) (2016) 042001. [arXiv:1606.03991](#), [doi:10.1103/PhysRevLett.117.042001](#).
- [232] A. Abulencia, et al., Measurement of the Inclusive Jet Cross Section using the k_T algorithm in $p\bar{p}$ Collisions at $\sqrt{s}=1.96$ TeV with the CDF II Detector, Phys. Rev. D75 (2007) 092006. [arXiv:hep-ex/0701051](#), [doi:10.1103/PhysRevD.75.092006](#).
- [233] T. Aaltonen, et al., Measurement of the Inclusive Jet Cross Section at the Fermilab Tevatron p-pbar Collider Using a Cone-Based Jet Algorithm, Phys. Rev. D78 (2008) 052006. [arXiv:0807.2204](#), [doi:10.1103/PhysRevD.78.052006](#).
- [234] V. M. Abazov, et al., Measurement of the inclusive jet cross-section in $p\bar{p}$ collisions at $s^{(1/2)} = 1.96$ -TeV, Phys. Rev. Lett. 101 (2008) 062001. [arXiv:0802.2400](#), [doi:10.1103/PhysRevLett.101.062001](#).
- [235] V. M. Abazov, et al., Measurement of the inclusive jet cross section in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. D85 (2012) 052006. [arXiv:1110.3771](#), [doi:10.1103/PhysRevD.85.052006](#).
- [236] G. Aad, et al., Measurement of inclusive jet and dijet cross sections in proton-proton collisions at 7 TeV centre-of-mass energy with the ATLAS detector, Eur. Phys. J. C71 (2011) 1512. [arXiv:1009.5908](#), [doi:10.1140/epjc/s10052-010-1512-2](#).
- [237] G. Aad, et al., Measurement of the inclusive jet cross section in pp collisions at $\sqrt{s}=2.76$ TeV and comparison to the inclusive jet cross section at $\sqrt{s}=7$ TeV using the ATLAS detector, Eur.Phys.J. C73 (2013) 2509. [arXiv:1304.4739](#), [doi:10.1140/epjc/s10052-013-2509-4](#).
- [238] G. Aad, et al., Measurement of the inclusive jet cross-section in proton-proton collisions at $\sqrt{s} = 7$ TeV using 4.5 fb¹ of data with the ATLAS detector, JHEP 02 (2015) 153, [Erratum: JHEP09,141(2015)]. [arXiv:1410.8857](#), [doi:10.1007/JHEP02\(2015\)153](#), [doi:10.1007/JHEP09\(2015\)141](#).
- [239] S. Chatrchyan, et al., Measurements of differential jet cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV with the CMS detector, Phys.Rev. D87 (2013) 112002. [arXiv:1212.6660](#), [doi:10.1103/PhysRevD.87.112002](#).
- [240] S. Chatrchyan, et al., Measurement of the ratio of inclusive jet cross sections using the anti- k_T algorithm with radius parameters $R=0.5$ and 0.7 in pp collisions at $\sqrt{s} = 7$ TeV, Phys. Rev. D90 (7) (2014) 072006. [arXiv:1406.0324](#), [doi:10.1103/PhysRevD.90.072006](#).
- [241] G. Aad, et al., Measurement of dijet cross sections in pp collisions at 7 TeV centre-of-mass energy using the ATLAS detector, JHEP 1405 (2014) 059. [arXiv:1312.3524](#), [doi:10.1007/JHEP05\(2014\)059](#).
- [242] S. Chatrchyan, et al., Measurement of the differential dijet production cross section in proton-proton collisions at $\sqrt{s} = 7$ TeV, Phys. Lett. B700 (2011) 187–206. [arXiv:1104.1693](#), [doi:10.1016/j.physletb.2011.05.027](#).
- [243] V. Khachatryan, et al., Measurement of the double-differential inclusive jet cross section in protonproton collisions at $\sqrt{s} = 13$ TeV, Eur. Phys. J. C76 (8) (2016) 451. [arXiv:1605.04436](#), [doi:10.1140/epjc/s10052-016-4286-3](#).
- [244] A. M. Sirunyan, et al., Measurement of the triple-differential dijet cross section in proton-proton collisions at $\sqrt{s} = 8$ TeV and constraints on parton distribution functions. [arXiv:1705.02628](#).

- [245] M. L. Mangano, J. Rojo, Cross Section Ratios between different CM energies at the LHC: opportunities for precision measurements and BSM sensitivity, JHEP 1208 (2012) 010. [arXiv:1206.3557](#), [doi:10.1007/JHEP08\(2012\)010](#).
- [246] S. D. Ellis, Z. Kunszt, D. E. Soper, Two jet production in hadron collisions at order α_s^3 in QCD, Phys.Rev.Lett. 69 (1992) 1496–1499. [doi:10.1103/PhysRevLett.69.1496](#).
- [247] Z. Kunszt, D. E. Soper, Calculation of jet cross-sections in hadron collisions at order α_s^3 , Phys. Rev. D46 (1992) 192–221. [doi:10.1103/PhysRevD.46.192](#).
- [248] Z. Nagy, Three jet cross-sections in hadron hadron collisions at next-to-leading order, Phys.Rev.Lett. 88 (2002) 122003. [arXiv:hep-ph/0110315](#), [doi:10.1103/PhysRevLett.88.122003](#).
- [249] Z. Nagy, Next-to-leading order calculation of three-jet observables in hadron hadron collision, Phys. Rev. D68 (2003) 094002. [arXiv:hep-ph/0307268](#), [doi:10.1103/PhysRevD.68.094002](#).
- [250] J. Gao, Z. Liang, D. E. Soper, H.-L. Lai, P. M. Nadolsky, C. P. Yuan, MEKS: a program for computation of inclusive jet cross sections at hadron colliders, Comput. Phys. Commun. 184 (2013) 1626–1642. [arXiv:1207.0513](#), [doi:10.1016/j.cpc.2013.01.022](#).
- [251] A. Gehrmann-De Ridder, T. Gehrmann, E. Glover, J. Pires, Second order QCD corrections to jet production at hadron colliders: the all-gluon contribution, Phys.Rev.Lett. 110 (2013) 162003. [arXiv:1301.7310](#), [doi:10.1103/PhysRevLett.110.162003](#).
- [252] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, N. Glover, J. Pires, S. Wells, Second order QCD corrections to gluonic jet production at hadron colliders, PoS LL2014 (2014) 001. [arXiv:1407.5558](#).
- [253] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, T. A. Morgan, Precise QCD predictions for the production of a Z boson in association with a hadronic jet, Phys. Rev. Lett. 117 (2) (2016) 022001. [arXiv:1507.02850](#), [doi:10.1103/PhysRevLett.117.022001](#).
- [254] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, T. A. Morgan, NNLO QCD corrections for Z boson plus jet production, in: Proceedings, 12th International Symposium on Radiative Corrections (Radcor 2015) and LoopFest XIV (Radiative Corrections for the LHC and Future Colliders): Los Angeles, CA, USA, June 15-19, 2015, 2016. [arXiv:1601.04569](#).
- [255] S. Dittmaier, A. Huss, C. Speckner, Weak radiative corrections to dijet production at hadron colliders, JHEP 1211 (2012) 095. [arXiv:1210.0438](#), [doi:10.1007/JHEP11\(2012\)095](#).
- [256] S. Carrazza, J. Pires, Perturbative QCD description of jet data from LHC Run-I and Tevatron Run-II, JHEP 10 (2014) 145. [arXiv:1407.7031](#), [doi:10.1007/JHEP10\(2014\)145](#).
- [257] N. Kidonakis, J. F. Owens, Effects of higher-order threshold corrections in high-E(T) jet production, Phys. Rev. D63 (2001) 054019. [arXiv:hep-ph/0007268](#), [doi:10.1103/PhysRevD.63.054019](#).
- [258] M. C. Kumar, S.-O. Moch, Phenomenology of threshold corrections for inclusive jet production at hadron colliders, Phys. Lett. B730 (2014) 122–129. [arXiv:1309.5311](#), [doi:10.1016/j.physletb.2014.01.034](#).
- [259] M. Klasen, G. Kramer, M. Michael, Next-to-next-to-leading order contributions to jet photoproduction and determination of α_s , Phys. Rev. D89 (7) (2014) 074032. [arXiv:1310.1724](#), [doi:10.1103/PhysRevD.89.074032](#).
- [260] D. de Florian, P. Hinderer, A. Mukherjee, F. Ringer, W. Vogelsang, Approximate next-to-next-to-leading order corrections to hadronic jet production, Phys.Rev.Lett. 112 (2014) 082001. [arXiv:1310.7192](#), [doi:10.1103/PhysRevLett.112.082001](#).
- [261] B. J. A. Watt, P. Motylinski, R. S. Thorne, The Effect of LHC Jet Data on MSTW PDFs, Eur.Phys.J. C74 (2014) 2934. [arXiv:1311.5703](#), [doi:10.1140/epjc/s10052-014-2934-z](#).
- [262] S. D. Drell, T.-M. Yan, Massive Lepton Pair Production in Hadron-Hadron Collisions at High-Energies, Phys. Rev. Lett. 25 (1970) 316–320, [Erratum: Phys. Rev. Lett.25,902(1970)]. [doi:10.1103/PhysRevLett.25.316](#).
- [263] P. M. Nadolsky, et al., Implications of CTEQ global analysis for collider observables, Phys. Rev. D78 (2008) 013004. [arXiv:0802.0007](#), [doi:10.1103/PhysRevD.78.013004](#).
- [264] S. D. Ellis, W. J. Stirling, Constraints on isospin breaking in the light quark sea from the Drell-Yan process, Phys. Lett. B256 (1991) 258–264. [doi:10.1016/0370-2693\(91\)90684-I](#).
- [265] A. Martin, A. T. Mathijssen, W. Stirling, R. Thorne, B. Watt, et al., Extended Parameterisations for MSTW PDFs and their effect on Lepton Charge Asymmetry from W Decays, Eur.Phys.J. C73 (2) (2013) 2318. [arXiv:1211.1215](#), [doi:10.1140/epjc/s10052-013-2318-9](#).
- [266] R. S. Towell, et al., Improved measurement of the anti-d/anti-u asymmetry in the nucleon sea, Phys. Rev. D64 (2001) 052002. [arXiv:hep-ex/0103030](#), [doi:10.1103/PhysRevD.64.052002](#).
- [267] Fermilab E906 experiment, Drell-Yan Measurements of Nucleon and Nuclear Structure with the Fermilab Main Injector, D. F. Geesaman and P. E. Reimer, spokespersons; <http://www.phy.anl.gov/mep/SeaQuest/index.html>.
- [268] V. M. Abazov, et al., Measurement of the shape of the boson rapidity distribution for $p\bar{p} \rightarrow Z/\gamma^* \rightarrow e^+e^- + X$ events produced at $\sqrt{s}=1.96$ -TeV, Phys. Rev. D76 (2007) 012003. [arXiv:hep-ex/0702025](#), [doi:10.1103/PhysRevD.76.012003](#).

- [269] T. A. Aaltonen, et al., Measurement of $d\sigma/dy$ of Drell-Yan e^+e^- pairs in the Z Mass Region from $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV, Phys. Lett. B692 (2010) 232–239. [arXiv:0908.3914](#), [doi:10.1016/j.physletb.2010.06.043](#).
- [270] V. M. Abazov, et al., Measurement of the muon charge asymmetry in $p\bar{p} \rightarrow W+X \rightarrow \mu\nu + X$ events at $\sqrt{s}=1.96$ TeV, Phys.Rev. D88 (2013) 091102. [arXiv:1309.2591](#), [doi:10.1103/PhysRevD.88.091102](#).
- [271] V. M. Abazov, et al., Measurement of the electron charge asymmetry in $p\bar{p} \rightarrow W + X \rightarrow e\nu + X$ decays in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. D91 (3) (2015) 032007, [Erratum: Phys. Rev.D91,no.7,079901(2015)]. [arXiv:1412.2862](#), [doi:10.1103/PhysRevD.91.032007](#), [doi:10.1103/PhysRevD.91.079901](#).
- [272] T. Aaltonen, et al., Direct Measurement of the W Production Charge Asymmetry in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. Lett. 102 (2009) 181801. [arXiv:0901.2169](#), [doi:10.1103/PhysRevLett.102.181801](#).
- [273] V. M. Abazov, et al., Measurement of the W Boson Production Charge Asymmetry in $p\bar{p} \rightarrow W + X \rightarrow e\nu + X$ Events at $\sqrt{s} = 1.96$ TeV, Phys. Rev. Lett. 112 (15) (2014) 151803, [Erratum: Phys. Rev. Lett.114,no.4,049901(2015)]. [arXiv:1312.2895](#), [doi:10.1103/PhysRevLett.114.049901](#), [doi:10.1103/PhysRevLett.112.151803](#).
- [274] S. Chatrchyan, et al., Measurement of the Rapidity and Transverse Momentum Distributions of Z Bosons in pp Collisions at $\sqrt{s} = 7$ TeV, Phys. Rev. D85 (2012) 032002. [arXiv:1110.4973](#), [doi:10.1103/PhysRevD.85.032002](#).
- [275] G. Aad, et al., Measurement of the inclusive W^\pm and Z/γ^* cross sections in the electron and muon decay channels in pp collisions at $\sqrt{s}=7$ TeV with the ATLAS detector, Phys.Rev. D85 (2012) 072004. [arXiv:1109.5141](#).
- [276] S. Chatrchyan, et al., Measurement of the differential and double-differential Drell-Yan cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV, JHEP 12 (2013) 030. [arXiv:1310.7291](#), [doi:10.1007/JHEP12\(2013\)030](#).
- [277] V. Khachatryan, et al., Measurements of differential and double-differential Drell-Yan cross sections in proton-proton collisions at 8 TeV, Eur. Phys. J. C75 (4) (2015) 147. [arXiv:1412.1115](#), [doi:10.1140/epjc/s10052-015-3364-2](#).
- [278] G. Aad, et al., Measurement of the high-mass Drell–Yan differential cross-section in pp collisions at $\sqrt{s}=7$ TeV with the ATLAS detector, Phys.Lett. B725 (2013) 223. [arXiv:1305.4192](#), [doi:10.1016/j.physletb.2013.07.049](#).
- [279] G. Aad, et al., Measurement of the low-mass Drell-Yan differential cross section at $\sqrt{s} = 7$ TeV using the ATLAS detector, JHEP 06 (2014) 112. [arXiv:1404.1212](#), [doi:10.1007/JHEP06\(2014\)112](#).
- [280] S. Chatrchyan, et al., Measurement of the electron charge asymmetry in inclusive W production in pp collisions at $\sqrt{s} = 7$ TeV, Phys.Rev.Lett. 109 (2012) 111806. [arXiv:1206.2598](#), [doi:10.1103/PhysRevLett.109.111806](#).
- [281] S. Chatrchyan, et al., Measurement of the muon charge asymmetry in inclusive pp to WX production at $\sqrt{s} = 7$ TeV and an improved determination of light parton distribution functions, Phys.Rev. D90 (2014) 032004. [arXiv:1312.6283](#), [doi:10.1103/PhysRevD.90.032004](#).
- [282] M. Aaboud, et al., Precision measurement and interpretation of inclusive W^+ , W^- and Z/γ^* production cross sections with the ATLAS detector, Eur. Phys. J. C77 (6) (2017) 367. [arXiv:1612.03016](#), [doi:10.1140/epjc/s10052-017-4911-9](#).
- [283] R. Aaij, et al., Measurement of the cross-section for $Z \rightarrow e^+e^-$ production in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 1302 (2013) 106. [arXiv:1212.4620](#), [doi:10.1007/JHEP02\(2013\)106](#).
- [284] R. Aaij, et al., Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 08 (2015) 039. [arXiv:1505.07024](#), [doi:10.1007/JHEP08\(2015\)039](#).
- [285] R. Aaij, et al., Measurement of forward $Z \rightarrow e^+e^-$ production at $\sqrt{s} = 8$ TeV, JHEP 05 (2015) 109. [arXiv:1503.00963](#), [doi:10.1007/JHEP05\(2015\)109](#).
- [286] R. Aaij, et al., Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 13$ TeV, JHEP 09 (2016) 136. [arXiv:1607.06495](#), [doi:10.1007/JHEP09\(2016\)136](#).
- [287] R. Aaij, et al., Measurement of the forward W boson cross-section in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 12 (2014) 079. [arXiv:1408.4354](#), [doi:10.1007/JHEP12\(2014\)079](#).
- [288] R. Aaij, et al., Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV, JHEP 01 (2016) 155. [arXiv:1511.08039](#), [doi:10.1007/JHEP01\(2016\)155](#).
- [289] R. Aaij, et al., Measurement of forward $W \rightarrow e\nu$ production in pp collisions at $\sqrt{s} = 8$ TeV, JHEP 10 (2016) 030. [arXiv:1608.01484](#), [doi:10.1007/JHEP10\(2016\)030](#).
- [290] R. V. Harlander, W. B. Kilgore, Next-to-next-to-leading order Higgs production at hadron colliders, Phys. Rev. Lett. 88 (2002) 201801. [arXiv:hep-ph/0201206](#), [doi:10.1103/PhysRevLett.88.201801](#).
- [291] C. Anastasiou, L. J. Dixon, K. Melnikov, F. Petriello, High precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at NNLO, Phys. Rev. D69 (2004) 094008. [arXiv:hep-ph/0312266](#), [doi:10.1103/PhysRevD.69.094008](#).
- [292] C. Anastasiou, L. J. Dixon, K. Melnikov, F. Petriello, Dilepton rapidity distribution in the Drell-Yan process at NNLO in QCD, Phys. Rev. Lett. 91 (2003) 182002. [arXiv:hep-ph/0306192](#), [doi:10.1103/PhysRevLett.91.182002](#).
- [293] K. Melnikov, F. Petriello, The W boson production cross section at the LHC through $O(\alpha_s^2)$, Phys. Rev. Lett. 96 (2006) 231803. [arXiv:hep-ph/0603182](#), [doi:10.1103/PhysRevLett.96.231803](#).
- [294] K. Melnikov, F. Petriello, Electroweak gauge boson production at hadron colliders through $O(\alpha_s^2)$, Phys. Rev. D74 (2006) 114017. [arXiv:hep-ph/0609070](#), [doi:10.1103/PhysRevD.74.114017](#).

- [295] R. Gavin, Y. Li, F. Petriello, S. Quackenbush, FEWZ 2.0: A code for hadronic Z production at next-to-next-to-leading order, *Comput. Phys. Commun.* 182 (2011) 2388–2403. [arXiv:1011.3540](#), [doi:10.1016/j.cpc.2011.06.008](#).
- [296] R. Gavin, Y. Li, F. Petriello, S. Quackenbush, W Physics at the LHC with FEWZ 2.1, *Comput.Phys.Commun.* 184 (2013) 208–214. [arXiv:1201.5896](#), [doi:10.1016/j.cpc.2012.09.005](#).
- [297] Y. Li, F. Petriello, Combining QCD and electroweak corrections to dilepton production in FEWZ, *Phys.Rev.* D86 (2012) 094034. [arXiv:1208.5967](#), [doi:10.1103/PhysRevD.86.094034](#).
- [298] F. A. Berends, R. Kleiss, Hard Photon Effects in W^\pm and Z^0 Decay, *Z.Phys.* C27 (1985) 365. [doi:10.1007/BF01548639](#).
- [299] F. A. Berends, R. Kleiss, J. Revol, J. Vialle, QED Radiative Corrections and Radiative Decays of the Intermediate Weak Bosons Produced in Proton - Anti-proton Collisions, *Z.Phys.* C27 (1985) 155. [doi:10.1007/BF01642494](#).
- [300] U. Baur, S. Keller, W. Sakumoto, QED radiative corrections to Z boson production and the forward backward asymmetry at hadron colliders, *Phys.Rev.* D57 (1998) 199–215. [arXiv:hep-ph/9707301](#), [doi:10.1103/PhysRevD.57.199](#).
- [301] U. Baur, O. Brein, W. Hollik, C. Schappacher, D. Wackerroth, Electroweak radiative corrections to neutral current Drell-Yan processes at hadron colliders, *Phys.Rev.* D65 (2002) 033007. [arXiv:hep-ph/0108274](#), [doi:10.1103/PhysRevD.65.033007](#).
- [302] C. Carloni Calame, G. Montagna, O. Nicrosini, A. Vicini, Precision electroweak calculation of the production of a high transverse-momentum lepton pair at hadron colliders, *JHEP* 0710 (2007) 109. [arXiv:0710.1722](#), [doi:10.1088/1126-6708/2007/10/109](#).
- [303] S. Dittmaier, M. Huber, Radiative corrections to the neutral-current Drell-Yan process in the Standard Model and its minimal supersymmetric extension, *JHEP* 1001 (2010) 060. [arXiv:0911.2329](#), [doi:10.1007/JHEP01\(2010\)060](#).
- [304] S. Catani, L. Cieri, G. Ferrera, D. de Florian, M. Grazzini, Vector boson production at hadron colliders: a fully exclusive QCD calculation at NNLO, *Phys.Rev.Lett.* 103 (2009) 082001. [arXiv:0903.2120](#), [doi:10.1103/PhysRevLett.103.082001](#).
- [305] S. Catani, 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. [arXiv:hep-ph/0703012](#), [doi:10.1103/PhysRevLett.98.222002](#).
- [306] J. M. Campbell, R. K. Ellis, C. Williams, Vector boson pair production at the LHC, *JHEP* 1107 (2011) 018. [arXiv:1105.0020](#), [doi:10.1007/JHEP07\(2011\)018](#).
- [307] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. Giele, X. Liu, F. Petriello, C. Williams, Color singlet production at NNLO in MCFM, *Eur. Phys. J.* C77 (1) (2017) 7. [arXiv:1605.08011](#), [doi:10.1140/epjc/s10052-016-4558-y](#).
- [308] G. Bozzi, S. Catani, G. Ferrera, D. de Florian, M. Grazzini, Production of Drell-Yan lepton pairs in hadron collisions: Transverse-momentum resummation at next-to-next-to-leading logarithmic accuracy, *Phys. Lett.* B696 (2011) 207–213. [arXiv:1007.2351](#), [doi:10.1016/j.physletb.2010.12.024](#).
- [309] C. Balazs, J.-w. Qiu, C. P. Yuan, Effects of QCD resummation on distributions of leptons from the decay of electroweak vector bosons, *Phys. Lett.* B355 (1995) 548–554. [arXiv:hep-ph/9505203](#), [doi:10.1016/0370-2693\(95\)00726-2](#).
- [310] S. Berge, P. M. Nadolsky, F. Olness, C. P. Yuan, Transverse momentum resummation at small x for the Tevatron and CERN LHC, *Phys. Rev.* D72 (2005) 033015. [arXiv:hep-ph/0410375](#), [doi:10.1103/PhysRevD.72.033015](#).
- [311] V. Khachatryan, et al., Measurement of the differential cross section and charge asymmetry for inclusive $pp \rightarrow W^\pm + X$ production at $\sqrt{s} = 8$ TeV, *Eur. Phys. J.* C76 (8) (2016) 469. [arXiv:1603.01803](#), [doi:10.1140/epjc/s10052-016-4293-4](#).
- [312] G. Aad, et al., Determination of the strange quark density of the proton from ATLAS measurements of the $W \rightarrow \ell\nu$ and $Z \rightarrow \ell\ell$ cross sections, *Phys. Rev. Lett.* 109 (2012) 012001. [arXiv:1203.4051](#), [doi:10.1103/PhysRevLett.109.012001](#).
- [313] J. Rojo, Improving quark flavor separation with forward W and Z production at LHCb, in: 25th International Workshop on Deep Inelastic Scattering and Related Topics (DIS 2017) Birmingham, UK, April 3-7, 2017, 2017. [arXiv:1705.04468](#).
- [314] J. C. Collins, D. E. Soper, G. F. Sterman, Transverse Momentum Distribution in Drell-Yan Pair and W and Z Boson Production, *Nucl. Phys.* B250 (1985) 199–224. [doi:10.1016/0550-3213\(85\)90479-1](#).
- [315] C. T. H. Davies, W. J. Stirling, Nonleading Corrections to the Drell-Yan Cross-Section at Small Transverse Momentum, *Nucl. Phys.* B244 (1984) 337–348. [doi:10.1016/0550-3213\(84\)90316-X](#).
- [316] R. K. Ellis, S. Veseli, W and Z transverse momentum distributions: Resummation in q_T space, *Nucl. Phys.* B511 (1998) 649–669. [arXiv:hep-ph/9706526](#), [doi:10.1016/S0550-3213\(97\)00655-X](#).
- [317] J.-w. Qiu, X.-f. Zhang, QCD prediction for heavy boson transverse momentum distributions, *Phys. Rev. Lett.* 86 (2001) 2724–2727. [arXiv:hep-ph/0012058](#), [doi:10.1103/PhysRevLett.86.2724](#).
- [318] F. Landry, R. Brock, P. M. Nadolsky, C. P. Yuan, Tevatron Run-1 Z boson data and Collins-Soper-Sterman resummation formalism, *Phys. Rev.* D67 (2003) 073016. [arXiv:hep-ph/0212159](#), [doi:10.1103/PhysRevD.67.073016](#).
- [319] S. Mantry, F. Petriello, Transverse Momentum Distributions from Effective Field Theory with Numerical Results, *Phys. Rev.* D83 (2011) 053007. [arXiv:1007.3773](#), [doi:10.1103/PhysRevD.83.053007](#).
- [320] T. Becher, M. Neubert, Drell-Yan Production at Small q_T , Transverse Parton Distributions and the Collinear Anomaly, *Eur. Phys. J.* C71 (2011) 1665. [arXiv:1007.4005](#), [doi:10.1140/epjc/s10052-011-1665-7](#).

- [321] S. Catani, L. Cieri, D. de Florian, G. Ferrera, M. Grazzini, Universality of transverse-momentum resummation and hard factors at the NNLO, Nucl. Phys. B881 (2014) 414–443. [arXiv:1311.1654](#), [doi:10.1016/j.nuclphysb.2014.02.011](#).
- [322] J.-w. Qiu, X.-f. Zhang, Role of the nonperturbative input in QCD resummed Drell-Yan Q_T distributions, Phys. Rev. D63 (2001) 114011. [arXiv:hep-ph/0012348](#), [doi:10.1103/PhysRevD.63.114011](#).
- [323] I. Scimemi, A. Vladimirov, Analysis of vector boson production within TMD factorization [arXiv:1706.01473](#).
- [324] T. Becher, G. Bell, C. Lorentzen, S. Marti, Transverse-momentum spectra of electroweak bosons near threshold at NNLO, JHEP 02 (2014) 004. [arXiv:1309.3245](#), [doi:10.1007/JHEP02\(2014\)004](#).
- [325] N. Kidonakis, R. J. Gonsalves, NNLO soft-gluon corrections for the Z-boson and W-boson transverse momentum distributions, Phys. Rev. D89 (9) (2014) 094022. [arXiv:1404.4302](#), [doi:10.1103/PhysRevD.89.094022](#).
- [326] R. Boughezal, A. Guffanti, F. Petriello, M. Ubiali, The impact of the LHC Z-boson transverse momentum data on PDF determinations, JHEP 07 (2017) 130. [arXiv:1705.00343](#), [doi:10.1007/JHEP07\(2017\)130](#).
- [327] G. Aad, et al., Measurement of the transverse momentum and ϕ_η^* distributions of Drell-Yan lepton pairs in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, Eur. Phys. J. C76 (5) (2016) 291. [arXiv:1512.02192](#), [doi:10.1140/epjc/s10052-016-4070-4](#).
- [328] V. Khachatryan, et al., Measurement of the Z boson differential cross section in transverse momentum and rapidity in proton-proton collisions at 8 TeV, Phys. Lett. B749 (2015) 187–209. [arXiv:1504.03511](#), [doi:10.1016/j.physletb.2015.07.065](#).
- [329] G. Aad, et al., Measurement of the transverse momentum distribution of Z/γ^* bosons in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, Phys. Lett. B705 (2011) 415–434. [arXiv:1107.2381](#), [doi:10.1016/j.physletb.2011.10.018](#).
- [330] G. Aad, et al., Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, JHEP 09 (2014) 145. [arXiv:1406.3660](#), [doi:10.1007/JHEP09\(2014\)145](#).
- [331] V. Khachatryan, et al., Measurement of the transverse momentum spectra of weak vector bosons produced in proton-proton collisions at $\sqrt{s} = 8$ TeV, JHEP 02 (2017) 096. [arXiv:1606.05864](#), [doi:10.1007/JHEP02\(2017\)096](#).
- [332] V. M. Abazov, et al., Measurement of the ϕ_η^* distribution of muon pairs with masses between 30 and 500 GeV in 10.4 fb^{-1} of $p\bar{p}$ collisions, Phys. Rev. D91 (7) (2015) 072002. [arXiv:1410.8052](#), [doi:10.1103/PhysRevD.91.072002](#).
- [333] R. J. Gonsalves, J. Pawlowski, C.-F. Wai, QCD Radiative Corrections to Electroweak Boson Production at Large Transverse Momentum in Hadron Collisions, Phys. Rev. D40 (1989) 2245. [doi:10.1103/PhysRevD.40.2245](#).
- [334] H. Baer, M. H. Reno, A Complete $O(\alpha_s)$ event generator for $p\bar{p} \rightarrow W^+X \rightarrow e^+$ neutrino X with parton showering, Phys. Rev. D44 (1991) 3375–3378. [doi:10.1103/PhysRevD.44.3375](#).
- [335] P. B. Arnold, R. P. Kauffman, W and Z production at next-to-leading order: From large $q(t)$ to small, Nucl. Phys. B349 (1991) 381–413. [doi:10.1016/0550-3213\(91\)90330-Z](#).
- [336] E. Maina, S. Moretti, D. A. Ross, One loop weak corrections to gamma / Z hadroproduction at finite transverse momentum, Phys. Lett. B593 (2004) 143–150, [Erratum: Phys. Lett. B614,216(2005)]. [arXiv:hep-ph/0403050](#), [doi:10.1016/j.physletb.2004.04.043](#), [doi:10.1016/j.physletb.2005.03.064](#).
- [337] J. H. Kuhn, A. Kulesza, S. Pozzorini, M. Schulze, Logarithmic electroweak corrections to hadronic Z+1 jet production at large transverse momentum, Phys. Lett. B609 (2005) 277–285. [arXiv:hep-ph/0408308](#), [doi:10.1016/j.physletb.2005.01.059](#).
- [338] J. H. Kuhn, A. Kulesza, S. Pozzorini, M. Schulze, One-loop weak corrections to hadronic production of Z bosons at large transverse momenta, Nucl. Phys. B727 (2005) 368–394. [arXiv:hep-ph/0507178](#), [doi:10.1016/j.nuclphysb.2005.08.019](#).
- [339] T. Becher, X. Garcia i Tormo, Electroweak Sudakov effects in W, Z and γ production at large transverse momentum, Phys. Rev. D88 (1) (2013) 013009. [arXiv:1305.4202](#), [doi:10.1103/PhysRevD.88.013009](#).
- [340] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, T. A. Morgan, NNLO QCD corrections for Drell-Yan p_T^Z and ϕ^* observables at the LHC, JHEP 11 (2016) 094. [arXiv:1610.01843](#), [doi:10.1007/JHEP11\(2016\)094](#).
- [341] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. T. Giele, X. Liu, F. Petriello, Z-boson production in association with a jet at next-to-next-to-leading order in perturbative QCD, Phys. Rev. Lett. 116 (15) (2016) 152001. [arXiv:1512.01291](#), [doi:10.1103/PhysRevLett.116.152001](#).
- [342] R. Boughezal, X. Liu, F. Petriello, Phenomenology of the Z-boson plus jet process at NNLO, Phys. Rev. D94 (7) (2016) 074015. [arXiv:1602.08140](#), [doi:10.1103/PhysRevD.94.074015](#).
- [343] D. d’Enterría, J. Rojo, Quantitative constraints on the gluon distribution function in the proton from collider isolated-photon data, Nucl. Phys. B860 (2012) 311–338. [arXiv:1202.1762](#).
- [344] L. Apanasevich, et al., Evidence for parton k_T effects in high p_T particle production, Phys. Rev. Lett. 81 (1998) 2642–2645. [arXiv:hep-ex/9711017](#), [doi:10.1103/PhysRevLett.81.2642](#).
- [345] L. Apanasevich, et al., Measurement of direct photon production at Tevatron fixed target energies, Phys. Rev. D70 (2004)

092009. [arXiv:hep-ex/0407011](#), [doi:10.1103/PhysRevD.70.092009](#).
- [346] A. D. Martin, R. G. Roberts, W. J. Stirling, R. S. Thorne, Parton distributions and the LHC: W and Z production, Eur. Phys. J. C14 (2000) 133–145. [arXiv:hep-ph/9907231](#), [doi:10.1007/s100520050740](#), [doi:10.1007/s100520000324](#).
- [347] R. Ichou, D. d’Enterria, Sensitivity of isolated photon production at TeV hadron colliders to the gluon distribution in the proton, Phys. Rev. D82 (2010) 014015. [arXiv:1005.4529](#), [doi:10.1103/PhysRevD.82.014015](#).
- [348] L. Carminati, G. Costa, D. D’Enterria, I. Koletsou, G. Marchiori, J. Rojo, M. Stockton, F. Tartarelli, Sensitivity of the LHC isolated-gamma+jet data to the parton distribution functions of the proton, Europhys. Lett. 101 (2013) 61002. [arXiv:1212.5511](#), [doi:10.1209/0295-5075/101/61002](#).
- [349] G. Aad, et al., Measurement of the inclusive isolated prompt photon cross section in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, JHEP 08 (2016) 005. [arXiv:1605.03495](#), [doi:10.1007/JHEP08\(2016\)005](#).
- [350] M. Aaboud, et al., Measurement of the cross section for inclusive isolated-photon production in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector, Phys. Lett. B770 (2017) 473–493. [arXiv:1701.06882](#), [doi:10.1016/j.physletb.2017.04.072](#).
- [351] T. A. Aaltonen, et al., Measurement of the Inclusive-Isolated Prompt-Photon Cross Section in $p\bar{p}$ Collisions using the full CDF Data Set. [arXiv:1703.00599](#).
- [352] V. M. Abazov, et al., Measurement of the isolated photon cross section in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ -TeV, Phys. Lett. B639 (2006) 151–158, [Erratum: Phys. Lett. B658,285(2008)]. [arXiv:hep-ex/0511054](#), [doi:10.1016/j.physletb.2007.06.047](#), [doi:10.1016/j.physletb.2006.04.048](#).
- [353] G. Aad, et al., Measurement of the inclusive isolated prompt photons cross section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector using 4.6fb^{-1} , Phys. Rev. D89 (5) (2014) 052004. [arXiv:1311.1440](#), [doi:10.1103/PhysRevD.89.052004](#).
- [354] S. Chatrchyan, et al., Measurement of the Differential Cross Section for Isolated Prompt Photon Production in pp Collisions at 7 TeV, Phys.Rev. D84 (2011) 052011. [arXiv:1108.2044](#), [doi:10.1103/PhysRevD.84.052011](#).
- [355] G. Aad, et al., Centrality, rapidity and transverse momentum dependence of isolated prompt photon production in lead-lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV measured with the ATLAS detector, Phys. Rev. C93 (3) (2016) 034914. [arXiv:1506.08552](#), [doi:10.1103/PhysRevC.93.034914](#).
- [356] S. Chatrchyan, et al., Measurement of isolated photon production in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, Phys. Lett. B710 (2012) 256–277. [arXiv:1201.3093](#), [doi:10.1016/j.physletb.2012.02.077](#).
- [357] S. Catani, M. Fontannaz, J. P. Guillet, E. Pilon, Cross-section of isolated prompt photons in hadron hadron collisions, JHEP 05 (2002) 028. [arXiv:hep-ph/0204023](#), [doi:10.1088/1126-6708/2002/05/028](#).
- [358] J. H. Kuhn, A. Kulesza, S. Pozzorini, M. Schulze, Electroweak corrections to hadronic photon production at large transverse momenta, JHEP 03 (2006) 059. [arXiv:hep-ph/0508253](#), [doi:10.1088/1126-6708/2006/03/059](#).
- [359] M. D. Schwartz, Precision direct photon spectra at high energy and comparison to the 8 TeV ATLAS data, JHEP 09 (2016) 005. [arXiv:1606.02313](#), [doi:10.1007/JHEP09\(2016\)005](#).
- [360] T. A. Aaltonen, et al., Combination of measurements of the top-quark pair production cross section from the Tevatron Collider, Phys.Rev. D89 (2014) 072001. [arXiv:1309.7570](#), [doi:10.1103/PhysRevD.89.072001](#).
- [361] G. Aad, et al., Measurement of the cross section for top-quark pair production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector using final states with two high-pt leptons, JHEP 1205 (2012) 059. [arXiv:1202.4892](#), [doi:10.1007/JHEP05\(2012\)059](#).
- [362] G. Aad, et al., Measurement of the $t\bar{t}$ production cross-section in pp collisions at $\sqrt{s} = 7$ TeV using kinematic information of lepton+jets events. ATLAS-CONF-2011-121, ATLAS-COM-CONF-2011-132.
- [363] G. Aad, et al., Measurement of the $t\bar{t}$ production cross-section in pp collisions at $\sqrt{s} = 8$ TeV using $e\mu$ events with b -tagged jets. ATLAS-CONF-2013-097, ATLAS-COM-CONF-2013-112.
- [364] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{s} = 8$ TeV, JHEP 1402 (2014) 024. [arXiv:1312.7582](#), [doi:10.1007/JHEP02\(2014\)024](#).
- [365] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 1211 (2012) 067. [arXiv:1208.2671](#), [doi:10.1007/JHEP11\(2012\)067](#).
- [366] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in pp collisions at $\sqrt{s} = 7$ TeV with lepton + jets final states, Phys.Lett. B720 (2013) 83–104. [arXiv:1212.6682](#), [doi:10.1016/j.physletb.2013.02.021](#).
- [367] G. Aad, et al., Measurements of top-quark pair differential cross-sections in the lepton+jets channel in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector, Eur. Phys. J. C76 (10) (2016) 538. [arXiv:1511.04716](#), [doi:10.1140/epjc/s10052-016-4366-4](#).
- [368] V. Khachatryan, et al., Measurement of the differential cross section for top quark pair production in pp collisions at $\sqrt{s} = 8$ TeV, Eur. Phys. J. C75 (11) (2015) 542. [arXiv:1505.04480](#), [doi:10.1140/epjc/s10052-015-3709-x](#).
- [369] A. M. Sirunyan, et al., Measurement of double-differential cross sections for top quark pair production in pp collisions at $\sqrt{s} = 8$ TeV and impact on parton distribution functions, Eur. Phys. J. C77 (7) (2017) 459. [arXiv:1703.01630](#),

[doi:10.1140/epjc/s10052-017-4984-5](https://doi.org/10.1140/epjc/s10052-017-4984-5).

- [370] M. Czakon, A. Mitov, NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction, JHEP 1301 (2013) 080. [arXiv:1210.6832](https://arxiv.org/abs/1210.6832), [doi:10.1007/JHEP01\(2013\)080](https://doi.org/10.1007/JHEP01(2013)080).
- [371] P. Brnreuther, M. Czakon, A. Mitov, Percent Level Precision Physics at the Tevatron: First Genuine NNLO QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$, Phys. Rev. Lett. 109 (2012) 132001. [arXiv:1204.5201](https://arxiv.org/abs/1204.5201), [doi:10.1103/PhysRevLett.109.132001](https://doi.org/10.1103/PhysRevLett.109.132001).
- [372] M. Czakon, A. Mitov, Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders, Comput. Phys. Commun. 185 (2014) 2930. [arXiv:1112.5675](https://arxiv.org/abs/1112.5675), [doi:10.1016/j.cpc.2014.06.021](https://doi.org/10.1016/j.cpc.2014.06.021).
- [373] M. Cacciari, M. Czakon, M. L. Mangano, A. Mitov, P. Nason, Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation, Phys.Lett. B710 (2012) 612–622. [arXiv:1111.5869](https://arxiv.org/abs/1111.5869), [doi:10.1016/j.physletb.2012.03.013](https://doi.org/10.1016/j.physletb.2012.03.013).
- [374] M. Czakon, D. Heymes, A. Mitov, High-precision differential predictions for top-quark pairs at the LHC, Phys. Rev. Lett. 116 (8) (2016) 082003. [arXiv:1511.00549](https://arxiv.org/abs/1511.00549), [doi:10.1103/PhysRevLett.116.082003](https://doi.org/10.1103/PhysRevLett.116.082003).
- [375] M. Czakon, P. Fiedler, A. Mitov, Resolving the Tevatron Top Quark Forward-Backward Asymmetry Puzzle: Fully Differential Next-to-Next-to-Leading-Order Calculation, Phys. Rev. Lett. 115 (5) (2015) 052001. [arXiv:1411.3007](https://arxiv.org/abs/1411.3007), [doi:10.1103/PhysRevLett.115.052001](https://doi.org/10.1103/PhysRevLett.115.052001).
- [376] M. Czakon, D. Heymes, A. Mitov, D. Pagani, I. Tsiniikos, M. Zaro, Top-pair production at the LHC through NNLO QCD and NLO EW. [arXiv:1705.04105](https://arxiv.org/abs/1705.04105).
- [377] D. Pagani, I. Tsiniikos, M. Zaro, The impact of the photon PDF and electroweak corrections on $t\bar{t}$ distributions, Eur. Phys. J. C76 (9) (2016) 479. [arXiv:1606.01915](https://arxiv.org/abs/1606.01915), [doi:10.1140/epjc/s10052-016-4318-z](https://doi.org/10.1140/epjc/s10052-016-4318-z).
- [378] J. Gao, A. S. Papanastasiou, Top-quark pair-production and decay at high precision. [arXiv:1705.08903](https://arxiv.org/abs/1705.08903).
- [379] M. Czakon, N. P. Hartland, A. Mitov, E. R. Nocera, J. Rojo, Pinning down the large-x gluon with NNLO top-quark pair differential distributions, JHEP 04 (2017) 044. [arXiv:1611.08609](https://arxiv.org/abs/1611.08609), [doi:10.1007/JHEP04\(2017\)044](https://doi.org/10.1007/JHEP04(2017)044).
- [380] M. Czakon, M. L. Mangano, A. Mitov, J. Rojo, Constraints on the gluon PDF from top quark pair production at hadron colliders, JHEP 1307 (2013) 167. [arXiv:1303.7215](https://arxiv.org/abs/1303.7215), [doi:10.1007/JHEP07\(2013\)167](https://doi.org/10.1007/JHEP07(2013)167).
- [381] M. Beneke, P. Falgari, S. Klein, J. Piclum, C. Schwinn, et al., Inclusive Top-Pair Production Phenomenology with TOPIXS, JHEP 1207 (2012) 194. [arXiv:1206.2454](https://arxiv.org/abs/1206.2454), [doi:10.1007/JHEP07\(2012\)194](https://doi.org/10.1007/JHEP07(2012)194).
- [382] M. Guzzi, K. Lipka, S.-O. Moch, Top-quark pair production at hadron colliders: differential cross section and phenomenological applications with DiffTop, JHEP 01 (2015) 082. [arXiv:1406.0386](https://arxiv.org/abs/1406.0386), [doi:10.1007/JHEP01\(2015\)082](https://doi.org/10.1007/JHEP01(2015)082).
- [383] F. Maltoni, G. Ridolfi, M. Ubiali, b-initiated processes at the LHC: a reappraisal, JHEP 07 (2012) 022, [Erratum: JHEP04,095(2013)]. [arXiv:1203.6393](https://arxiv.org/abs/1203.6393), [doi:10.1007/JHEP04\(2013\)095](https://doi.org/10.1007/JHEP04(2013)095), [doi:10.1007/JHEP07\(2012\)022](https://doi.org/10.1007/JHEP07(2012)022).
- [384] S. Forte, D. Napoletano, M. Ubiali, Higgs production in bottom-quark fusion in a matched scheme, Phys. Lett. B751 (2015) 331–337. [arXiv:1508.01529](https://arxiv.org/abs/1508.01529), [doi:10.1016/j.physletb.2015.10.051](https://doi.org/10.1016/j.physletb.2015.10.051).
- [385] M. Brucherseifer, F. Caola, K. Melnikov, On the NNLO QCD corrections to single-top production at the LHC, Phys. Lett. B736 (2014) 58–63. [arXiv:1404.7116](https://arxiv.org/abs/1404.7116), [doi:10.1016/j.physletb.2014.06.075](https://doi.org/10.1016/j.physletb.2014.06.075).
- [386] E. L. Berger, J. Gao, C. P. Yuan, H. X. Zhu, NNLO QCD Corrections to t-channel Single Top-Quark Production and Decay, Phys. Rev. D94 (7) (2016) 071501. [arXiv:1606.08463](https://arxiv.org/abs/1606.08463), [doi:10.1103/PhysRevD.94.071501](https://doi.org/10.1103/PhysRevD.94.071501).
- [387] M. Aaboud, et al., Fiducial, total and differential cross-section measurements of t-channel single top-quark production in pp collisions at 8 TeV using data collected by the ATLAS detector, Eur. Phys. J. C77 (8) (2017) 531. [arXiv:1702.02859](https://arxiv.org/abs/1702.02859), [doi:10.1140/epjc/s10052-017-5061-9](https://doi.org/10.1140/epjc/s10052-017-5061-9).
- [388] M. Aaboud, et al., Measurement of the inclusive cross-sections of single top-quark and top-antiquark t-channel production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, JHEP 04 (2017) 086. [arXiv:1609.03920](https://arxiv.org/abs/1609.03920), [doi:10.1007/JHEP04\(2017\)086](https://doi.org/10.1007/JHEP04(2017)086).
- [389] R. Aaij, et al., Prompt charm production in pp collisions at $\sqrt{s}=7$ TeV, Nucl. Phys. B871 (2013) 1–20. [arXiv:1302.2864](https://arxiv.org/abs/1302.2864), [doi:10.1016/j.nuclphysb.2013.02.010](https://doi.org/10.1016/j.nuclphysb.2013.02.010).
- [390] R. Aaij, et al., Measurements of prompt charm production cross-sections in pp collisions at $\sqrt{s} = 5$ TeV, JHEP 06 (2017) 147. [arXiv:1610.02230](https://arxiv.org/abs/1610.02230), [doi:10.1007/JHEP06\(2017\)147](https://doi.org/10.1007/JHEP06(2017)147).
- [391] R. Aaij, et al., Measurements of prompt charm production cross-sections in pp collisions at $\sqrt{s} = 13$ TeV, JHEP 03 (2016) 159, [Erratum: JHEP05,074(2017)]. [arXiv:1510.01707](https://arxiv.org/abs/1510.01707), [doi:10.1007/JHEP03\(2016\)159](https://doi.org/10.1007/JHEP03(2016)159), [doi:10.1007/JHEP09\(2016\)013](https://doi.org/10.1007/JHEP09(2016)013), [doi:10.1007/JHEP05\(2017\)074](https://doi.org/10.1007/JHEP05(2017)074).
- [392] P. Skands, S. Carrazza, J. Rojo, Tuning PYTHIA 8.1: the Monash 2013 Tune, European Physical Journal 74 (2014) 3024. [arXiv:1404.5630](https://arxiv.org/abs/1404.5630), [doi:10.1140/epjc/s10052-014-3024-y](https://doi.org/10.1140/epjc/s10052-014-3024-y).
- [393] M. Cacciari, M. L. Mangano, P. Nason, Gluon PDF constraints from the ratio of forward heavy-quark production at the LHC at $\sqrt{s} = 7$ and 13 TeV, Eur. Phys. J. C75 (12) (2015) 610. [arXiv:1507.06197](https://arxiv.org/abs/1507.06197), [doi:10.1140/epjc/s10052-015-3814-x](https://doi.org/10.1140/epjc/s10052-015-3814-x).
- [394] O. Zenaiev, Charm Production and QCD Analysis at HERA and LHC, Eur. Phys. J. C77 (3) (2017) 151. [arXiv:1612.02371](https://arxiv.org/abs/1612.02371), [doi:10.3204/PUBDB-2017-01474](https://doi.org/10.3204/PUBDB-2017-01474), [doi:10.1140/epjc/s10052-017-4620-4](https://doi.org/10.1140/epjc/s10052-017-4620-4).

- [395] R. Gauld, Understanding forward B hadron production, JHEP 05 (2017) 084. [arXiv:1703.03636](#), [doi:10.1007/JHEP05\(2017\)084](#).
- [396] P. Nason, S. Dawson, R. K. Ellis, The Total Cross-Section for the Production of Heavy Quarks in Hadronic Collisions, Nucl. Phys. B303 (1988) 607–633. [doi:10.1016/0550-3213\(88\)90422-1](#).
- [397] P. Nason, S. Dawson, R. K. Ellis, The One Particle Inclusive Differential Cross-Section for Heavy Quark Production in Hadronic Collisions, Nucl. Phys. B327 (1989) 49–92, [Erratum: Nucl. Phys. B335,260(1990)]. [doi:10.1016/0550-3213\(90\)90180-L](#), [10.1016/0550-3213\(89\)90286-1](#).
- [398] M. Cacciari, S. Frixione, N. Houdeau, M. L. Mangano, P. Nason, G. Ridolfi, Theoretical predictions for charm and bottom production at the LHC, JHEP 10 (2012) 137. [arXiv:1205.6344](#), [doi:10.1007/JHEP10\(2012\)137](#).
- [399] T. Sjostrand, S. Mrenna, P. Z. Skands, A Brief Introduction to PYTHIA 8.1, Comput. Phys. Commun. 178 (2008) 852–867. [arXiv:0710.3820](#), [doi:10.1016/j.cpc.2008.01.036](#).
- [400] S. Alioli, P. Nason, C. Oleari, E. Re, A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX, JHEP 1006 (2010) 043. [arXiv:1002.2581](#), [doi:10.1007/JHEP06\(2010\)043](#).
- [401] R. Gauld, J. Rojo, L. Rottoli, J. Talbert, Charm production in the forward region: constraints on the small- x gluon and backgrounds for neutrino astronomy, JHEP 11 (2015) 009. [arXiv:1506.08025](#), [doi:10.1007/JHEP11\(2015\)009](#).
- [402] W. Stirling, E. Vryonidou, Charm production in association with an electroweak gauge boson at the LHC, Phys.Rev.Lett. 109 (2012) 082002. [arXiv:1203.6781](#), [doi:10.1103/PhysRevLett.109.082002](#).
- [403] U. Baur, F. Halzen, S. Keller, M. L. Mangano, K. Riesselmann, The Charm content of $W + 1$ jet events as a probe of the strange quark distribution function, Phys.Lett. B318 (1993) 544–548. [arXiv:hep-ph/9308370](#), [doi:10.1016/0370-2693\(93\)91553-Y](#).
- [404] S. Chatrchyan, et al., Measurement of associated $W +$ charm production in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 02 (2014) 013. [arXiv:1310.1138](#), [doi:10.1007/JHEP02\(2014\)013](#).
- [405] G. Aad, et al., Measurement of the production of a W boson in association with a charm quark in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, JHEP 1405 (2014) 068. [arXiv:1402.6263](#), [doi:10.1007/JHEP05\(2014\)068](#).
- [406] R. Aaij, et al., Study of W boson production in association with beauty and charm, Phys. Rev. D92 (5) (2015) 052001. [arXiv:1505.04051](#), [doi:10.1103/PhysRevD.92.052001](#).
- [407] S. P. Jones, A. D. Martin, M. G. Ryskin, T. Teubner, Probes of the small x gluon via exclusive J/ψ and Υ production at HERA and the LHC, JHEP 11 (2013) 085. [arXiv:1307.7099](#), [doi:10.1007/JHEP11\(2013\)085](#).
- [408] R. Aaij, et al., Updated measurements of exclusive J/ψ and $\psi(2S)$ production cross-sections in pp collisions at $\sqrt{s} = 7$ TeV, J. Phys. G41 (2014) 055002. [arXiv:1401.3288](#), [doi:10.1088/0954-3899/41/5/055002](#).
- [409] T. L. Collaboration, Central exclusive production of J/ψ and $\psi(2S)$ mesons in pp collisions at $\sqrt{s} = 13$ TeV. LHCb-CONF-2016-007, CERN-LHCb-CONF-2016-007.
- [410] B. B. Abelev, et al., Exclusive J/ψ photoproduction off protons in ultra-peripheral p-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, Phys. Rev. Lett. 113 (23) (2014) 232504. [arXiv:1406.7819](#), [doi:10.1103/PhysRevLett.113.232504](#).
- [411] S. P. Jones, A. D. Martin, M. G. Ryskin, T. Teubner, Exclusive J/ψ and Υ photoproduction and the low x gluon, J. Phys. G43 (3) (2016) 035002. [arXiv:1507.06942](#), [doi:10.1088/0954-3899/43/3/035002](#).
- [412] D. Yu. Ivanov, A. Schafer, L. Szymanowski, G. Krasnikov, Exclusive photoproduction of a heavy vector meson in QCD, Eur. Phys. J. C34 (3) (2004) 297–316, [Erratum: Eur. Phys. J. C75,no.2,75(2015)]. [arXiv:hep-ph/0401131](#), [doi:10.1140/epjc/s2004-01712-x](#), [10.1140/epjc/s10052-015-3298-8](#).
- [413] S. P. Jones, A. D. Martin, M. G. Ryskin, T. Teubner, The exclusive J/ψ process at the LHC tamed to probe the low x gluon, Eur. Phys. J. C76 (11) (2016) 633. [arXiv:1610.02272](#), [doi:10.1140/epjc/s10052-016-4493-y](#).
- [414] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, et al., The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, JHEP 1407 (2014) 079. [arXiv:1405.0301](#), [doi:10.1007/JHEP07\(2014\)079](#).
- [415] S. J. Brodsky, G. R. Farrar, Scaling Laws at Large Transverse Momentum, Phys. Rev. Lett. 31 (1973) 1153–1156. [doi:10.1103/PhysRevLett.31.1153](#).
- [416] R. D. Ball, E. R. Nocera, J. Rojo, The asymptotic behaviour of parton distributions at small and large x , Eur. Phys. J. C76 (7) (2016) 383. [arXiv:1604.00024](#), [doi:10.1140/epjc/s10052-016-4240-4](#).
- [417] J. Pumplin, Parametrization dependence and $\Delta\chi^2$ in parton distribution fitting, Phys.Rev. D82 (2010) 114020. [arXiv:0909.5176](#), [doi:10.1103/PhysRevD.82.114020](#).
- [418] A. I. Signal, A. W. Thomas, Possible Strength of the Nonperturbative Strange Sea of the Nucleon, Phys. Lett. B191 (1987) 205. [doi:10.1016/0370-2693\(87\)91348-7](#).
- [419] S. Catani, D. de Florian, G. Rodrigo, W. Vogelsang, Perturbative generation of a strange-quark asymmetry in the nucleon, Phys. Rev. Lett. 93 (2004) 152003. [arXiv:hep-ph/0404240](#), [doi:10.1103/PhysRevLett.93.152003](#).
- [420] R. D. Ball, et al., Fitting Parton Distribution Data with Multiplicative Normalization Uncertainties, JHEP 05 (2010) 075. [arXiv:0912.2276](#), [doi:10.1007/JHEP05\(2010\)075](#).

- [421] F. James, M. Roos, Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations, *Comput. Phys. Commun.* 10 (1975) 343–367. doi:10.1016/0010-4655(75)90039-9.
- [422] J. Pumplin, D. R. Stump, W. K. Tung, Multivariate fitting and the error matrix in global analysis of data, *Phys. Rev. D* 65 (2001) 014011. arXiv:hep-ph/0008191, doi:10.1103/PhysRevD.65.014011.
- [423] J. Rojo, J. I. Latorre, Neural network parametrization of spectral functions from hadronic tau decays and determination of qcd vacuum condensates, *JHEP* 01 (2004) 055. arXiv:hep-ph/0401047.
- [424] M. C. Gonzalez-Garcia, M. Maltoni, J. Rojo, Determination of the atmospheric neutrino fluxes from atmospheric neutrino data, *JHEP* 10 (2006) 075. arXiv:hep-ph/0607324.
- [425] B. C. Allanach, D. Grellscheid, F. Quevedo, Genetic algorithms and experimental discrimination of susy models, *JHEP* 07 (2004) 069. arXiv:hep-ph/0406277.
- [426] P. M. Nadolsky, Z. Sullivan, PDF uncertainties in WH production at Tevatron, eConf C010630 (2001) P510. arXiv:hep-ph/0110378.
- [427] L. Del Debbio, S. Forte, J. I. Latorre, A. Piccione, J. Rojo, Neural network determination of parton distributions: The nonsinglet case, *JHEP* 03 (2007) 039. arXiv:hep-ph/0701127.
- [428] L. Del Debbio, S. Forte, J. I. Latorre, A. Piccione, J. Rojo, Unbiased determination of the proton structure function $f_2(p)$ with estimation, *JHEP* 03 (2005) 080. arXiv:hep-ph/0501067.
- [429] S. Carrazza, S. Forte, Z. Kassabov, J. Rojo, Specialized minimal PDFs for optimized LHC calculations, *Eur. Phys. J. C* 76 (4) (2016) 205. arXiv:1602.00005, doi:10.1140/epjc/s10052-016-4042-8.
- [430] D. Stump, J. Pumplin, R. Brock, D. Casey, J. Huston, J. Kalk, H. L. Lai, W. K. Tung, Uncertainties of predictions from parton distribution functions. 1. The Lagrange multiplier method, *Phys. Rev. D* 65 (2001) 014012. arXiv:hep-ph/0101051, doi:10.1103/PhysRevD.65.014012.
- [431] J. Beringer, et al., Review of Particle Physics (RPP), *Phys. Rev. D* 86 (2012) 010001. doi:10.1103/PhysRevD.86.010001.
- [432] J. Pumplin, J. Huston, H. L. Lai, W.-K. Tung, C. P. Yuan, Collider Inclusive Jet Data and the Gluon Distribution, *Phys. Rev. D* 80 (2009) 014019. arXiv:0904.2424, doi:10.1103/PhysRevD.80.014019.
- [433] J. Gao, Applications of the META parton distribution functions, in: Proceedings, 49th Rencontres de Moriond on QCD and High Energy Interactions: La Thuile, Italy, March 22-29, 2014, 2014, pp. 193–196. URL http://moriond.in2p3.fr/Proceedings/2014/Moriond_QCD_2014.pdf
- [434] S. Carrazza, Z. Kassabov, SMPDF Web: a web-based application for specialized minimal parton distribution functions, 2016. arXiv:1606.09248.
- [435] C. Patrignani, et al., Review of Particle Physics, *Chin. Phys.* C40 (10) (2016) 100001. doi:10.1088/1674-1137/40/10/100001.
- [436] A. D. Martin, W. J. Stirling, R. S. Thorne, G. Watt, Uncertainties on α_s in global PDF analyses, *Eur. Phys. J. C* 64 (2009) 653–680. arXiv:0905.3531, doi:10.1140/epjc/s10052-009-1164-2.
- [437] R. D. Ball, V. Bertone, L. Del Debbio, S. Forte, A. Guffanti, et al., Precision NNLO determination of $\alpha_s(M_Z)$ using an unbiased global parton set, *Phys. Lett. B* 707 (2012) 66–71. arXiv:1110.2483, doi:10.1016/j.physletb.2011.11.053.
- [438] S. Lionetti, et al., Precision determination of α_s using an unbiased global NLO parton set, *Phys. Lett. B* 701 (2011) 346–352. arXiv:1103.2369, doi:10.1016/j.physletb.2011.05.071.
- [439] H.-L. Lai, et al., Uncertainty induced by QCD coupling in the CTEQ global analysis of parton distributions, *Phys. Rev. D* 82 (2010) 054021. arXiv:1004.4624, doi:10.1103/PhysRevD.82.054021.
- [440] L. A. Harland-Lang, A. D. Martin, P. Motylinski, R. S. Thorne, Charm and beauty quark masses in the MMHT2014 global PDF analysis, *Eur. Phys. J. C* 76 (1) (2016) 10. arXiv:1510.02332, doi:10.1140/epjc/s10052-015-3843-5.
- [441] A. H. Hoang, A. V. Manohar, Charm quark mass from inclusive semileptonic B decays, *Phys. Lett. B* 633 (2006) 526–532. arXiv:hep-ph/0509195, doi:10.1016/j.physletb.2005.12.020.
- [442] R. D. Ball, V. Bertone, F. Cerutti, L. Del Debbio, S. Forte, et al., Reweighting and Unweighting of Parton Distributions and the LHC W lepton asymmetry data, *Nucl. Phys. B* 855 (2012) 608–638. arXiv:1108.1758, doi:10.1016/j.nuclphysb.2011.10.018.
- [443] R. D. Ball, et al., Reweighting NNPDFs: the W lepton asymmetry, *Nucl. Phys. B* 849 (2011) 112–143. arXiv:1012.0836, doi:10.1016/j.nuclphysb.2011.03.017.
- [444] S. Camarda, et al., QCD analysis of W- and Z-boson production at Tevatron, *Eur. Phys. J. C* 75 (9) (2015) 458. arXiv:1503.05221, doi:10.1140/epjc/s10052-015-3655-7.
- [445] N. Armesto, J. Rojo, C. A. Salgado, P. Zurita, Bayesian reweighting of nuclear PDFs and constraints from proton-lead collisions at the LHC, *JHEP* 11 (2013) 015. arXiv:1309.5371, doi:10.1007/JHEP11(2013)015.
- [446] W. T. Giele, S. Keller, Implications of hadron collider observables on parton distribution function uncertainties, *Phys. Rev. D* 58 (1998) 094023. arXiv:hep-ph/9803393, doi:10.1103/PhysRevD.58.094023.
- [447] N. Sato, J. F. Owens, H. Prosper, Bayesian Reweighting for Global Fits, *Phys. Rev. D* 89 (11) (2014) 114020. arXiv:

- 1310.1089, doi:10.1103/PhysRevD.89.114020.
- [448] H. Paukkunen, P. Zurita, PDF reweighting in the Hessian matrix approach, JHEP 12 (2014) 100. [arXiv:1402.6623](#), doi:10.1007/JHEP12(2014)100.
- [449] H. Plathow-Besch, PDFLIB: A Library of all available parton density functions of the nucleon, the pion and the photon and the corresponding α_s calculations, Comput. Phys. Commun. 75 (1993) 396–416. doi:10.1016/0010-4655(93)90051-D.
- [450] D. Bourilkov, R. C. Group, M. R. Whalley, LHAPDF: PDF use from the Tevatron to the LHC, in: TeV4LHC Workshop - 4th meeting Batavia, Illinois, October 20-22, 2005, 2006. [arXiv:hep-ph/0605240](#).
- [451] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, et al., LHAPDF6: parton density access in the LHC precision era, Eur.Phys.J. C75 (2015) 132. [arXiv:1412.7420](#), doi:10.1140/epjc/s10052-015-3318-8.
- [452] T. Gleisberg, et al., Event generation with SHERPA 1.1, JHEP 02 (2009) 007. [arXiv:0811.4622](#), doi:10.1088/1126-6708/2009/02/007.
- [453] V. Bertone, S. Carrazza, N. P. Hartland, E. R. Nocera, J. Rojo, A determination of the fragmentation functions of pions, kaons, and protons with faithful uncertainties, Eur. Phys. J. C77 (8) (2017) 516. [arXiv:1706.07049](#), doi:10.1140/epjc/s10052-017-5088-y.
- [454] J. P. Berge, et al., A Measurement of Differential Cross-Sections and Nucleon Structure Functions in Charged Current Neutrino Interactions on Iron, Z. Phys. C49 (1991) 187–224. doi:10.1007/BF01555493.
- [455] A. Aktas, et al., Measurement of $F_2(c\bar{c})$ and $F_2(b\bar{b})$ at high Q^2 using the H1 vertex detector at HERA, Eur. Phys. J. C40 (2005) 349–359. [arXiv:hep-ex/0411046](#), doi:10.1140/epjc/s2005-02154-8.
- [456] G. Moreno, et al., Dimuon production in proton - copper collisions at $\sqrt{s} = 38.8$ -GeV, Phys. Rev. D43 (1991) 2815–2836. doi:10.1103/PhysRevD.43.2815.
- [457] F. Abe, et al., Forward-backward charge asymmetry of electron pairs above the Z^0 pole, Phys. Rev. Lett. 77 (1996) 2616–2621. doi:10.1103/PhysRevLett.77.2616.
- [458] D. E. Acosta, et al., Measurement of the forward-backward charge asymmetry from $W \rightarrow e\nu$ production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. D71 (2005) 051104. [arXiv:hep-ex/0501023](#), doi:10.1103/PhysRevD.71.051104.
- [459] V. M. Abazov, et al., Measurement of the muon charge asymmetry from W boson decays, Phys. Rev. D77 (2008) 011106. [arXiv:0709.4254](#), doi:10.1103/PhysRevD.77.011106.
- [460] V. M. Abazov, et al., Measurement of the ratios of the $Z/\gamma^* + \geq n$ jet production cross sections to the total inclusive Z/γ^* cross section in p anti- p collisions at $\sqrt{s} = 1.96$ -TeV, Phys. Lett. B658 (2008) 112–119. [arXiv:hep-ex/0608052](#), doi:10.1016/j.physletb.2007.10.046.
- [461] C. Balazs, C. P. Yuan, Soft gluon effects on lepton pairs at hadron colliders, Phys. Rev. D56 (1997) 5558–5583. [arXiv:hep-ph/9704258](#), doi:10.1103/PhysRevD.56.5558.
- [462] M. Guzzi, P. M. Nadolsky, B. Wang, Nonperturbative contributions to a resummed leptonic angular distribution in inclusive neutral vector boson production, Phys. Rev. D90 (1) (2014) 014030. [arXiv:1309.1393](#), doi:10.1103/PhysRevD.90.014030.
- [463] H.-L. Lai, et al., New parton distributions for collider physics, Phys. Rev. D82 (2010) 074024. [arXiv:1007.2241](#), doi:10.1103/PhysRevD.82.074024.
- [464] J. F. Owens, A. Accardi, W. Melnitchouk, Global parton distributions with nuclear and finite- Q^2 corrections, Phys. Rev. D87 (9) (2013) 094012. [arXiv:1212.1702](#), doi:10.1103/PhysRevD.87.094012.
- [465] R. Aaij, et al., Inclusive W and Z production in the forward region at $\sqrt{s} = 7$ TeV, JHEP 1206 (2012) 058. [arXiv:1204.1620](#), doi:10.1007/JHEP06(2012)058.
- [466] S. Dulat, et al., Intrinsic Charm Parton Distribution Functions from CTEQ-TEA Global Analysis, Phys.Rev. D89 (7) (2014) 073004. [arXiv:1309.0025](#), doi:10.1103/PhysRevD.89.073004.
- [467] T.-J. Hou, S. Dulat, J. Gao, M. Guzzi, J. W. Huston, P. Nadolsky, J. Pumplin, C. R. Schmidt, D. Stump, C. P. Yuan, Heavy Flavors on CT14, PoS DIS2015 (2015) 166.
- [468] S. Chekanov, et al., Measurement of isolated photon production in deep inelastic ep scattering, Phys.Lett. B687 (2010) 16–25. [arXiv:0909.4223](#), doi:10.1016/j.physletb.2010.02.045.
- [469] A. D. Martin, R. G. Roberts, W. J. Stirling, R. S. Thorne, Parton distributions: A New global analysis, Eur. Phys. J. C4 (1998) 463–496. [arXiv:hep-ph/9803445](#), doi:10.1007/s100529800904, 10.1007/s100520050220.
- [470] A. D. Martin, R. G. Roberts, W. J. Stirling, R. S. Thorne, Physical gluons and high-E(T) jets, Phys. Lett. B604 (2004) 61–68. [arXiv:hep-ph/0410230](#), doi:10.1016/j.physletb.2004.10.040.
- [471] A. D. Martin, W. J. Stirling, R. S. Thorne, G. Watt, Update of Parton Distributions at NNLO, Phys. Lett. B652 (2007) 292–299. [arXiv:0706.0459](#), doi:10.1016/j.physletb.2007.07.040.
- [472] T. Bolton, Determining the CKM parameter V_{cd} from νN charm production. [arXiv:hep-ex/9708014](#).
- [473] L. A. Harland-Lang, A. D. Martin, P. Motylinski, R. S. Thorne, Uncertainties on α_s in the MMHT2014 global PDF analysis and implications for SM predictions, Eur. Phys. J. C75 (9) (2015) 435. [arXiv:1506.05682](#), doi:10.1140/epjc/s10052-015-3630-3.

- [474] R. Thorne, Effect of changes of variable flavor number scheme on parton distribution functions and predicted cross sections, *Phys.Rev. D*86 (2012) 074017. [arXiv:1201.6180](#), [doi:10.1103/PhysRevD.86.074017](#).
- [475] R. S. Thorne, L. A. Harland-Lang, A. D. Martin, P. Motylinski, The Effect of Final HERA inclusive Cross Section Data MMHT2014 PDFs, in: Proceedings, 2015 European Physical Society Conference on High Energy Physics (EPS-HEP 2015), 2015. [arXiv:1508.06621](#).
- [476] L. A. Harland-Lang, R. Nathvani, R. S. Thorne, A. D. Martin, MMHT PDFs: updates and outlook, *Acta Phys. Polon. B*48 (2017) 1011–1024. [arXiv:1704.00162](#), [doi:10.5506/APhysPolB.48.1011](#).
- [477] S. Forte, J. I. Latorre, L. Magnea, A. Piccione, Determination of α_s from scaling violations of truncated moments of structure functions, *Nucl. Phys. B*643 (2002) 477–500. [arXiv:hep-ph/0205286](#), [doi:10.1016/S0550-3213\(02\)00688-0](#).
- [478] R. D. Ball, et al., A determination of parton distributions with faithful uncertainty estimation, *Nucl. Phys. B*809 (2009) 1–63. [arXiv:0808.1231](#), [doi:10.1016/j.nuclphysb.2008.09.037](#).
- [479] S. Davidson, S. Forte, P. Gambino, N. Rius, A. Strumia, Old and new physics interpretations of the NuTeV anomaly, *JHEP* 02 (2002) 037. [arXiv:hep-ph/0112302](#).
- [480] T. Sjstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, P. Z. Skands, An Introduction to PYTHIA 8.2, *Comput. Phys. Commun.* 191 (2015) 159–177. [arXiv:1410.3012](#), [doi:10.1016/j.cpc.2015.01.024](#).
- [481] R. D. Ball, et al., Theoretical issues in PDF determination and associated uncertainties, *Phys.Lett. B*723 (2013) 330. [arXiv:1303.1189](#), [doi:10.1016/j.physletb.2013.05.019](#).
- [482] L. Rottoli, Fitting EMC structure functions with intrinsic charm, *PoS DIS2016* (2016) 032. [arXiv:1606.09289](#).
- [483] S. J. Brodsky, P. Hoyer, C. Peterson, N. Sakai, The Intrinsic Charm of the Proton, *Phys. Lett. B*93 (1980) 451–455. [doi:10.1016/0370-2693\(80\)90364-0](#).
- [484] S. Alekhin, J. Blümlein, S. Moch, Parton Distribution Functions and Benchmark Cross Sections at NNLO, *Phys.Rev. D*86 (2012) 054009. [arXiv:1202.2281](#), [doi:10.1103/PhysRevD.86.054009](#).
- [485] H. Georgi, H. D. Politzer, Freedom at Moderate Energies: Masses in Color Dynamics, *Phys. Rev. D*14 (1976) 1829. [doi:10.1103/PhysRevD.14.1829](#).
- [486] L. T. Brady, A. Accardi, T. J. Hobbs, W. Melnitchouk, Next-to leading order analysis of target mass corrections to structure functions and asymmetries, *Phys. Rev. D*84 (2011) 074008, [Erratum: *Phys. Rev.D*85,039902(2012)]. [arXiv:1108.4734](#), [doi:10.1103/PhysRevD.84.074008](#), [doi:10.1103/PhysRevD.85.039902](#).
- [487] W. Melnitchouk, A. W. Schreiber, A. W. Thomas, Deep inelastic scattering from off-shell nucleons, *Phys. Rev. D*49 (1994) 1183–1198. [arXiv:nucl-th/9311008](#), [doi:10.1103/PhysRevD.49.1183](#).
- [488] S. A. Kulagin, G. Piller, W. Weise, Shadowing, binding and off-shell effects in nuclear deep inelastic scattering, *Phys. Rev. C*50 (1994) 1154–1169. [arXiv:nucl-th/9402015](#), [doi:10.1103/PhysRevC.50.1154](#).
- [489] S. A. Kulagin, R. Petti, Global study of nuclear structure functions, *Nucl. Phys. A*765 (2006) 126–187. [arXiv:hep-ph/0412425](#), [doi:10.1016/j.nuclphysa.2005.10.011](#).
- [490] Y. Kahn, W. Melnitchouk, S. A. Kulagin, New method for extracting neutron structure functions from nuclear data, *Phys. Rev. C*79 (2009) 035205. [arXiv:0809.4308](#), [doi:10.1103/PhysRevC.79.035205](#).
- [491] R. B. Wiringa, V. G. J. Stoks, R. Schiavilla, An Accurate nucleon-nucleon potential with charge independence breaking, *Phys. Rev. C*51 (1995) 38–51. [arXiv:nucl-th/9408016](#), [doi:10.1103/PhysRevC.51.38](#).
- [492] R. Machleidt, The High precision, charge dependent Bonn nucleon-nucleon potential (CD-Bonn), *Phys. Rev. C*63 (2001) 024001. [arXiv:nucl-th/0006014](#), [doi:10.1103/PhysRevC.63.024001](#).
- [493] F. Gross, A. Stadler, Covariant spectator theory of np scattering: Phase shifts obtained from precision fits to data below 350-MeV, *Phys. Rev. C*78 (2008) 014005. [arXiv:0802.1552](#), [doi:10.1103/PhysRevC.78.014005](#).
- [494] N. Baillie, et al., Measurement of the neutron F2 structure function via spectator tagging with CLAS, *Phys. Rev. Lett.* 108 (2012) 142001, [Erratum: *Phys. Rev. Lett.*108,199902(2012)]. [arXiv:1110.2770](#), [doi:10.1103/PhysRevLett.108.199902](#), [doi:10.1103/PhysRevLett.108.142001](#).
- [495] P. Belov, et al., Parton distribution functions at LO, NLO and NNLO with correlated uncertainties between orders, *Eur. Phys. J. C*74 (10) (2014) 3039. [arXiv:1404.4234](#), [doi:10.1140/epjc/s10052-014-3039-4](#).
- [496] F. Hautmann, H. Jung, A. Lelek, V. Radescu, R. Zlebcik, Soft-gluon resolution scale in QCD evolution equations, *Phys. Lett. B*772 (2017) 446–451. [arXiv:1704.01757](#), [doi:10.1016/j.physletb.2017.07.005](#).
- [497] V. Bertone, et al., Impact of the heavy quark matching scales in PDF fits. [arXiv:1707.05343](#).
- [498] C. Collaboration, Measurement of Triple-Differential Dijet Cross Sections at $\sqrt{s} = 8$ TeV with the CMS Detector and Constraints on Parton Distribution Functions. CMS-PAS-SMP-16-011.
- [499] C. Collaboration, Measurement of the inclusive $t\bar{t}$ cross section at $\sqrt{s} = 5.02$ TeV. CMS-PAS-TOP-16-023.
- [500] S. Carrazza, A. Ferrara, D. Palazzo, J. Rojo, APFEL Web: a web-based application for the graphical visualization of parton distribution functions, *J.Phys. G*42 (2015) 057001. [arXiv:1410.5456](#), [doi:10.1088/0954-3899/42/5/057001](#).
- [501] I. Borsa, R. Sassot, M. Stratmann, Probing the Sea Quark Content of the Proton with One-Particle-Inclusive Processes.

[arXiv:1708.01630](#).

- [502] S. J. Brodsky, A. Kusina, F. Lyonnet, I. Schienbein, H. Spiesberger, R. Vogt, A review of the intrinsic heavy quark content of the nucleon, *Adv. High Energy Phys.* 2015 (2015) 231547. [arXiv:1504.06287](#), [doi:10.1155/2015/231547](#).
- [503] J. Pumplin, H. L. Lai, W. K. Tung, The charm parton content of the nucleon, *Phys. Rev. D* 75 (2007) 054029. [arXiv:hep-ph/0701220](#), [doi:10.1103/PhysRevD.75.054029](#).
- [504] P. Jimenez-Delgado, T. Hobbs, J. Londergan, W. Melnitchouk, New limits on intrinsic charm in the nucleon from global analysis of parton distributions, *Phys.Rev.Lett.* 114 (8) (2015) 082002. [arXiv:1408.1708](#), [doi:10.1103/PhysRevLett.114.082002](#).
- [505] T. Boettcher, P. Ilten, M. Williams, Direct probe of the intrinsic charm content of the proton, *Phys. Rev. D* 93 (7) (2016) 074008. [arXiv:1512.06666](#), [doi:10.1103/PhysRevD.93.074008](#).
- [506] G. Bailas, V. P. Goncalves, Phenomenological implications of the intrinsic charm in the Z boson production at the LHC, *Eur. Phys. J. C* 76 (3) (2016) 105. [arXiv:1512.06007](#), [doi:10.1140/epjc/s10052-016-3941-z](#).
- [507] P.-H. Beauchemin, V. A. Bednyakov, G. I. Lykasov, Yu. Yu. Stepanenko, Search for intrinsic charm in vector boson production accompanied by heavy flavor jets, *Phys. Rev. D* 92 (3) (2015) 034014. [arXiv:1410.2616](#), [doi:10.1103/PhysRevD.92.034014](#).
- [508] T. P. Stavreva, J. F. Owens, Direct Photon Production in Association With A Heavy Quark At Hadron Colliders, *Phys. Rev. D* 79 (2009) 054017. [arXiv:0901.3791](#), [doi:10.1103/PhysRevD.79.054017](#).
- [509] V. A. Bednyakov, M. A. Demichev, G. I. Lykasov, T. Stavreva, M. Stockton, Searching for intrinsic charm in the proton at the LHC, *Phys. Lett. B* 728 (2014) 602–606. [arXiv:1305.3548](#), [doi:10.1016/j.physletb.2013.12.031](#).
- [510] R. Vogt, S. J. Brodsky, QCD and intrinsic heavy quark predictions for leading charm and beauty hadroproduction, *Nucl. Phys. B* 438 (1995) 261–277. [arXiv:hep-ph/9405236](#), [doi:10.1016/0550-3213\(94\)00543-N](#).
- [511] B. A. Kniehl, G. Kramer, I. Schienbein, H. Spiesberger, Inclusive Charmed-Meson Production at the CERN LHC, *Eur. Phys. J. C* 72 (2012) 2082. [arXiv:1202.0439](#), [doi:10.1140/epjc/s10052-012-2082-2](#).
- [512] M. Gluck, E. Reya, I. Schienbein, Radiatively generated parton distributions of real and virtual photons, *Phys. Rev. D* 60 (1999) 054019, [Erratum: *Phys. Rev. D* 62, 019902 (2000)]. [arXiv:hep-ph/9903337](#), [doi:10.1103/PhysRevD.60.054019](#), [doi:10.1103/PhysRevD.62.019902](#).
- [513] V. Bertone, S. Carrazza, D. Pagani, M. Zaro, On the Impact of Lepton PDFs, *JHEP* 11 (2015) 194. [arXiv:1508.07002](#), [doi:10.1007/JHEP11\(2015\)194](#).
- [514] D. de Florian, G. F. R. Sborlini, G. Rodrigo, QED corrections to the Altarelli-Parisi splitting functions, *Eur. Phys. J. C* 76 (5) (2016) 282. [arXiv:1512.00612](#), [doi:10.1140/epjc/s10052-016-4131-8](#).
- [515] D. de Florian, G. F. R. Sborlini, G. Rodrigo, Two-loop QED corrections to the Altarelli-Parisi splitting functions, *JHEP* 10 (2016) 056. [arXiv:1606.02887](#), [doi:10.1007/JHEP10\(2016\)056](#).
- [516] R. Sadykov, Impact of QED radiative corrections on Parton Distribution Functions. [arXiv:1401.1133](#).
- [517] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, Sudakov effects in photon-initiated processes, *Phys. Lett. B* 761 (2016) 20–24. [arXiv:1605.04935](#), [doi:10.1016/j.physletb.2016.08.004](#).
- [518] S. Kallweit, J. M. Lindert, S. Pozzorini, M. Schonherr, NLO QCD+EW predictions for $2\ell 2\nu$ diboson signatures at the LHC. [arXiv:1705.00598](#).
- [519] J. C. Bernauer, et al., Electric and magnetic form factors of the proton, *Phys. Rev. C* 90 (1) (2014) 015206. [arXiv:1307.6227](#), [doi:10.1103/PhysRevC.90.015206](#).
- [520] V. M. Budnev, I. F. Ginzburg, G. V. Meledin, V. G. Serbo, The Two photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation, *Phys. Rept.* 15 (1975) 181–281. [doi:10.1016/0370-1573\(75\)90009-5](#).
- [521] M. Gluck, C. Pisano, E. Reya, The Polarized and unpolarized photon content of the nucleon, *Phys. Lett. B* 540 (2002) 75–80. [arXiv:hep-ph/0206126](#), [doi:10.1016/S0370-2693\(02\)02125-1](#).
- [522] A. Martin, M. Ryskin, The photon PDF of the proton, *Eur.Phys.J. C* 74 (2014) 3040. [arXiv:1406.2118](#), [doi:10.1140/epjc/s10052-014-3040-y](#).
- [523] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, The production of a diphoton resonance via photon-photon fusion, *JHEP* 03 (2016) 182. [arXiv:1601.07187](#), [doi:10.1007/JHEP03\(2016\)182](#).
- [524] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, The photon PDF in events with rapidity gaps, *Eur. Phys. J. C* 76 (5) (2016) 255. [arXiv:1601.03772](#), [doi:10.1140/epjc/s10052-016-4100-2](#).
- [525] H. Anlauf, H. D. Dahmen, P. Manakos, T. Mannel, T. Ohl, KRONOS: A Monte Carlo event generator for higher order electromagnetic radiative corrections to deep inelastic scattering at HERA, *Comput. Phys. Commun.* 70 (1992) 97–119. [doi:10.1016/0010-4655\(92\)90095-G](#).
- [526] A. Mukherjee, C. Pisano, Manifestly covariant analysis of the QED Compton process in $e p \rightarrow e \gamma p$ and $e p \rightarrow e \gamma X$, *Eur. Phys. J. C* 30 (2003) 477–486. [arXiv:hep-ph/0306275](#), [doi:10.1140/epjc/s2003-01308-0](#).
- [527] V. Bertone, S. Carrazza, N. P. Hartland, J. Rojo, Illuminating the photon content of the proton within a global PDF analy-

sisarXiv:1712.07053.

- [528] G. Aad, et al., Measurement of the double-differential high-mass Drell-Yan cross section in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, JHEP 08 (2016) 009. [arXiv:1606.01736](#), [doi:10.1007/JHEP08\(2016\)009](#).
- [529] D. Bourilkov, Photon-induced Background for Dilepton Searches and Measurements in pp Collisions at 13 TeV. [arXiv:1606.00523](#).
- [530] E. Accomando, J. Fiaschi, F. Hautmann, S. Moretti, C. H. Shepherd-Themistocleous, Photon-initiated production of a dilepton final state at the LHC: Cross section versus forward-backward asymmetry studies, Phys. Rev. D95 (3) (2017) 035014. [arXiv:1606.06646](#), [doi:10.1103/PhysRevD.95.035014](#).
- [531] K. Mishra, et al., Electroweak Corrections at High Energies, in: Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013, 2013. [arXiv:1308.1430](#).
- [532] J. M. Campbell, D. Wackerth, J. Zhou, Study of weak corrections to Drell-Yan, top-quark pair, and dijet production at high energies with MCFM, Phys. Rev. D94 (9) (2016) 093009. [arXiv:1608.03356](#), [doi:10.1103/PhysRevD.94.093009](#).
- [533] S. Frixione, V. Hirschi, D. Pagani, H. S. Shao, M. Zaro, Electroweak and QCD corrections to top-pair hadroproduction in association with heavy bosons, JHEP 06 (2015) 184. [arXiv:1504.03446](#), [doi:10.1007/JHEP06\(2015\)184](#).
- [534] S. Kallweit, J. M. Lindert, P. Maierhofer, S. Pozzorini, M. Schnherr, NLO electroweak automation and precise predictions for W+multijet production at the LHC, JHEP 04 (2015) 012. [arXiv:1412.5157](#), [doi:10.1007/JHEP04\(2015\)012](#).
- [535] **LHC Higgs Cross Section Working Group**, S. Dittmaier, et al., Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables. [arXiv:1101.0593](#).
- [536] **LHC Higgs Cross Section Working Group**, S. Dittmaier, et al., Handbook of LHC Higgs Cross Sections: 2. Differential Distributions. [arXiv:1201.3084](#).
- [537] S. Dawson, et al., Working Group Report: Higgs Boson, in: Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013, 2013. [arXiv:1310.8361](#).
- [538] J. K. Behr, D. Bortoletto, J. A. Frost, N. P. Hartland, C. Issever, J. Rojo, Boosting Higgs pair production in the $b\bar{b}b\bar{b}$ final state with multivariate techniques, Eur. Phys. J. C76 (7) (2016) 386. [arXiv:1512.08928](#), [doi:10.1140/epjc/s10052-016-4215-5](#).
- [539] A. Papaefstathiou, L. L. Yang, J. Zurita, Higgs boson pair production at the LHC in the $b\bar{b}W^+W^-$ channel, Phys. Rev. D87 (1) (2013) 011301. [arXiv:1209.1489](#), [doi:10.1103/PhysRevD.87.011301](#).
- [540] M. J. Dolan, C. Englert, M. Spannowsky, New Physics in LHC Higgs boson pair production, Phys. Rev. D87 (5) (2013) 055002. [arXiv:1210.8166](#), [doi:10.1103/PhysRevD.87.055002](#).
- [541] F. Bishara, R. Contino, J. Rojo, Higgs pair production in vector-boson fusion at the LHC and beyond, Eur. Phys. J. C77 (7) (2017) 481. [arXiv:1611.03860](#), [doi:10.1140/epjc/s10052-017-5037-9](#).
- [542] R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, P. Torrielli, E. Vryonidou, M. Zaro, Higgs pair production at the LHC with NLO and parton-shower effects, Phys. Lett. B732 (2014) 142–149. [arXiv:1401.7340](#), [doi:10.1016/j.physletb.2014.03.026](#).
- [543] R. D. Ball, M. Bonvini, S. Forte, S. Marzani, G. Ridolfi, Higgs production in gluon fusion beyond NNLO, Nucl.Phys. B874 (2013) 746. [arXiv:1303.3590](#), [doi:10.1016/j.nuclphysb.2013.06.012](#).
- [544] F. A. Dreyer, A. Karlberg, Vector-Boson Fusion Higgs Production at Three Loops in QCD, Phys. Rev. Lett. 117 (7) (2016) 072001. [arXiv:1606.00840](#), [doi:10.1103/PhysRevLett.117.072001](#).
- [545] G. Aad, et al., Measurement of fiducial differential cross sections of gluon-fusion production of Higgs bosons decaying to $WW^* \rightarrow e\nu\mu\nu$ with the ATLAS detector at $\sqrt{s} = 8$ TeV, JHEP 08 (2016) 104. [arXiv:1604.02997](#), [doi:10.1007/JHEP08\(2016\)104](#).
- [546] V. Khachatryan, et al., Measurement of differential and integrated fiducial cross sections for Higgs boson production in the four-lepton decay channel in pp collisions at $\sqrt{s} = 7$ and 8 TeV, JHEP 04 (2016) 005. [arXiv:1512.08377](#), [doi:10.1007/JHEP04\(2016\)005](#).
- [547] L. Randall, R. Sundrum, An Alternative to compactification, Phys.Rev.Lett. 83 (1999) 4690–4693. [arXiv:hep-th/9906064](#), [doi:10.1103/PhysRevLett.83.4690](#).
- [548] L. Randall, R. Sundrum, A Large mass hierarchy from a small extra dimension, Phys.Rev.Lett. 83 (1999) 3370–3373. [arXiv:hep-ph/9905221](#), [doi:10.1103/PhysRevLett.83.3370](#).
- [549] J. Alwall, et al., MadGraph/MadEvent v4: The New Web Generation, JHEP 09 (2007) 028. [arXiv:0706.2334](#).
- [550] M. Aaboud, et al., Search for heavy Higgs bosons A/H decaying to a top quark pair in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. [arXiv:1707.06025](#).
- [551] I. Brivio, M. Trott, The Standard Model as an Effective Field Theory. [arXiv:1706.08945](#).
- [552] E. L. Berger, M. Guzzi, H.-L. Lai, P. M. Nadolsky, F. I. Olness, Constraints on color-octet fermions from a global parton distribution analysis, Phys. Rev. D82 (2010) 114023. [arXiv:1010.4315](#), [doi:10.1103/PhysRevD.82.114023](#).

- [553] M. Baak, J. Cth, J. Haller, A. Hoecker, R. Kogler, K. Mnig, M. Schott, J. Stelzer, The global electroweak fit at NNLO and prospects for the LHC and ILC, *Eur. Phys. J. C* 74 (2014) 3046. [arXiv:1407.3792](#), [doi:10.1140/epjc/s10052-014-3046-5](#).
- [554] M. Benayoun, P. David, L. DelBuono, F. Jegerlehner, Muon $g - 2$ estimates: can one trust effective Lagrangians and global fits?, *Eur. Phys. J. C* 75 (12) (2015) 613. [arXiv:1507.02943](#), [doi:10.1140/epjc/s10052-015-3830-x](#).
- [555] D. S. M. Alves, J. Galloway, J. T. Ruderman, J. R. Walsh, Running Electroweak Couplings as a Probe of New Physics, *JHEP* 02 (2015) 007. [arXiv:1410.6810](#), [doi:10.1007/JHEP02\(2015\)007](#).
- [556] M. Mojaza, S. J. Brodsky, X.-G. Wu, Systematic All-Orders Method to Eliminate Renormalization-Scale and Scheme Ambiguities in Perturbative QCD, *Phys. Rev. Lett.* 110 (2013) 192001. [arXiv:1212.0049](#), [doi:10.1103/PhysRevLett.110.192001](#).
- [557] F. I. Olness, D. E. Soper, Correlated theoretical uncertainties for the one-jet inclusive cross section, *Phys. Rev. D* 81 (2010) 035018. [arXiv:0907.5052](#), [doi:10.1103/PhysRevD.81.035018](#).
- [558] M. Cacciari, N. Houdeau, Meaningful characterisation of perturbative theoretical uncertainties, *JHEP* 1109 (2011) 039. [arXiv:1105.5152](#), [doi:10.1007/JHEP09\(2011\)039](#).
- [559] E. Bagnaschi, M. Cacciari, A. Guffanti, L. Jenniches, An extensive survey of the estimation of uncertainties from missing higher orders in perturbative calculations, *JHEP* 02 (2015) 133. [arXiv:1409.5036](#), [doi:10.1007/JHEP02\(2015\)133](#).
- [560] A. David, G. Passarino, How well can we guess theoretical uncertainties?, *Phys. Lett. B* 726 (2013) 266–272. [arXiv:1307.1843](#), [doi:10.1016/j.physletb.2013.08.025](#).
- [561] S. Forte, A. Isgur, G. Vita, Do we need N^3 LO Parton Distributions?, *Phys.Lett. B* 731 (2014) 136–140. [arXiv:1312.6688](#), [doi:10.1016/j.physletb.2014.02.027](#).
- [562] R. Gupta, Introduction to lattice QCD: Course, in: *Probing the standard model of particle interactions. Proceedings, Summer School in Theoretical Physics, NATO Advanced Study Institute, 68th session, Les Houches, France, July 28-September 5, 1997. Pt. 1, 2, 1997, pp. 83–219.* [arXiv:hep-lat/9807028](#).
- [563] H.-W. Lin, et al., Parton distributions and lattice QCD calculations: a community white paper [arXiv:1711.07916](#).
- [564] E. R. Nocera, H.-W. Lin, F. Olness, K. Orginos, J. Rojo, The PDFLattice2017 workshop: a summary report, 2017. [arXiv:1709.01511](#).
- [565] H.-W. Lin, Lattice Hadron Structure: Applications within and beyond QCD, *PoS LATTICE2012* (2012) 013. [arXiv:1212.6849](#).
- [566] M. Constantinou, Hadron Structure, *PoS LATTICE2014* (2015) 001. [arXiv:1411.0078](#).
- [567] S. Syritsyn, Review of Hadron Structure Calculations on a Lattice, *PoS LATTICE2013* (2014) 009. [arXiv:1403.4686](#).
- [568] H.-W. Lin, J.-W. Chen, S. D. Cohen, X. Ji, Flavor Structure of the Nucleon Sea from Lattice QCD, *Phys. Rev. D* 91 (2015) 054510. [arXiv:1402.1462](#), [doi:10.1103/PhysRevD.91.054510](#).
- [569] C. Alexandrou, K. Cichy, V. Drach, E. Garcia-Ramos, K. Hadjiyiannakou, K. Jansen, F. Steffens, C. Wiese, Lattice calculation of parton distributions, *Phys. Rev. D* 92 (2015) 014502. [arXiv:1504.07455](#), [doi:10.1103/PhysRevD.92.014502](#).
- [570] J.-W. Chen, S. D. Cohen, X. Ji, H.-W. Lin, J.-H. Zhang, Nucleon Helicity and Transversity Parton Distributions from Lattice QCD, *Nucl. Phys. B* 911 (2016) 246–273. [arXiv:1603.06664](#), [doi:10.1016/j.nuclphysb.2016.07.033](#).
- [571] C. Alexandrou, K. Cichy, M. Constantinou, K. Hadjiyiannakou, K. Jansen, F. Steffens, C. Wiese, Updated Lattice Results for Parton Distributions, *Phys. Rev. D* 96 (1) (2017) 014513. [arXiv:1610.03689](#), [doi:10.1103/PhysRevD.96.014513](#).
- [572] S. Aoki, et al., Review of lattice results concerning low-energy particle physics, *Eur. Phys. J. C* 77 (2) (2017) 112. [arXiv:1607.00299](#), [doi:10.1140/epjc/s10052-016-4509-7](#).
- [573] W. Detmold, W. Melnitchouk, A. W. Thomas, Extraction of parton distributions from lattice QCD, *Mod.Phys.Lett. A* 18 (2003) 2681–2698. [arXiv:hep-lat/0310003](#), [doi:10.1142/S0217732303012209](#).
- [574] H.-W. Lin, et al., Parton Distributions and Lattice QCD calculations. [arXiv:1710.aaaaa](#).
- [575] J. R. Green, M. Engelhardt, S. Krieg, J. W. Negele, A. V. Pochinsky, S. N. Syritsyn, Nucleon Structure from Lattice QCD Using a Nearly Physical Pion Mass, *Phys. Lett. B* 734 (2014) 290–295. [arXiv:1209.1687](#), [doi:10.1016/j.physletb.2014.05.075](#).
- [576] C. Alexandrou, M. Constantinou, K. Hadjiyiannakou, K. Jansen, C. Kallidonis, G. Koutsou, A. V. Avils-Casco, C. Wiese, The nucleon spin explained using lattice QCD simulations. [arXiv:1706.02973](#).
- [577] G. S. Bali, S. Collins, B. Gille, M. Gckeler, J. Najjar, R. H. Rdl, A. Schfer, R. W. Schiel, A. Sternbeck, W. Sldner, The moment $\langle x \rangle_{u-d}$ of the nucleon from $N_f = 2$ lattice QCD down to nearly physical quark masses, *Phys. Rev. D* 90 (7) (2014) 074510. [arXiv:1408.6850](#), [doi:10.1103/PhysRevD.90.074510](#).
- [578] X. Ji, Parton Physics on a Euclidean Lattice, *Phys.Rev.Lett.* 110 (26) (2013) 262002. [arXiv:1305.1539](#), [doi:10.1103/PhysRevLett.110.262002](#).
- [579] X. Ji, Parton Physics from Large-Momentum Effective Field Theory, *Sci. China Phys. Mech. Astron.* 57 (2014) 1407–1412. [arXiv:1404.6680](#), [doi:10.1007/s11433-014-5492-3](#).
- [580] Y.-Q. Ma, J.-W. Qiu, Extracting Parton Distribution Functions from Lattice QCD Calculations. [arXiv:1404.6860](#).

- [581] Y.-Q. Ma, J.-W. Qiu, QCD Factorization and PDFs from Lattice QCD Calculation, *Int. J. Mod. Phys. Conf. Ser.* 37 (2015) 1560041. [arXiv:1412.2688](#), [doi:10.1142/S2010194515600411](#).
- [582] A. V. Radyushkin, Quasi-parton distribution functions, momentum distributions, and pseudo-parton distribution functions, *Phys. Rev. D* 96 (3) (2017) 034025. [arXiv:1705.01488](#), [doi:10.1103/PhysRevD.96.034025](#).
- [583] K. Orginos, A. Radyushkin, J. Karpie, S. Zafeiropoulos, Lattice QCD exploration of pseudo-PDFs [arXiv:1706.05373](#).
- [584] H. Baer, T. Barklow, K. Fujii, Y. Gao, A. Hoang, S. Kanemura, J. List, H. E. Logan, A. Nomerotski, M. Perelstein, et al., The International Linear Collider Technical Design Report - Volume 2: Physics. [arXiv:1306.6352](#).
- [585] L. Linssen, A. Miyamoto, M. Stanitzki, H. Weerts, Physics and Detectors at CLIC: CLIC Conceptual Design Report. [arXiv:1202.5940](#), [doi:10.5170/CERN-2012-003](#).
- [586] M. Bicer, et al., First Look at the Physics Case of TLEP, *JHEP* 01 (2014) 164. [arXiv:1308.6176](#), [doi:10.1007/JHEP01\(2014\)164](#).
- [587] C.-S. S. Group, CEPC-SPPC Preliminary Conceptual Design Report. 1. Physics and Detector IHEP-CEPC-DR-2015-01, IHEP-TH-2015-01, IHEP-EP-2015-01.
- [588] R. Contino, et al., Physics at a 100 TeV pp collider: Higgs and EW symmetry breaking studies, CERN Yellow Report (3) (2017) 255–440. [arXiv:1606.09408](#), [doi:10.23731/CYRM-2017-003.255](#).
- [589] D. Boer, M. Diehl, R. Milner, R. Venugopalan, W. Vogelsang, et al., Gluons and the quark sea at high energies: Distributions, polarization, tomography. [arXiv:1108.1713](#).
- [590] J. L. Abelleira Fernandez, et al., On the Relation of the LHeC and the LHC. [arXiv:1211.5102](#).
- [591] N. Armesto, Physics at low x_{Bj} and in eA at the LHeC, *J. Phys. Conf. Ser.* 422 (2013) 012030. [doi:10.1088/1742-6596/422/1/012030](#).
- [592] M. Ciafaloni, D. Colferai, G. Salam, A. Stasto, A Matrix formulation for small- x singlet evolution, *JHEP* 0708 (2007) 046. [arXiv:0707.1453](#), [doi:10.1088/1126-6708/2007/08/046](#).
- [593] G. Altarelli, R. D. Ball, S. Forte, Small x Resummation with Quarks: Deep-Inelastic Scattering, *Nucl. Phys. B* 799 (2008) 199–240. [arXiv:0802.0032](#), [doi:10.1016/j.nuclphysb.2008.03.003](#).
- [594] M. Bonvini, S. Marzani, T. Peraro, Small- x resummation from HELL, *Eur. Phys. J. C* 76 (11) (2016) 597. [arXiv:1607.02153](#), [doi:10.1140/epjc/s10052-016-4445-6](#).
- [595] L. N. Lipatov, Reggeization of the Vector Meson and the Vacuum Singularity in Nonabelian Gauge Theories, *Sov. J. Nucl. Phys.* 23 (1976) 338–345, [*Yad. Fiz.* 23,642(1976)].
- [596] V. S. Fadin, E. A. Kuraev, L. N. Lipatov, On the Pomernanchuk Singularity in Asymptotically Free Theories, *Phys. Lett.* 60B (1975) 50–52. [doi:10.1016/0370-2693\(75\)90524-9](#).
- [597] E. A. Kuraev, L. N. Lipatov, V. S. Fadin, Multi - Reggeon Processes in the Yang-Mills Theory, *Sov. Phys. JETP* 44 (1976) 443–450, [*Zh. Eksp. Teor. Fiz.* 71,840(1976)].
- [598] F. Caola, S. Forte, J. Rojo, HERA data and DGLAP evolution: Theory and phenomenology, *Nucl. Phys. A* 854 (2011) 32–44. [arXiv:1007.5405](#), [doi:10.1016/j.nuclphysa.2010.08.009](#).
- [599] F. Caola, S. Forte, J. Rojo, Deviations from NLO QCD evolution in inclusive HERA data, *Phys. Lett. B* 686 (2010) 127–135. [arXiv:0910.3143](#), [doi:10.1016/j.physletb.2010.02.043](#).
- [600] J. L. Albacete, J. G. Milhano, P. Quiroga-Arias, J. Rojo, Linear vs Non-Linear QCD Evolution: From HERA Data to LHC Phenomenology, *Eur. Phys. J. C* 72 (2012) 2131. [arXiv:1203.1043](#), [doi:10.1140/epjc/s10052-012-2131-x](#).
- [601] J. L. Albacete, N. Armesto, J. G. Milhano, P. Quiroga-Arias, C. A. Salgado, AAMQS: A non-linear QCD analysis of new HERA data at small- x including heavy quarks, *Eur. Phys. J. C* 71 (2011) 1705. [arXiv:1012.4408](#), [doi:10.1140/epjc/s10052-011-1705-3](#).
- [602] J. Rojo, Parton Distributions at a 100 TeV Hadron Collider, *PoS DIS2016* (2016) 275. [arXiv:1605.08302](#).
- [603] R. Gauld, J. Rojo, E. Slade, The small- x gluon from forward charm production: implications for a 100 TeV proton collider, in: 25th International Workshop on Deep Inelastic Scattering and Related Topics (DIS 2017) Birmingham, UK, April 3-7, 2017, 2017. [arXiv:1705.04217](#).
- [604] T. Han, J. Sayre, S. Westhoff, Top-Quark Initiated Processes at High-Energy Hadron Colliders, *JHEP* 04 (2015) 145. [arXiv:1411.2588](#), [doi:10.1007/JHEP04\(2015\)145](#).
- [605] S. Dawson, A. Ismail, I. Low, A Redux on “When is the Top Quark a Parton?”, *Phys. Rev. D* 90 (2014) 014005. [arXiv:1405.6211](#), [doi:10.1103/PhysRevD.90.014005](#).
- [606] C. W. Bauer, N. Ferland, B. R. Webber, Standard Model Parton Distributions at Very High Energies, *JHEP* 08 (2017) 036. [arXiv:1703.08562](#), [doi:10.1007/JHEP08\(2017\)036](#).
- [607] M. Ciafaloni, P. Ciafaloni, D. Comelli, Towards collinear evolution equations in electroweak theory, *Phys. Rev. Lett.* 88 (2002) 102001. [arXiv:hep-ph/0111109](#), [doi:10.1103/PhysRevLett.88.102001](#).
- [608] J. Chen, T. Han, B. Tweedie, Electroweak Splitting Functions and High Energy Showering. [arXiv:1611.00788](#).
- [609] B. Fuks, J. Proudom, J. Rojo, I. Schienbein, Characterizing New Physics with Polarized Beams at High-Energy Hadron Colliders, *JHEP* 05 (2014) 045. [arXiv:1403.2383](#), [doi:10.1007/JHEP05\(2014\)045](#).