

Closure Testing the MSHT PDFs and a First Direct Comparison to the Neural Net Approach

L.A. Harland-Lang,^{a,*} R.S. Thorne^a and T.Cridge^b

^a*Department of Physics and Astronomy, University College London, London, WC1E 6BT, UK*

^b*Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany*

E-mail: robert.thorne@ucl.ac.uk, thomas.cridge@desy.de,

l.harland-lang@ucl.ac.uk

We present a brief overview of the first global closure test of the fixed parameterisation (MSHT) approach to PDF fitting. We find that the default MSHT20 parameterisation can reproduce the features of the input set in such a closure test to well within the textbook uncertainties. This provides strong evidence that parameterisation inflexibility in the MSHT20 fit is not a significant issue in the data region. We also present the first completely like-for-like comparison between two global PDF fits, namely MSHT and NNPDF, where the only difference is guaranteed to be due to the fitting methodology. We find that this gives a moderately, but noticeably, better fit quality than the central NNPDF4.0 fits and that this difference persists at the level of the PDFs and benchmark cross sections. The NNPDF4.0 uncertainties are found to be broadly in line with the MSHT results if a textbook $T^2 = 1$ tolerance is applied, but to be significantly smaller if a tolerance typical of the MSHT20 fit is applied.

DIS2024

8-12 April 2024

Grenoble, France

*Speaker

1. Introduction

Global PDF fits now combine high precision theory, such that NNLO is the standard and now even approximate N³LO QCD precision is being accounted for [1, 2], with a wealth of data from multiple experiments and, increasingly, the LHC. Combing these ingredients, we can hope to constrain proton structure with high levels of precision and accuracy. However, as well as differing in the data and the relevant theoretical ingredients that enter the fit, these global PDF analyses also differ rather significantly in the methodologies they apply. Moreover, these methodological effects alone are observed to have a significant impact on the resulting PDFs and, crucially, their uncertainties.

It is the aim of the study presented here, and described in more detail in [3], to begin to address this issue directly. Namely, we present the first full global closure test of a fixed parameterisation approach to PDF fitting, focusing on the MSHT20 [4] case, and the first completely like-for-like comparison between two global PDF fits, namely MSHT and NNPDF [5], where the only difference is guaranteed to be due to the fitting methodology, and PDF parameterisation in particular. We present a brief selection of results relating to these two studies below.

2. Closure Testing the MSHT Framework

This study makes use of the publicly available NNPDF fitting code [6]. By developing a suitable interface to this code, it is possible for us to evaluate the fit quality within the NNPDF framework (i.e. with their theory and choice of dataset, or any subset of it) but for an arbitrary PDF set. Here, we will take this to be parameterised using the fixed basis of Chebyshev polynomials used by the MSHT collaboration. This interface then also allows for a PDF fit (either as a closure test or to real data) to be performed to the NNPDF theory and datasets, but whilst using the MSHT parameterisation of the underlying PDFs at the input scale Q_0 .

Two types of closure test are performed, namely ‘unfluctuated’ and ‘fluctuated’, where the pseudodata are not, and are fluctuated by the experimental uncertainties, respectively. For both cases we generate pseudodata corresponding to the NNPDF4.0 global dataset, using the central replica of the NNPDF4.0 (perturbative charm) set [5]. In the unfluctuated test, if the PDF parameterisation in the fit is flexible enough, then a perfect fit (with $\chi^2 = 0$) is in principle achievable, and any deviations from this will provide an assessment of any lack of flexibility in the fixed MSHT20 parameterisation. We find a χ^2/N_{pts} of 0.0005, i.e. a very close level of agreement and indeed somewhat lower than that found in the corresponding NNPDF4.0 tests [7].

A representative comparison between the fit PDFs and the input is shown in Fig. 1. To guide this, PDF uncertainties according to the textbook $\Delta\chi^2 = T^2 = 1$ and the enlarged $T^2 = 10$ case, similar in size to that taken by MSHT, are indicated. The agreement between the input and fit is found to be very good in the data region, at the per mille level, corresponding to $\sim 10\%$ of the overall $T^2 = 1$ uncertainty. In the extrapolation region at high x (and in the poorly constrained low to intermediate x valence distributions, not shown) the agreement is less good, as we might expect.

We next consider some results of the fluctuated closure test. We in particular generate 100 pseudodata sets due to the same NNPDF4.0 (p. charm) input set as before, in each case with the pseudodata fluctuated with a different random number seed. Fitting the ensemble of these we can therefore generate a set of 100 fits with which to perform the closure test. While the unfluctuated

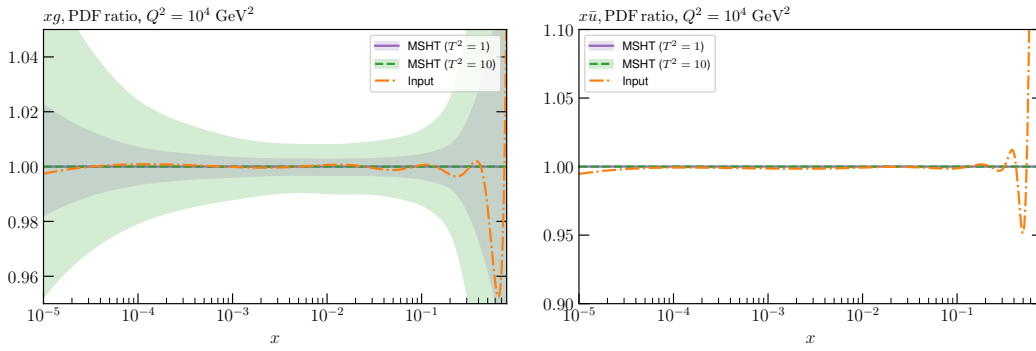


Figure 1: A selection of PDFs at $Q^2 = 10^4 \text{ GeV}^2$ that result from an unfluctuated closure test fit to the NNPDF4.0 dataset, using the MSHT20 parameterisation. The PDF uncertainties calculated with a $T^2 = 1$ ($T^2 = 10$) fixed tolerance are shown in purple (green) and the NNPDF4.0 (p. charm) input is given by the dashed orange line. Results are shown as a ratio to the MSHT fits.

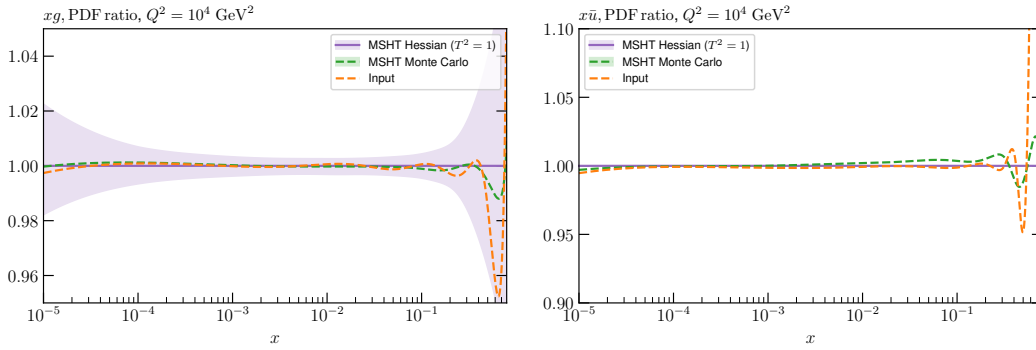


Figure 2: A selection of PDFs at $Q^2 = 10^4 \text{ GeV}^2$ that result from a fluctuated closure test fit to the NNPDF4.0 dataset, using the MSHT20 parameterisation. The Hessian PDF uncertainties calculated with a $T^2 = 1$ fixed tolerance that result from the unfluctuated closure test are shown in purple (and are as in Fig. 1), while the result of fluctuated closure test are shown in green. The NNPDF4.0 (p. charm) input is given by the dashed orange line. Results are shown as a ratio to the MSHT ($T^2 = 1$) fit.

test assesses the faithfulness of the MSHT fixed parameterisation approach in fitting the central value of the input set, the fluctuated test assesses the faithfulness of the PDF uncertainties that come out of the fit. In particular, if the uncertainties are faithful, then we expect the input PDF set to lie within the 1σ uncertainty band of the fitted PDFs with close to 68% frequency across the fits.

A visual indication that this is achieved is shown in Fig. 2, which compares the result of the unfluctuated fit to the 68% C.L. ensemble of best fits to the 100 fluctuated pseudodatasets. We can see that in the data region the matching between the MC replica and Hessian result is rather close. Combining this with the observation that the Hessian uncertainty on any of the 100 fluctuated fits is very stable, and closely matches this $T^2 = 1$ Hessian uncertainty on the unfluctuated fit, is then sufficient to demonstrate that the closure test has been passed at the level of each individual PDF and at any given x point in the data region. In other words, each of these 100 fits will be

	NNPDF4.0 pch	MSHT fit	MSHT fit (w positivity)
Perturbative Charm (4626)	5928.3 (1.282)	5736.7 (1.240)	5837.8 (1.262)
Fitted Charm, (4666)	5692.1 (1.233)	5645.2 (1.222)	5651.0 (1.224)

Table 1: Global χ^2 (t_0 definition) values for the NNPDF4.0 fit and the MSHT fits to the NNPDF dataset/theory settings, with both perturbative and fitted charm, and with and without positivity imposed for the MSHT fits.

consistent within its 1σ Hessian uncertainty with the input set with close to 68% frequency.

3. A Direct Comparison to NNPDF4.0

We now fit the real data corresponding to the same NNPDF4.0 dataset; we will refer to these as ‘MSHT fits’ for brevity. We recall in particular that both the data in the fit, but also the theory settings, are completely identical to those used in the corresponding NNPDF fit. The comparison is therefore completely like–for–like at the level of the full global PDF fit, with the only difference being due to the PDF parameterisation. The global fit qualities are shown in Table 1, for both perturbative and fitted charm. In the latter case the charm quark PDF is freely parameterised using the same basis of Chebyshev polynomials taken for other PDFs in the MSHT analysis, and this represents the first model independent fitted charm determination outside of the NNPDF approach.

We can see that, despite the nominally larger flexibility of the NNPDF4.0 fit, the MSHT fit produces a moderately, but noticeably, better fit quality than the central NNPDF results, both with perturbative and fitted charm. In the former case, this can in part be explained by positivity of the PDFs imposed by NNPDF, in particular of the gluon at low x and Q^2 , and which is not by default imposed in the MSHT fit. Accounting for this leads to some deterioration in the perturbative charm fit, but with a result that remains better than the NNPDF value. In [3] the question of why this occurs is considered in detail, and the possibility that the MSHT parameterisation might lead to overfitting closely examined, but no evidence of this found. Conversely, it is demonstrated that if a more restrictive parameterisation is taken, then the closure tests described in the previous section are not passed successfully, and a poorer description in the full fit is arrived at.

We also find some significant differences in the PDFs, such that those due to the NNPDF4.0 releases are not always compatible with these MSHT fits within the nominal NNPDF uncertainties. These differences tend to be larger in the perturbative charm fits, but are also present with fitted charm. An example of this is shown in Fig. 3. In the left plot the ratio to the NNPDF result is shown, with $T^2 = 1$ uncertainties taken in the MSHT results in order to demonstrate the effect more clearly, and a clear difference is observed. In the right plot the PDF uncertainties are given, and we can see that the NNPDF uncertainty is similar in size to the MSHT case, if $T^2 = 1$ is taken. Similar results are found for other PDFs, most notably with respect to the quark flavour decomposition.

The PDF uncertainty comparison in particular demonstrates that the NNPDF4.0 fit, as a result of changes in methodology alone, produces PDF uncertainties that are rather more closely in line with those of the fixed MSHT20 parameterisation if the textbook $T^2 = 1$ criterion is applied. This at a minimum points to an inherent inconsistency between the approaches, although as discussed in more detail in [3] there are very good reasons for including an enlarged, $T^2 > 1$, tolerance in

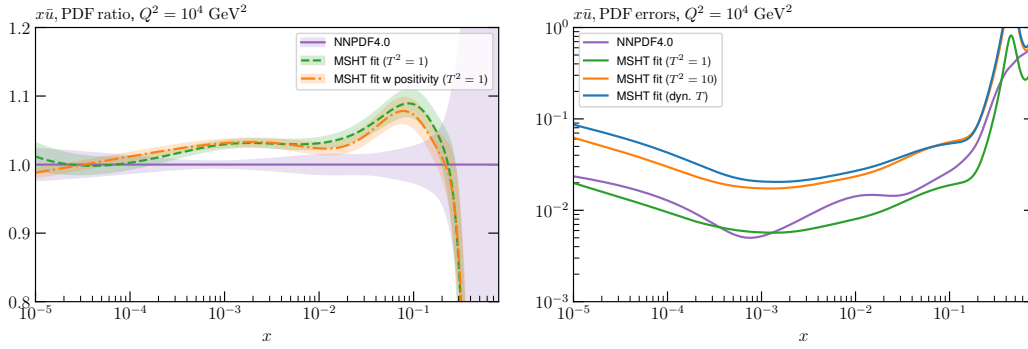


Figure 3: The \bar{u} at $Q^2 = 10^4 \text{ GeV}^2$ that results from a global PDF fit to the NNPDF4.0 dataset/theory (fitted charm) setting, but using the MSHT20 parameterisation. Results with and without a positivity constraint applied are shown. In the left plot the ratio to the NNPDF4.0 set is shown, with PDF uncertainties for the MSHT fits corresponding to a fixed $T^2 = 1$ tolerance. In the right plot the corresponding PDF uncertainties are shown, for different fixed and dynamic tolerances in the MSHT case.

global PDF fits. Evidence for this is demonstrated in this work at the level of closure tests, where it is shown that if data/theory inconsistencies are included in the closure test, then these do not lead to any increase in the $T^2 = 1$ PDF uncertainty.

In summary, we have presented a selection of results from the first global closure test of the fixed parameterisation (MSHT) approach to PDF fitting and first completely like-for-like comparison between two global PDF fits, namely MSHT and NNPDF. We find that the default MSHT20 parameterisation can reproduce the features of the input set in such a closure test well, while for the full fit we find a moderately, but noticeably, better fit quality than the central NNPDF4.0 fits.

References

- [1] J. McGowan, T. Cridge, L. A. Harland-Lang and R. S. Thorne, *Eur. Phys. J. C* **83** (2023) no.3, 185 doi:10.1140/epjc/s10052-023-11236-0 [arXiv:2207.04739 [hep-ph]].
- [2] R. D. Ball *et al.* [NNPDF], *Eur. Phys. J. C* **84** (2024) no.7, 659 doi:10.1140/epjc/s10052-024-12891-7 [arXiv:2402.18635 [hep-ph]].
- [3] L. A. Harland-Lang, T. Cridge and R. S. Thorne, [arXiv:2407.07944 [hep-ph]].
- [4] S. Bailey, T. Cridge, L. A. Harland-Lang, A. D. Martin and R. S. Thorne, *Eur. Phys. J. C* **81** (2021) no.4, 341 doi:10.1140/epjc/s10052-021-09057-0 [arXiv:2012.04684 [hep-ph]].
- [5] R. D. Ball *et al.* [NNPDF], *Eur. Phys. J. C* **82** (2022) no.5, 428 doi:10.1140/epjc/s10052-022-10328-7 [arXiv:2109.02653 [hep-ph]].
- [6] R. D. Ball *et al.* [NNPDF], *Eur. Phys. J. C* **81** (2021) no.10, 958 doi:10.1140/epjc/s10052-021-09747-9 [arXiv:2109.02671 [hep-ph]].
- [7] L. Del Debbio, T. Giani and M. Wilson, *Eur. Phys. J. C* **82** (2022) no.4, 330 doi:10.1140/epjc/s10052-022-10297-x [arXiv:2111.05787 [hep-ph]].