

Chapter II

Parton distribution functions

1 Consistency of LHC top pair production data and their impact on parton distributions ¹

We revisit the impact of the ATLAS and CMS top pair production measurements at $\sqrt{s} = 8$ TeV on a global determination of parton distribution functions (PDFs). Our analysis includes all the differential distributions from the ATLAS 8 TeV $t\bar{t}$ lepton+jet data set, together with their cross-correlations, in a PDF determination akin to the published NNPDF3.1 set. We study the mutual consistency of these distributions and their consistency with the rest of the data sets of the global fit. We specifically address the relative impact of the normalized and unnormalized data, the consequences of fitting the charm PDF, the role of the top quark transverse momentum distributions, and the effects of partially decorrelating experimental systematic uncertainties.

1.1 Top pair production data and parton distributions

The set of processes used for the accurate determination of the parton distribution functions (PDFs) [797] has been steadily widening over time, beyond the traditional combination of deep-inelastic scattering (DIS), Drell-Yan (DY) and jet production data which has now been used for more than thirty years [798]. Top pair production data were suggested as an effective way to constrain the gluon distribution at large x since the early days of the LHC (see e.g. Ref. [799]); and their impact on PDF fits was studied both at the level of total cross-sections [800] and, subsequently, of differential distributions [801]. Several measurements have been published by both ATLAS and CMS at center-of-mass energies of $\sqrt{s} = 5.02, 7, 8$ and 13 TeV: total cross-sections and differential distributions with respect to a variety of kinematic variables, including the top transverse momentum p_T^t , the top rapidity y_t , the rapidity of the top pair $y_{t\bar{t}}$ and the invariant mass of the top pair $m_{t\bar{t}}$, both normalized and unnormalized with respect to the total cross-section (see Sect. 67.3.1 in Ref. [794] for an updated review of all of the available measurements).

Differential cross-sections for top-quark pair production measured by ATLAS [680] and CMS [679] in the lepton+jets channel at a center-of-mass energy of 8 TeV were used in the NNPDF3.1 global PDF determination [769], complementing selected total cross-sections (already used in the ABM12 [806] and NNPDF3.0 [748] PDF fits). In order to include these observables in the PDF fit, the choice of a specific kinematic distribution had to be made because the information on correlations across distributions was not available at that time: their simultaneous inclusion would have otherwise amounted to double counting, as they come from the same underlying data. The particular choice of observables adopted for NNPDF3.1 (see Table II.1), namely the normalized rapidity distribution of the top quark (for ATLAS) and of the top pair (for CMS), was based on the results of a previous study [801], which analyzed the impact of different observables on PDFs and their consistency.

In the NNPDF3.1 global analysis the impact of the inclusion of the top rapidity distributions at 8 TeV on the resulting PDFs was assessed, and found to be significant on the gluon PDF (in the region $0.1 \lesssim x \lesssim 0.5$) and negligible for other PDFs. The consistency of the constraint imposed on the large- x gluon by this data with those coming from other data included in NNPDF3.1 was further studied in Ref. [807], where top, Z transverse momentum and single-

¹ S. Forte, E. R. Nocera, J. Rojo

Dataset	N_{dat}	Fit 1	Fit 2	Fit 3	Fit 4	Fit 5	Fit 6	Fit 7	Fit 8	Fit 9	
ATLAS $t\bar{t}$ norm. diff. (cor.) [680]	21	[2.74]	[2.90]	[2.64]	2.28	[4.60]	2.29	[3.49]	2.23	[2.31]	
ATLAS $t\bar{t}$ norm. diff. (unc.) [680]	21	[2.08]	[2.07]	[2.05]	[1.94]	[4.16]	[1.89]	[3.07]	[1.92]	[1.97]	
ATLAS $1/\sigma d\sigma/dp_T^t$	7	[3.50]	[3.57]	[3.25]	2.94	[2.45]	2.95	[2.54]	2.92	[3.10]	
ATLAS $1/\sigma d\sigma/dy_t$	4	1.45	1.33	1.08	1.20	[4.81]	1.10	[2.98]	1.10	[1.05]	
ATLAS $1/\sigma d\sigma/dy_{t\bar{t}}$	4	[1.26]	[1.13]	[1.66]	1.55	[10.2]	1.40	[6.30]	1.31	[1.59]	
ATLAS $1/\sigma d\sigma/dm_{t\bar{t}}$	6	[1.78]	[1.83]	[1.67]	1.59	[1.36]	1.61	[1.42]	1.61	[1.61]	
ATLAS $t\bar{t}$ unnorm. diff. (cor.) [680]	25	[7.96]	[8.33]	[7.52]	[6.88]	5.76	[7.18]	5.23	[2.25]	2.16	
ATLAS $t\bar{t}$ unnorm. diff. (unc.) [680]	25	[2.09]	[2.08]	[2.09]	[2.06]	[2.28]	[2.25]	[2.17]	[2.07]	[2.02]	
ATLAS $d\sigma/dp_T^t$	† ‡	8	[2.41]	[2.46]	[2.50]	[2.43]	2.50	[2.46]	2.54	[2.50]	2.42
ATLAS $d\sigma/dy_t$	†	5	[0.87]	[0.78]	[0.73]	[0.76]	1.14	[0.73]	0.87	[0.73]	0.66
ATLAS $d\sigma/dy_{t\bar{t}}$	†	5	[1.21]	[1.11]	[1.32]	[1.19]	2.36	[1.15]	1.86	[1.14]	1.16
ATLAS $d\sigma/dm_{t\bar{t}}$	† ‡	7	[3.27]	[3.30]	[3.18]	[3.24]	2.85	[3.25]	2.94	[3.25]	3.16
Fixed target DIS (NC)	1039	1.21	1.20	1.20	1.20	1.20	1.27	1.27	1.26	1.26	
HERA DIS (NC)	1064	1.15	1.14	1.15	1.14	1.15	1.20	1.20	1.20	1.20	
Fixed target DIS (CC)	908	1.08	1.09	1.08	1.09	1.08	1.11	1.11	1.11	1.10	
HERA DIS (CC)	81	1.19	1.18	1.19	1.18	1.19	1.15	1.14	1.14	1.15	
Fermilab DY	189	1.24	1.22	1.23	1.24	1.23	1.14	1.14	1.15	1.14	
Tevatron DY	74	1.29	1.26	1.28	1.27	1.30	1.25	1.25	1.25	1.24	
ATLAS DY	75	1.55	1.50	1.53	1.48	1.49	1.80	1.81	1.81	1.77	
ATLAS W/Z rap. 2011	34	2.14	2.19	2.15	2.07	2.10	2.69	2.72	2.71	2.63	
CMS DY	154	1.23	1.23	1.22	1.23	1.22	1.24	1.23	1.24	1.25	
LHCb DY	85	1.47	1.52	1.57	1.49	1.55	1.43	1.47	1.41	1.42	
ATLAS jets	31	0.90	1.13	1.09	1.08	1.01	1.11	1.05	1.11	1.12	
CMS jets	133	0.88	0.96	0.93	0.95	1.05	0.94	0.99	0.93	0.93	
ATLAS $Z p_T$	92	0.90	0.92	0.93	0.91	0.95	0.95	0.97	0.94	0.96	
CMS $Z p_T$	28	1.33	1.30	1.31	1.35	1.33	1.30	1.31	1.31	1.28	
ATLAS $\sigma_{t\bar{t}}$ [674, 675]	3/2	0.86	0.84	0.77	0.74	0.68	0.72	0.83	0.72	0.71	
CMS $\sigma_{t\bar{t}}$ [676, 802]	3	0.20	0.21	0.23	0.29	0.56	0.35	0.34	0.28	0.27	
CMS $1/\sigma d\sigma/dy_{t\bar{t}}$ [679]	9	0.94	1.08	1.05	0.98	1.01	0.97	1.46	0.97	1.01	
Total		1.148	1.163	1.163	1.171	1.195	1.204	1.227	1.204	1.203	

Table II.1: Value of χ^2 per data point for all of the ATLAS top pair-production distributions, with respect to the top transverse momentum p_T^t , the top rapidity y_t , the rapidity of the top pair $y_{t\bar{t}}$ and the invariant mass of the top pair $m_{t\bar{t}}$, both normalized and unnormalized, and for all of the above combined; χ^2 values are also shown for all other data from the NNPDF3.1 data set [803], which are included in all the PDF fits considered here: charged- and neutral-current (CC and NC) DIS structure functions from both fixed-target and HERA combined experiments; fixed-target and collider DY rapidity and invariant mass distributions; single-jet inclusive cross-sections; Z transverse momentum distributions; total top-pair cross-sections; and the CMS top rapidity distribution. For each data set the total number of data points is shown; note that indented data sets are subsets of the preceding non-indented data set. For the ATLAS top-pair differential distributions we indicate whether they were included in NNPDF3.1 [769] (–), in CT18 [804] (‡), or in the MMHT-based study of Ref. [805] (†). Each column corresponds to a different PDF fit (see text); χ^2 values for the data sets not included in the fit are quoted in brackets.

inclusive jet distributions were added in turn to a baseline global data set. Excellent consistency was found, with all data sets pulling the gluon in the same direction, and the top and jet data having the biggest impact.

There are a number of reasons why the impact of top data on PDF determination is worth revisiting.

- Since the publication of the original ATLAS paper [680], the covariance matrix of the individual measurements was updated [808, 809]: it is advisable to check whether the previous results of Ref. [769] are affected by this update.
- In the same Refs. [808, 809] full information on the correlation between pairs of different kinematic distributions was made available: it is now possible to include all distributions at once and check the comparative impact on PDFs, and how it affects the conclusions of Refs. [769, 807].
- Some recent studies, specifically an ATLAS study [808, 809] within the xFITTER [810] framework, and a study [805] based on the MMHT [750] framework, found that there are serious difficulties in simultaneously including all of the differential distributions from the ATLAS 8 TeV lepton+jets top data set in a PDF determination: it is worth investigating whether similar conclusions also apply when analyzing these data within the current

NNPDF framework [769].

- The recently published CT18 PDF set [804] also includes top-quark pair differential distributions, but with a different choice of observables in comparison to NNPDF3.1 (see Table II.1). It is interesting to compare and assess the impact of different choices at the PDF level.

We will address all these issues by performing a number of PDF determinations based on the NNPDF3.1 methodology and data set, adding the full set of ATLAS top data either in normalized or unnormalized form to the baseline data set in various ways, and studying the fit quality and the impact on PDFs.

1.2 The ATLAS top production data and their impact on PDFs

All the PDF fits presented here are based on the NNPDF3.1 methodology of Ref. [769], with the slightly modified data set used in Ref. [803]. The latter differs from the original NNPDF3.1 data set in that only processes for which full NNLO computations are available are included (in particular, in the NNPDF3.1 fit some jet data were included using NLO theory). This data set will be supplemented with a number of top pair differential distributions measured in the lepton+jet channel at a center-of-mass energy of 8 TeV by ATLAS (see Table II.1). We refer to Ref. [769] for a detailed discussion of the NNPDF3.1 data set and the associated fitting methodology. In Sect. 1.3 we will present the various PDF sets and discuss how their features vary as the underlying data set is changed, in a series of pairwise comparisons between PDFs. In Sect. 1.3.1 we will then discuss these results and our best understanding of them, also using information from PDF determinations in which methodological changes are made, either in the treatment of theory or of correlated experimental uncertainties, as well as a few further auxiliary PDF determinations based on special subsets of data.

1.3 The impact of the data set choice on PDFs

In Table II.1 we list all the ATLAS top-quark pair observables corresponding to the 8 TeV lepton+jets data set whose inclusion we consider and compare. In each case we provide the total number of data points, and we indicate whether the observable was included in NNPDF3.1 [769] (–), CT18 [804] (‡) or the recent MMHT-based study [805] (†). In the case of the CMS measurements, the same top-quark data as in the NNPDF3.1 PDF determination are included: on the one hand, full correlations are not available for CMS, hence the inclusion of all differential distributions at once is not possible; also, there appear to be no specific issues with the CMS 8 TeV top-quark data, with consensus [801,805,811] that all observables show a similar pull on the gluon distribution, with only the invariant mass distribution providing a poor fit quality. Hence, as in the related study of Ref. [805], here we will focus our attention on the ATLAS data.

Our results in terms of fit quality are summarized in Table II.1, where we show the χ^2 per datapoint for the ATLAS top normalized and unnormalized (or, equivalently, absolute) distributions, as well as (indented) the breakdown of each of them into the four individual observables that make up the normalized or absolute top data sets: transverse momentum distribution, rapidity distribution, pair rapidity distribution, and pair invariant mass distribution. Note that the normalized distributions always have one fewer datapoint, because the last data bin is fixed by the normalization condition. Also, to avoid double counting, the corresponding measurement of the total cross-section is removed from the fit whenever unnormalized distributions are included (while other total cross-sections at 7 and 13 TeV are retained). Note that in principle, when fitting the normalized distribution, the correlation between the differential distribution and total cross-section should be included. This is not done because this information is not

available to us. In principle, for full consistency we should therefore exclude the total cross-section from the data set: note however that because this is a single data point, with $\chi^2 \lesssim 1$, so this exclusion would in practice make no difference. We have also computed χ^2 values for the full ATLAS top normalized or unnormalized data set decorrelating different distributions, i.e., using a block-diagonal covariance matrix that only correlates data points that belong to the same distribution: these values are provided for illustration as a separate row in the Table (and not used for fitting).

In Table II.1 we also provide χ^2 values for all other data sets in the global fit: DIS structure functions from both fixed-target and HERA combined experiments; fixed-target and collider DY rapidity and invariant mass distributions; single-jet inclusive cross-sections; Z transverse momentum distributions; total top-pair cross-sections; the CMS top rapidity distribution; and finally the global fit quality for the complete fitted data set. In view of the discussion in Sect. 1.3.1 below for the ATLAS DY data we also show (indented) the χ^2 value corresponding to the specific subset of this data corresponding to the 2011 W/Z rapidity distribution [517]. Fit quality is always shown both for data which are and for those which are not included in the fit. All χ^2 values not used for fitting are shown in square brackets in Table II.1.

Each column in Table II.1 corresponds to a separate PDF fit. All these fits have been carried out using NNLO theory, and they only differ in the treatment of the ATLAS top data, with all the rest of the data set being identical to Ref. [803]. We also have performed the corresponding NLO fits, but these are not shown here because they do not appear to add anything to the discussion: they merely exhibit somewhat worse fit quality but with identical qualitative features.

The fits included in Table II.1 and which we will discuss below are the following:

- Fit 1** NNP3.1: this simply reproduces for reference the published [769] NNP3.1 results.
- Fit 2** Baseline: this is our baseline fit, which only differs from the published NNP3.1 because it is based on the slightly different data set already adopted in Ref. [803].²
- Fit 3** Baseline, corrected: this is the same as the baseline, but now using the new, updated covariance matrix for the top data from Ref. [808, 809].
- Fit 4** All, normalized: this includes all of the different ATLAS top observables, and correlations across different distributions [808, 809, 812, 813].³
- Fit 5** All, unnormalized: this is the same as #4, but now including all of the observables in the absolute, rather than normalized version.
- Fit 6** Perturbative charm, normalized: this is the same as #4, but now with the charm PDF being generated through perturbative matching (as e.g. CT18 and MMHT do), rather than independently parametrized and fitted [814] as in the default NNP3.1 set.
- Fit 7** Perturbative charm, unnormalized: this is the same as #6, but now using absolute observables.

²Note that the number of data points for the ATLAS top rapidity distribution is $N_{\text{dat}} = 10$ in Table 3 of Ref. [769], while it is $N_{\text{dat}} = 5$ in Table II.1. This is due to the fact that both the distribution with respect to rapidity ($N_{\text{dat}} = 10$) and with respect to the absolute value of the rapidity ($N_{\text{dat}} = 5$) were published by ATLAS in Ref. [680]. The former was used in Ref. [769], but information on correlations was only made available in Ref. [808, 809] for the latter, which is therefore used here.

³We have checked with the authors [813] of Ref. [805] that our implementation of correlations is in agreement with their own.

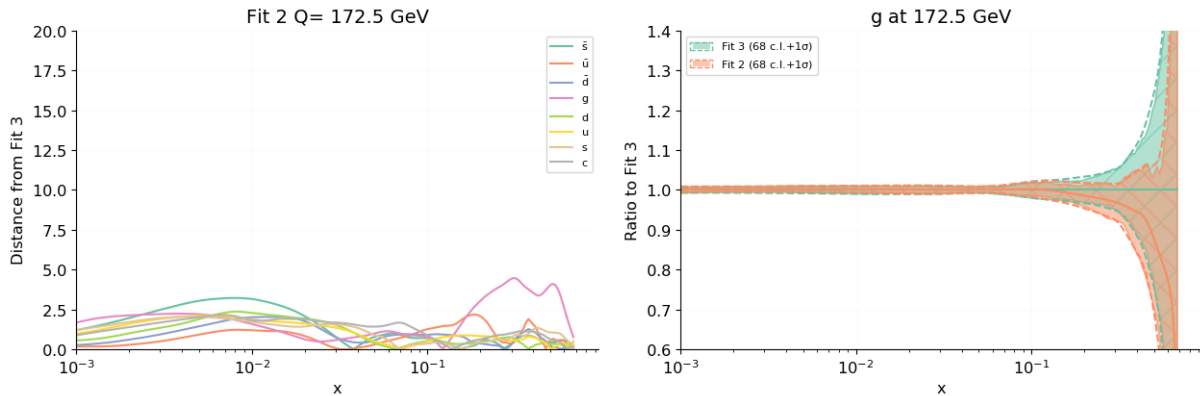


Fig. II.1: Comparison between the baseline (set #2) and PDFs determined using the same data set but the updated covariance matrix (set #3). The distances between PDFs (left) and the gluon distributions (right) are shown.

Fit 8 Perturbative charm, normalized and decorrelated: this is the same as #6, but now decorrelating parton-shower systematic uncertainties across bins belonging to different distributions, as suggested in Ref. [805].

Fit 9 Perturbative charm, unnormalized and decorrelated: this is the same as #8, but now using absolute observables.

We now discuss and compare the PDF determinations #1-#5, which correspond to different choices of underlying data set, in order to address the various issues listed at the end of Sect. 1.1; these comparisons have been generated using the REPORTENGINE software [815]. For each comparison, we show (as a function of x and at the scale corresponding to the top quark mass, $Q = 172.5$ GeV) the distances between all PDFs and we compare the gluon PDF, which is mostly affected by the top data. Recall that the distance d is defined as the difference in units of the standard deviation of the mean, so for a sample of 100 replicas $d \sim 1$ corresponds to statistically identical PDFs (replicas extracted from the same underlying distribution) and $d \sim 10$ corresponds to PDFs that differ by one- σ . In Sect. 1.3.1 we will discuss the PDF determinations #6-#9, which correspond to changes in methodology which we have performed in order to correctly interpret these results.

First, we assess the impact of the update in the ATLAS 8 TeV lepton+jets covariance matrix presented in Ref. [808, 809] on the results of Ref. [769]. We start from the baseline PDF set #2, which is essentially the same as NNPDF3.1, as can be seen from the χ^2 values in Table II.1. In Fig. II.1 we compare the baseline to PDFs determined using the same data set (i.e. essentially the NNPDF3.1 data set) but with the new updated covariance matrix. It is clear from the figure that the two sets of PDFs are very close to being statistically indistinguishable: the updated covariance matrix has essentially no effect on the PDF determination. Interestingly, it does however lead to an improved value of the χ^2 for the top rapidity distribution, which now corresponds to a near-perfect fit, $\chi^2 = 1.08$.

Next, we enlarge the top data set to include all of the ATLAS distributions, in normalized form; PDFs before and after this enlargement of the baseline data set are compared in Fig. II.2. It is clear that also in this case no significant effect is seen: the simultaneous inclusion of the four differential distributions carries effectively no new information as compared to fitting only the y_t distribution. The fit quality for individual observables is poor for the transverse momentum distribution ($\chi^2 = 2.94$) but fair to good for all other distributions, with the rapidity distribution




Fig. II.2: As Fig. II.1 but now comparing PDFs determined from the baseline data set (but improved covariance matrix, set #3) and PDFs determined including the full set of ATLAS normalized distributions (set #4).




Fig. II.3: As Fig. II.1 but now comparing PDFs determined including the full set of ATLAS normalized distributions respectively in normalized (set #4) or unnormalized (set #5) form.

being fitted best; the fit quality to the whole set of top data is fair, and it does not significantly improve by decorrelating systematic uncertainties. The fit quality to the global data set is essentially the same as that of the baseline, as it must be, given that the PDFs are unchanged. The reasons for the poor fit of the transverse momentum distributions will be discussed in Sect. 1.3.1. Note that the fit quality to each of the unnormalized distributions is similar (and sometimes better) to that of the corresponding normalized ones, despite the fact that these distributions are not being fitted; only the fit quality to the invariant mass distribution exhibits a significant deterioration. However, the fit quality to the full set of unnormalized distributions is very poor ($\chi^2 > 6$). Nonetheless, if we recompute this χ^2 value decorrelating uncertainties as discussed above (by taking a block-diagonal covariance matrix and thus neglecting cross-correlations between distributions), it becomes fair ($\chi^2 = 2.06$). This suggests an issue with correlated uncertainties for unnormalized observables. We will revisit this point when discussing PDF sets #8 and #9.

We now repeat the PDF determination with all ATLAS top distributions included, but using the unnormalized distributions. The resulting PDFs are compared to those obtained using the normalized distributions in Fig. II.3. It is clear that now a shift by more than one σ is observed between the two gluon PDFs in the large- x region $0.1 \lesssim x \lesssim 1$, with some smaller shift

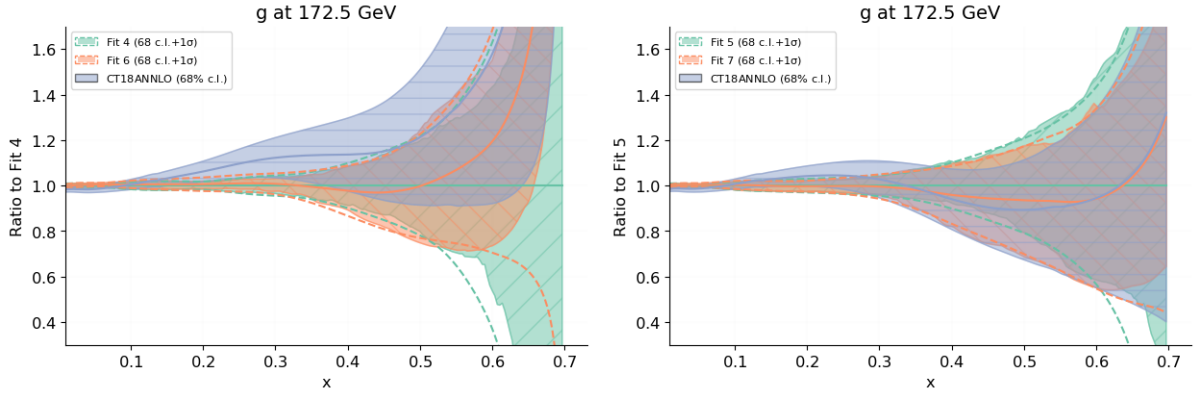


Fig. II.4: Comparison between the gluon PDFs in the sets in which all distributions are fitted, with fitted charm or perturbative charm. The comparison is shown both for fits to normalized distributions (left: PDF set #4 vs. set #6) and to unnormalized distributions (right: PDF set #5 vs. set #7). The gluon from the CT18A PDF set is also shown for comparison.

also seen for some quark PDFs: the absolute top-pair distributions appears to pull the large- x gluon upwards in comparison to the normalized distributions. The fit quality to the individual absolute distributions however turns out to be similar (and sometimes even worse) in comparison to the case in which they were not fitted, and the pattern is unchanged: it is only the fit quality to the correlated set of top observables that improves somewhat (from $\chi^2 = 6.88$ to $\chi^2 = 5.76$), though it remains very poor. Just like in the case in which normalized distributions were fitted, we find that this value improves considerably if it is recomputed decorrelating experimental systematics: from $\chi^2 = 5.76$ to $\chi^2 = 2.28$, a value similar (in fact slightly worse) to the value found when the normalized distributions were fitted. It is important to observe that in fit #5, in which top observables are included in unnormalized form, the fit quality to the global data set deteriorates somewhat in comparison to fit #4, in which normalized observables were used.

This concludes our presentation of results from PDF fits corresponding to the variations of underlying data set that we consider here. We now turn to their interpretation.

1.3.1 Interpretation and dependence of PDFs on the methodology

The PDF determinations presented in the previous section lead to the following immediate conclusions:

- The parton distributions determined using the NNPDF3.1 data set and methodology are unaffected if the ATLAS normalized y_t distribution is supplemented by the full set of ATLAS normalized differential distributions; the fit quality is generally good except for the transverse momentum distribution which is poorly fitted.
- If the normalized distributions are replaced by the absolute ones, the large- x gluon is pushed upwards in the large x region. The fit quality to the individual ATLAS top distributions is similar to the one found in the normalized case, but the fit quality to the full set of correlated observables is very poor, and the fit quality to the rest of the global data set deteriorates somewhat.

In order to compare with the results obtained by other groups, it is important to recall that a notable aspect of the NNPDF3.1 methodology is that the charm PDF is fitted, instead of being obtained from perturbative matching conditions. In Ref. [769] this choice was found to

be crucial in order to achieve a reasonable fit to the high-precision ATLAS 2011 W/Z rapidity distributions [517]. In view of the fact that the ATLAS study of Refs. [808,809] appears to find tension between their 2011 W, Z data set and their 8 TeV top observables, it is interesting to investigate whether these conclusions are affected if the NNPDF3.1 methodology is modified by deriving the charm PDF from the corresponding perturbative matching conditions, rather than being fitted. This choice corresponds to fits #6 and #7.

It is apparent from the χ^2 values of Table II.1 that, if the charm PDF is no longer fitted, the quality to the fit to the ATLAS 2011 W/Z rapidity data significantly deteriorates, consistently with the results of Ref. [769]. However, the fit quality to the top observables remains essentially the same as that found in the corresponding fitted-charm PDF sets. Interestingly, however, the quality of the global χ^2 in all these fits somewhat deteriorates, and it is similar (though somewhat worse) to that found using PDF set #5, namely when using top absolute rather than normalized distributions. The origin of this state of affairs can be understood by comparing the gluon distribution which is found in each of these cases. The comparison is displayed in Fig. II.4, where the CT18 gluon PDF is also shown for reference. The gluon from the CT18A set is shown, because it is based on a data set which also includes the ATLAS 2011 W/Z data, which are excluded in the baseline CT18 determination.

Recall from Fig. II.3 that we found that, when fitting the absolute top distributions, the gluon PDF was pushed upwards somewhat, and that this led to some deterioration of the global fit quality, suggesting that this enhanced gluon is disfavored by the global fit. It appears from Fig. II.4 that when fitting the normalized distribution, and replacing fitted charm with perturbative charm, the gluon is similarly pushed upwards. If the unnormalized distribution is fitted instead, it makes essentially no difference whether charm is fitted or not: a similar fit quality is found with either choice. Also, the gluon determined from a fit to unnormalized distributions is found to be in very good agreement with the CT18 gluon.

From this comparison it therefore appears that the good simultaneous fit of the ATLAS top data set and the global fit achieved in fit #4 (which is in turn essentially identical to NNPDF3.1) relies on two ingredients: using the normalized distributions, and fitting charm. If the unnormalized distributions are used, an enhanced gluon is found, with a worse global fit quality, and very little dependence on whether charm is fitted or not. This gluon is in excellent agreement with the CT18 gluon. If charm PDF is not fitted, but rather generated from perturbative matching, a poor fit to the ATLAS 2011 W/Z rapidity distribution is obtained.

In the MMHT-based study of Ref. [805], where only top-pair absolute distributions were considered, it was suggested that the poor fit quality to these distributions could be improved by decorrelating the parton shower (PS) uncertainties, and it was found that such a decorrelation affects the gluon. We have therefore checked whether our results would also be affected by decorrelating uncertainties as suggested in Ref. [805], by producing PDF sets #8 and #9. These PDF sets differ only in the treatment of correlated uncertainties from PDF sets #6 and #7 respectively. Note that this decorrelation is milder than that used in the computation of the decorrelated rows of Table II.1 (not used for fitting), in which the covariance matrix for the ATLAS top distributions was taken to be block-diagonal. The PDF sets #8 and #9, with perturbative charm and decorrelated uncertainties, are directly comparable to Ref. [805].

As should be clear from Table II.1, we find that indeed decorrelating uncertainties as suggested in Ref. [805] does lead to an acceptable fit quality for the full set of unnormalized top data, both when they are fitted, or when the absolute data are fitted instead. Interestingly, the value of the χ^2 found for the full set of unnormalized top data in these fits is almost identical to the value found in the corresponding fits in which the correlations were kept when fitting, but the χ^2 was fully decorrelated: for the unnormalized data in fit #9 we find $\chi^2 = 2.16$, while in fit #7 the uncorrelated χ^2 value is $\chi^2 = 2.17$. This suggests that the bulk of the correlation

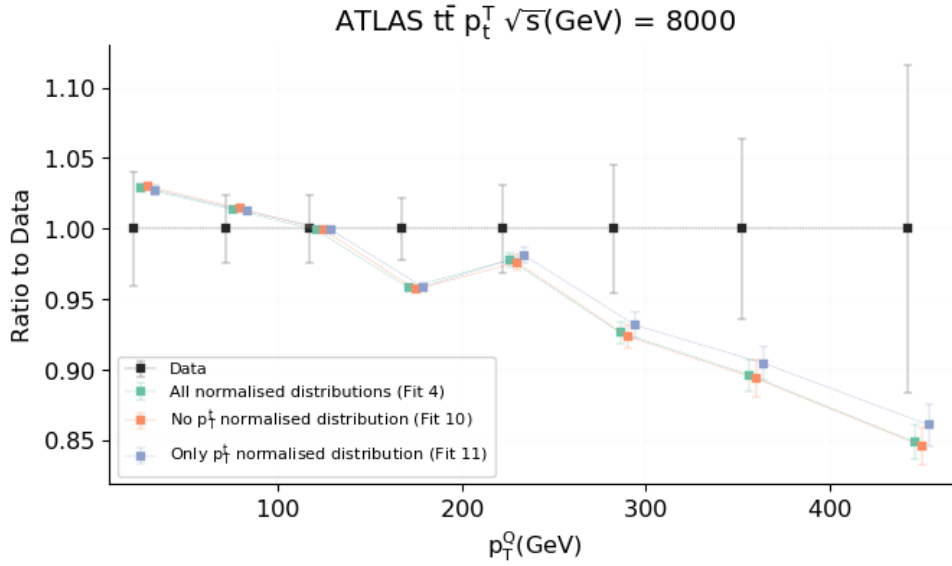


Fig. II.5: Comparison between data and theory predictions for the ATLAS top transverse momentum distribution. The predictions shown correspond to PDF sets #4, #10 and #11 (see text).

is indeed coming from the PS uncertainties singled out in Ref. [805]: removing them leads to same answer as removing correlations between different observables altogether. Also, it suggests that whether one fits the correlated or uncorrelated quantities makes very little difference at the level of PDFs, since the uncorrelated χ^2 value remains the same. This conclusion is supported by the fact that indeed all the χ^2 values for other data sets are essentially unchanged by having performed the decorrelation of Ref. [805]: χ^2 values for the other data sets are the same in fit #7 and fit #9, and also the same in fit #6 and fit #8. We have also checked explicitly that this is the case at the PDF level: when decorrelating uncertainties PDFs change very little. We conclude that fit results obtained in our framework are stable upon decorrelation.

We finally turn to our preferred PDF set #4, which, as mentioned, achieves good fit quality to both the ATLAS top data set and the global fit. We have seen that this PDF set is extremely stable, in that fitting just the top rapidity distribution, or the whole data set, leads to essentially the same PDFs. One may however note that even though the fit quality to the full data set is fair, the fit quality to the transverse momentum distribution remains poor. One may then ask first, whether this is again due to issues with the correlation matrix, and furthermore, if these data might favor a different PDF shape. In order to answer this question, we have performed two more PDF fits:

Fit 10 Same as Fit 4, but excluding the transverse momentum distribution from the ATLAS top data set.

Fit 11 Same as Fit 4, but now only including the transverse momentum distribution in the ATLAS top data set.

The ATLAS data for the transverse momentum distribution are compared to predictions obtained using our preferred Fit #4 as well as these two PDF sets in Fig. II.5. It is clear that the poor fit quality to these data is due to the fact that they have a genuinely different shape in comparison to the theory prediction, and thus it cannot be due to a treatment of correlations. However, when excluding these data, or only including them, nothing changes: the global fit remains perfectly stable upon their inclusion or exclusion, as we have also verified at

the PDF level. Indeed, interestingly, even if only the top transverse momentum distribution is fitted, the best-fit PDFs are indistinguishable from those obtained fitting all of the (normalized) distributions. In sum, while the reason of this data-theory discrepancy is unclear, it seems to be immaterial for the purposes of PDF determinations.

1.4 Conclusions

We have studied the effect of including the full set of differential top-quark pair distributions from the ATLAS 8 TeV lepton+jets data set in the NNPDF3.1 global PDF determination. Our main conclusions are the following:

- Inclusion of the normalized observables yields results which are essentially identical to those of the NNPDF3.1 PDF determination, in which only the top rapidity distribution was included.
- Good fit quality to all top observables except the transverse momentum distribution is found.
- The top transverse momentum distribution appears to have a somewhat different shape in comparison to the theory prediction; however, fit results are stable upon its inclusion or exclusion, and in fact even a fit in which only the transverse momentum distribution is included leads to PDFs which are the same as those when all distributions are fitted.
- Fitting the charm PDF, rather than obtaining it from perturbative matching, is crucial in order to achieve compatibility of the top production data and the ATLAS 2011 W/Z rapidity distribution data; if charm is not fitted the gluon PDF is affected and the global fit quality deteriorates.
- If unnormalized observables are used instead, the gluon PDF is somewhat enhanced in the large x region; this leads to a deterioration of the global fit quality. The fit quality to the full set of top observables is extremely poor, but it can be brought to be similar to what is found when fitting normalized observables by decorrelating different distributions; PDFs are stable upon this decorrelation.
- If unnormalized observables are fitted, it makes little difference to the gluon PDF if charm is fitted or not, though if charm is not fitted, all the light quark PDFs change by an amount which is small but sufficient to lead to considerable deterioration of the fit to the ATLAS 2011 W/Z rapidity distribution data.

We conclude that normalized top observables, together with fitted charm, are necessary ingredients in order to achieve good fit quality to the ATLAS 8 TeV lepton+jets top production data within the framework of the NNPDF3.1 global determination. Best-fit results obtained with these choices are extremely stable upon variations of the data set and treatment of uncertainties. A detailed benchmarking against results found in a CT, MMHT and ATLAS-xFITTER framework would be extremely beneficial for a complete understanding and validation of our findings. Also, it will be interesting to see to what extent these conclusions remain true when additional top-quark pair data sets are included in the fit, in particular the ATLAS and CMS $\sqrt{s} = 13$ TeV measurements, as well as other gluon-sensitive observables such as the jet and dijet cross-sections.

Acknowledgements

We are grateful to Shaun Bailey and Lucian Harland-Lang for correspondence and detailed clarifications concerning Refs. [805, 808] and for comments. We thank Joey Huston, Robert Thorne

and Mandy Cooper-Sarkar for discussions and comments on the manuscript.

Stefano Forte is supported by the European Research Council under the European Union’s Horizon 2020 research and innovation Programme (grant agreement ERC-AdG-740006). Emanuele R. Nocera is supported by the European Commission through the Marie Skłodowska-Curie Action ParDHonS FFs.TMDs (grant number 752748). Juan Rojo is partially supported by the Dutch National Science Foundation (NWO).

2 Assessing the compatibility of experimental pulls on LHC parton luminosities with the L_2 sensitivity ⁴

2.1 Introduction

The present lack of complete knowledge of the proton’s parton distribution functions (PDFs) forms one of the most significant uncertainties for crucial physics processes at the LHC, such as the gg fusion cross section for Higgs boson production. Information on the PDFs comes from global fits to a wide variety of high-energy data, including those taken by various LHC experiments. Global PDF analyses involve a subtle interplay among all of the fitted data sets in order to determine an optimal set of central PDFs and their associated uncertainties. Complicating the realization of these optimal PDFs are systematic tensions, which can exist among data sets, and which tend to resist the reduction in PDF uncertainties suggested by the precision of the fitted data.

Any program to comprehend and resolve these tensions necessarily requires a set of tools to determine the PDF sensitivities and pulls of the data in a given global analysis. The CTEQ-TEA (CT) group has pioneered a number of techniques to establish the sensitivity of a particular data set for constraining a particular PDF or observable, such as the Lagrange Multiplier (LM) scans [816], which, though robustly informative, are computationally costly [804], and evaluated for specific, fixed values of the parton momentum fraction, x and factorization scale, Q . On the other hand, the outcomes of popular fast techniques based on Monte-Carlo PDF reweighting [817–820], or Hessian profiling [821] and updating [822], sensitively depend on the choice of either statistical weights or tolerance.

A technique that does not involve the computational overhead of the LM method or the ambiguities of the reweighting approach is the L_2 sensitivity technique, as defined in Ref. [823] and deployed in the recently-released CT18 global fit [804]. The L_2 sensitivity is inexpensive to compute and provides an informative approximation to the $\Delta\chi^2$ trends in a given global analysis. Moreover, the L_2 sensitivity can also be readily calculated across a wide range of x , allowing the $\Delta\chi^2$ variations shown in the LM scans to be visualized and interpreted for multiple, simultaneous x values. We stress that the qualitative conclusions revealed by consideration of the L_2 sensitivities, discussed and presented below, are consistent with the picture based on the LM scans themselves. Although the L_2 sensitivities do not always provide the same numerical ordering as the LM scans for the subdominant experiments, they offer complementary information over broader reaches of x that are not completely captured by the LM scans.

While the L_2 sensitivity was used to analyze the pulls of data on the PDFs themselves in recent CT fits, the method is sufficiently flexible that it may be applied to other phenomenologically relevant quantities, including the parton-parton luminosities used in predictions for processes at hadron colliders. In this note, we demonstrate this application, highlighting a number of phenomenological consequences.

⁴ T. J. Hobbs, J. Huston, P. Nadolsky

2.2 Definition of the L_2 sensitivity

We work in the Hessian formalism [824–826] and compute the L_2 sensitivity, $S_{f,L_2}(E)$, for each experiment, E , as

$$S_{f,L_2}(E) = \vec{\nabla}\chi_E^2 \cdot \frac{\vec{\nabla}f}{|\vec{\nabla}f|} = \Delta\chi_E^2 \cos\varphi(f, \chi_E^2), \quad (\text{II.1})$$

which yields the variation of the log-likelihood function χ_E^2 due to a unit-length displacement of the fitted PDF parameters away from the global minimum \vec{a}_0 of $\chi^2(\vec{a})$ in the direction of $\vec{\nabla}f$. The PDF parameters \vec{a} are normalized so that a unit displacement from the best fit in any direction corresponds to the default confidence level of the Hessian error set (90% for CT18, on average corresponding to slightly less than $\Delta\chi_{tot}^2 = 100$ in a given direction.)

This displacement increases the PDF $f(x, Q)$ by its Hessian PDF error, Δf , and, to the extent its PDF variation is correlated with that of $f(x, Q)$ through the correlation angle

$$\varphi(f, \chi_E^2) = \cos^{-1} \left(\frac{\vec{\nabla}f}{|\vec{\nabla}f|} \cdot \frac{\vec{\nabla}\chi_E^2}{|\vec{\nabla}\chi_E^2|} \right), \quad (\text{II.2})$$

it changes χ_E^2 by $\Delta\chi_E^2(\hat{a}_f) = \Delta\chi_E^2 \cos\varphi(f, \chi_E^2) = S_{f,L_2}(E)$. The L_2 sensitivity, $S_{f,L_2}(E)$, therefore quantifies the impact that uncertainty-driven variations of PDFs at fixed x and Q have upon the description of fitted data sets. Plotting $S_{f,L_2}(E)$ against x furnishes useful information regarding the pulls of the CT18(Z) data sets upon the PDFs fitted in the global analysis, as well as various PDF combinations of interest. This also permits the rapid visualization of possible tensions within the global fit, since the PDF variations of some parton densities of given flavor are correlated with the variation of χ_E^2 (*i.e.*, $S_{f,L_2}(E) > 0$), while others are anti-correlated ($S_{f,L_2}(E) < 0$), at similar values of (x, Q) .

The terms on the right-hand side of Eq. (II.1) for S_{f,L_2} are computed as

$$\Delta X = |\vec{\nabla}X| = \frac{1}{2} \sqrt{\sum_{i=1}^{N_{eig}} (X_i^{(+)} - X_i^{(-)})^2}, \quad (\text{II.3})$$

and

$$\cos\varphi = \frac{\vec{\nabla}X \cdot \vec{\nabla}Y}{\Delta X \Delta Y} = \frac{1}{4\Delta X \Delta Y} \sum_{i=1}^{N_{eig}} (X_i^{(+)} - X_i^{(-)}) (Y_i^{(+)} - Y_i^{(-)}), \quad (\text{II.4})$$

from the values $X_i^{(+)}$ and $X_i^{(-)}$ that a quantity X takes for the parameter displacements along the (\pm) direction of the i -th eigenvector. With these symmetric master formulas, the sum of $S_{f,L_2}(E)$ over all experiments E should be within a few tens from zero, since the tolerance boundary for the total χ^2 is close to being spherically symmetric. The $S_{f,L_2}(E)$ variables for individual experiments tend to cancel among themselves to this accuracy; the order of magnitude of $S_{f,L_2}(E)$ can be also interpreted as a measure of tension of E against the rest of the experiments.

2.3 Application to parton luminosities

The L_2 sensitivity was explored in the CT18 paper [804], from which the description above is largely borrowed. It was applied to the determination of experimental sensitivities to specific PDFs, $f_a(x, Q)$, at a chosen factorization scale, Q , and as a function of the parton momentum fraction, x . For example, the L_2 sensitivity for the gluon distribution at a Q value of 100 GeV is shown in Fig. II.6. The pulls of a particular experiment on the gluon distribution can vary as a function of x . As stated previously, the larger the absolute value of $S_{f,L_2}(E)$, the greater

CT18 NNLO, $g(x, 100 \text{ GeV})$




Fig. II.6: The L_2 sensitivity of the most important experiments in the CT18 global PDF fit for the gluon distribution, $g(x, Q = 100 \text{ GeV})$, as a function of parton momentum fraction, x .

the sensitivity of that experiment to the determination of the PDF at that x (and Q) value. The L_2 sensitivity can be positive or negative. If positive, the upward variation of $f_a(x, Q)$ leads to an increase in the χ^2 for the specified experiment. A negative L_2 sensitivity indicates that the variation will lead to a decrease in χ^2 for this experiment. A large collection of figures illustrating L_2 sensitivities for various flavors of PDFs and parton luminosities in the CT18 and CT18Z NNLO analyses can be viewed online at [827].

So, for example, in Fig. II.6, at an x value near 0.01, a region sensitive to Higgs boson production through gluon-gluon fusion at 14 TeV, the strongest preference for a smaller gluon density [signaled by $\Delta\chi^2 > 0$ when $g(x, Q)$ is increased] comes from CDF jet data, F_2 measurements from CCFR, H1 heavy-flavor production, and the combined HERA1+II inclusive DIS data, followed by the ATLAS 7 TeV jet data. Descriptions of these and other quoted experimental data sets can be found in Ref. [804]. The two most important experiments that pull in the opposite direction are E866/NuSea, a fixed-target Drell-Yan experiment from Fermilab, and the ATLAS 8 TeV Z-boson p_T measurement. Note that only the most sensitive of the experiments for the determination of the gluon distribution have been plotted. There are 39 experimental data sets in the CT18 fit.

The CT18 analysis demonstrated using the L_2 sensitivity and other methods that the combination of the most extensive DIS experimental data sets – HERA, BCDMS, NMC, CCFR,...




Fig. II.7: Parton luminosities for processes at the LHC at $\sqrt{s} = 14$ TeV, in the central rapidity region $|y| < 5$: L_{gg} (upper-left), L_{gq} (upper-right), L_{qq} (lower-left), and $L_{q\bar{q}}$ (lower-right); evaluated using the CT18 (solid violet), CT18Z (short-dashed gray), and CT14_{HERAII} (long-dashed magenta) NNLO PDFs. In each instance, we display the luminosity ratios normalized to CT18.

– at the moment imposes the dominant constraints on the CT18 gluon PDF $g(x, Q)$ through Q dependence of DIS cross sections over a wide region of x and Q . At hadron-hadron colliders, the most sensitive measurements of $g(x, Q)$ are provided by inclusive jet production, especially by CMS and ATLAS. There can be no further improvement of the HERA data, or of E866/NuSea, but the importance shown by the LHC measurements provides an indication of where future, more precise, measurements at the LHC may improve the PDF uncertainties for the Higgs boson cross section, or for any other LHC measurement. Fig. II.6 indicates the L_2 sensitivities only at particular x values. This would correspond to one particular rapidity value for the Higgs boson, near zero. As Higgs bosons are produced over a reasonably wide rapidity range, production will be sensitive to a wide partonic x range, approximately, over $0.001 \lesssim x \lesssim 0.1$.

A more succinct understanding of the importance of each experiment to the production of a particle of a particular mass can be gained by showing the L_2 sensitivity to the parton-parton luminosity for a pair of initial partons a and b , defined as in [828] for production of a final state with invariant mass M_X at collider energy \sqrt{s} ; we apply an additional constraint that the rapidity of the final state, $y = \frac{1}{2} \ln(x_2/x_1)$, does not exceed y_{cut} in its absolute value, resulting in the parton luminosity definition:

$$L_{ab}(s, M_X^2, y_{cut}) = \frac{1}{1 + \delta_{ab}} \left[\int_{\frac{M_X}{\sqrt{s}} e^{-y_{cut}}}^{\frac{M_X}{\sqrt{s}} e^{y_{cut}}} \frac{d\xi}{\xi} f_a(\xi, M_X) f_b\left(\frac{M_X}{\xi\sqrt{s}}, M_X\right) + (a \leftrightarrow b) \right]. \quad (\text{II.5})$$

The uncertainty bands for the gluon-gluon, gluon-quark, quark-quark, and quark-antiquark lu-




Fig. II.8: The L_2 sensitivity of the most important experiments in the CT18 global PDF fit for the gg parton luminosity as a function of the mass of the final state. Here, we show the experimental pulls on the parton luminosity computed with a less restrictive rapidity cut, $|y_{cut}| < 5$, as compared with the $|y_{cut}| < 2.5$ selection more appropriate for LHC measurements shown in subsequent plots.

minisities at 14 TeV, as relevant for the LHC, are shown in Fig. II.7 based upon the CT18, CT18Z, and CT14_{HERAII} NNLO PDFs. The respective L_2 sensitivity for the gg parton luminosity can be viewed in Fig. II.8. Again, a more complete collection of L_2 sensitivities for parton luminosities can be viewed at [827].

Integrating over a larger range of parton x values, for the Higgs boson mass of 125 GeV, increases the importance for the HERAI+II data set (in the positive direction), with the L_2 sensitivity approaching a value of 6, and $\bar{\nu}$ DIS dimuon production [NuTvNbChXN] and the E866pp data (in the negative direction), with a value on the order of -5. BCDMS data on F_2^d and the ATLAS 8 TeV Z p_T distribution are also important on the negative side in this case. An astute reader will notice that the plot above was made by applying a rapidity cut of ± 5 on the produced Higgs boson. However, the precision coverage for ATLAS and CMS does not run past a rapidity of $|y| < 2.5$. A similar plot, but now imposing a rapidity cut of 2.5 is shown in Fig. II.10. A comparison between the two plots shows little difference, because most of the Higgs boson production in the gg fusion channel occurs within a rapidity of 2.5 in any case. From now on, a maximal rapidity cut of 2.5 will be applied.

All of the above L_2 sensitivity plots have been computed using the CT18 PDFs. It is instructive to also examine similar plots with CT18Z, which adds the precision ATLAS 7 TeV W/Z boson data to the fit, and, most importantly for the purposes of the gg parton luminosity, changes the scale used for low- x DIS production [804]. This has the impact of significantly increasing the low- x gluon distribution. The gg parton luminosity for CT18Z is shown in Figure II.10, and the changes leading to CT18Z have a marked effect on the pulls of the CT18Z experiments upon




Fig. II.9: The analog of the L_2 sensitivity plot for the gg parton luminosity shown in Fig. II.8, but in this case, calculated using a more restrictive rapidity cut of $|y_{cut}| < 2.5$.

L_{gg} . Most notably, this is true of the HERAI+II data, which under CT18Z exhibit pulls on the glue-gluon luminosity with significantly different dependence on M_X as compared to CT18. For instance, whereas the HERAI+II data resisted increases to the gluon distribution relevant for the lighter-mass, $M_X \lesssim 100$ GeV, region under CT18, for CT18Z, these pulls are essentially reversed, with the HERA data preferring the larger gluon at low x , leading to reductions in χ_E^2 in this light mass region. This feature is consistent with the large rise observed for the gluon PDF at low x with CT18Z relative to CT18 shown in Ref. [804]. In the immediate, $M_X \sim 125$ GeV, neighborhood of the Higgs production region, the HERA information has an L_2 sensitivity of approximately +5 under both fits.

Of course, we have mainly concentrated on the gg parton luminosity and its impact on the Higgs boson production. The technique can provide useful information for other mass values for the gg parton luminosity, and for other processes using other PDF luminosities. For example, the $q\bar{q}$ parton luminosity is plotted in Fig. II.11. At the mass of the W/Z boson, the primary influences in the positive direction are the NuTeV $\bar{\nu}$ data (NuTvNbChXN), BCDMSF2d, and CDHSWF2, and in the negative direction, the ATLAS 8 TeV $Z p_T$ distribution, LHCb 8 TeV W/Z data, NMC structure function ratios, and CMS 8 TeV jet data. For higher masses, on the order of 1 TeV, BCDMSF2d and CDHSWF2 are again most important on the positive side, while the HERAI+II experiment dominates in the negative L_2 direction.

It is also useful to examine the L_2 sensitivities for different categories of data. For example, in Fig. II.12, the L_2 sensitivities are shown for all experiments with an L_2 sensitivity exceeding 8 in some interval of M_X . It turns out that there are only two such experiments, the HERAI+II data and the CMS 8 TeV inclusive jet data, CMS8jtR7T. The other DIS data are added together (labeled as 1), as are all of the Drell-Yan data (2) and all of the $t\bar{t}$ and jets data (5). At the Higgs boson mass, the sum of all Drell-Yan data (2) has a pronounced pull in the negative direction,




Fig. II.10: The L_2 sensitivity of the most important experiments in the CT18Z global PDF fit for the gg parton luminosity.

in contrast to the sum of all $t\bar{t}$ + jets data (5) and HERAI+II, which pull more moderately in the opposing, positive direction. The other combined DIS data (excluding the inclusive HERA data) and CMS jet information have more modest pulls in this region. If we add all DIS data together, all Drell-Yan data together and all $t\bar{t}$ +jets data together, we get the result in Fig. II.13.

2.4 Conclusion

The L_2 sensitivity, plotted as a function of the invariant mass of the final state, is a useful indicator to understand the pulls on parton luminosity combinations from different experimental inputs, and the size of any tensions that may exist between experimental data sets, especially those from the LHC. The studies shown here have been created for the CT18 and CT18Z PDF sets. Comparable constructions for the other global PDF sets will help with the combination of such PDFs for the ongoing PDF4LHC20 benchmarking exercise. Ameliorating the tensions examined in this discussion will be critical to achieving the PDF precision required for the discovery program at the High-Luminosity LHC and beyond.

In addition to community benchmarking and other explorations in PDF fitting, future high-precision experiments will also be helpful. As an example, the Electron-Ion Collider (EIC) [829] will perform extremely precise measurements that are likely to substantially supersede the current fixed-target experimental data fitted in CT18. Independent EIC measurements will be valuable, for instance, given the competing pulls especially evident in Figs. II.8 and II.9 of the F_2^p and F_2^d data from BCDMS on L_{gg} in the $300 \text{ GeV} \lesssim M_X \lesssim 2 \text{ TeV}$ region. Precision measurements from the EIC will also have the potential to extend sensitivity to the higher-mass $M_X \gtrsim 2 \text{ TeV}$ region, where the L_2 sensitivities of current experiments are rapidly vanishing. By measuring inclusive cross sections with high precision over a wide sweep of x and Q , DIS experiments have the capacity to constrain scaling violations and provide access to the gluonic




Fig. II.11: The L_2 sensitivity of the most important experiments in the CT18 global PDF fit, as in Fig. II.9, but for the $q\bar{q}$ parton luminosity.

structure of the nucleon or of nuclei. For the EIC, the expected luminosities (10^2 – 10^3 times that of HERA) are sufficiently great that the resulting improvements in the gluon PDF can in turn significantly reduce the PDF uncertainty on the LHC $gg \rightarrow$ Higgs production. This conclusion has been demonstrated by computing the L_1 sensitivity of the EIC pseudodata to the PDF uncertainty of the 14 TeV Higgs-production cross section, as presented in the right panel of Fig. 2 in Ref. [830].

This work is partially supported by the U.S. Department of Energy under Grant No. DE-SC0010129 and by the U.S. National Science Foundation under Grant No. PHY-1719914. T. J. Hobbs acknowledges support from a JLab EIC Center Fellowship.




Fig. II.12: The L_2 sensitivity to the gg parton luminosity of all experiments with an L_2 sensitivity greater than 8, plus the combination of all other DIS data, all Drell-Yan data, and all $t\bar{t}$ + jets data.




Fig. II.13: The L_2 sensitivity to the gg parton luminosity of all the data fitted in CT18, now collected into categories for the DIS, Drell-Yan, and jets+ $t\bar{t}$ data.