

# A new generation of simultaneous fits to LHC data using deep learning

---

**Shayan Iranipour and Maria Ubiali**

*DAMTP, University of Cambridge,  
Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

*E-mail:* [s.iranipour@damtp.cam.ac.uk](mailto:s.iranipour@damtp.cam.ac.uk), [m.ubiali@damtp.cam.ac.uk](mailto:m.ubiali@damtp.cam.ac.uk)

**ABSTRACT:** We present a new methodology that is able to yield a simultaneous determination of the Parton Distribution Functions (PDFs) of the proton alongside any set of parameters that determine the theory predictions; whether within the Standard Model (SM) or beyond it. The *SIMUnet* methodology is based on an extension of the *MNPDF4.0* neural network architecture, which allows the addition of an extra layer to simultaneously determine PDFs alongside an arbitrary number of such parameters. We illustrate its capabilities by simultaneously fitting PDFs with a subset of Wilson coefficients within the Standard Model Effective Field Theory framework and show how the methodology extends naturally to larger subsets of Wilson coefficients and to other SM precision parameters, such as the strong coupling constant or the heavy quark masses.

**KEYWORDS:** Parton Distributions, SMEFT

**ARXIV EPRINT:** [2201.07240](https://arxiv.org/abs/2201.07240)

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Theoretical predictions in a simultaneous fit</b>	<b>4</b>
2.1	Fast interface for theory predictions	4
2.2	Observable dependence on the SMEFT Wilson coefficients	6
<b>3</b>	<b>Methodology</b>	<b>8</b>
3.1	Neural network design	8
3.2	Parameter fitting using linearisation	10
3.3	Incorporating non-linear effects	12
3.4	Fixed PDF analysis	14
<b>4</b>	<b>Results</b>	<b>14</b>
4.1	Benchmark scenarios	15
4.2	Experimental data	16
4.3	Results for Benchmark Scenario I	16
4.4	Inclusion of the HL-LHC projections	19
4.5	Results for Benchmark Scenario II	22
4.6	Results overview	24
<b>5</b>	<b>Methodology validation and closure testing</b>	<b>26</b>
5.1	Closure test results on the Wilson coefficients	27
5.2	Closure test results on the simultaneous fit	29
<b>6</b>	<b>Conclusions and outlook</b>	<b>30</b>
<b>A</b>	<b>FK table dependence on the strong coupling</b>	<b>32</b>
<b>B</b>	<b>Fit quality</b>	<b>34</b>
<b>C</b>	<b>Stability on replica number</b>	<b>37</b>

---

## 1 Introduction

The successful operation of the Large Hadron Collider (LHC) at CERN is enabling us to scrutinise the fundamental laws of nature to an unprecedented degree and for a wide range of energy scales. Despite many indications that the Standard Model (SM) of particle physics cannot be a complete description of nature and despite theoretical arguments pointing towards physics beyond the Standard Model (BSM) at the TeV scale, no direct evidence for new physics at the TeV scale has been gathered so far at colliders.

Far from being discouraging, it is an exciting time for particle physics: the precision level reached by the LHC experiments gives us the unique chance of investigating the effects of new particles whose masses are far above the TeV scale, but still produce observable effects at the scales within the direct kinematical reach of the LHC. Unlike for direct searches, which are limited by the energy reach of the collider, indirect searches are limited only by the theoretical and experimental control over the processes under inspection. As such, they are becoming increasingly relevant at the LHC as more data is collected and better theoretical predictions are devised.

Despite the broad consensus on the need for precision, the paradigm of indirect searches is often reduced to performing better measurements and improving the accuracy of theoretical predictions. However, for an indirect detection of new physics, it is pivotal to have a robust framework that is able to globally interpret all subtle deviations from the SM predictions that might arise. While huge progress has been made in determining key ingredients of theoretical predictions from the data [1], such as the Parton Distribution Functions (PDFs) of the proton [2–5], the strong coupling constant [6], and the coefficients of a suitable parametrisation of the effects of heavy new states via the addition of higher dimensional operators to the SM Lagrangian, such as the SMEFT [7, 8], it is not yet evident how to combine all these partial fits into a global interpretation of the LHC data.

These simultaneous determinations are sometimes essential to retrieve a correct interpretation of the LHC data. For example, in ref. [9] it was shown that any determination of the strong coupling constant  $\alpha_s$  from a process which depends on PDFs, such as hadronic processes or deep-inelastic scattering (DIS), generally does not lead to a correct result unless the PDFs are determined simultaneously along with  $\alpha_s$ . The case of the determination of the strong coupling constant is particularly relevant because of its strong correlation with the PDFs. However, similar considerations apply to the simultaneous determination of any physical parameter in PDF-dependent processes, such as for example the determination of the top quark mass [10], or of the electroweak (EW) parameters, such as the  $W$ -mass [11]. In the latter case, the correlation of PDFs and the EW parameters is in principle weaker than in the case of the strong coupling constant, but the very high accuracy which is sought suggests that currently available results, specifically in  $W$ -mass determination, should be reconsidered with care [12]. Going beyond unpolarised PDFs, if more exclusive or spin-dependent observables were included in a fit, then the simultaneous extraction of polarised and unpolarised PDFs and of fragmentation functions and PDFs are the next obvious frontiers. In the former case, the analysis of semi-inclusive measurements involving the production of an identified hadron in the final state employs observables that are often presented as ratios of spin-dependent to spin-averaged cross sections, thus unpolarised and polarised PDFs should be treated simultaneously in an universal QCD analysis [13, 14]. In the second case, the analysis of semi-inclusive measurements involving the production of an identified hadron in the final state can be accurately performed only if the final-state fragmentation functions are determined alongside the initial-state PDFs [13, 15].

Recent findings indicate that caution should be taken not only in the determination of SM precision parameters from hadronic processes, but also for the Wilson coefficients that parametrise the effects of heavy new physics via Effective Field Theory (EFT) expan-

sions. Such extensions of the SM Lagrangian determine the effect of physics, that lives well above the energy scale reached by the LHC, by adding higher dimensional operators to the SM Lagrangian, whose coefficients are suppressed by powers of the new physics scale. Although the proton structure parametrised by PDFs is intrinsically a low-energy quantity and, as such, it should in principle be separable from the high-energy new physics imprints, the complexity of the LHC environment might well intertwine them, first and foremost because the same high-energy data is used to constrain both the PDFs and the SMEFT parametrisation. The effects of a simultaneous determination of the Wilson coefficients of the SMEFT [16] and of the proton PDFs has been pioneered in several recent studies [17–20]. These studies reveal that, while with current DIS and Drell-Yan data the interplay can be kept under control, once High Luminosity LHC (HL-LHC) data are considered, neglecting the PDF interplay could potentially miss new physics manifestations or misinterpret them.

In this paper we present a new methodology, dubbed **SIMUnet**, which allows for a truly simultaneous determination of the PDFs alongside any physical parameter that enters theoretical predictions, whether a precision SM parameter, or the Wilson coefficients of some EFT expansion. **SIMUnet** is based on an extension of the **n3fit** methodology and the **NNPDF4.0** neural network architecture [2, 21], which treats both the PDFs and the parameters fitted alongside PDFs on a completely equal footing. We employ analogous deep learning techniques to **NNPDF4.0**, however, here we introduce an additional layer to the network that captures the sensitivity of the data to the parameters considered, leaving the first two layers of the architecture to capture the data dependence upon the underlying PDFs. This work was made possible thanks to the open-source availability of the full **NNPDF4.0** framework [21].

To display the potential of **SIMUnet**, in this work we apply it on the simultaneous fit of the PDFs and of the SMEFT Wilson coefficients considered in ref. [17]. The rationale behind this choice stems from the fact that the parameter dependence for the linear and quadratic contributions of the Wilson coefficients of the dim-6 SMEFT expansion can be easily incorporated within our framework. Furthermore, the results of this study can be benchmarked against those obtained earlier, thus also assessing the effect of using a methodology that overcomes the limitations of the methodology employed in ref. [17].

In what follows, in section 2 we outline a general framework that allows for the dependence of the theoretical predictions on the parameters considered to be obtained in an efficient way. We also describe how a simple method using  $K$ -factors can be employed in specific cases, such as EFT expansions. In section 3 we outline how our methodology works and how the neural network architecture can exploit the fast interface discussed in section 2 so as to constrain external parameters alongside PDFs. In section 4 we turn to discuss the results we obtain, showcasing the methodology’s ability to use current high-mass Drell-Yan (DY) tails from the LHC to constrain not only the PDF, but also to simultaneously determine the best fit Wilson coefficients to the data. We demonstrate how our methodology can recover degenerate (flat) directions in EFT space and how the introduction of HL-LHC projected data is able to eliminate this flat direction. In section 5 we stress test our proposed methodology using the closure testing strategy, whereby we artificially contaminate

the input data with a particular choice of Wilson coefficients and demonstrate how our approach is not only able to retrieve this choice, but is sufficiently robust to also simultaneously replicate a known underlying PDF law. Finally, in section 6 we conclude this study and discuss possible future avenues using the methodology presented in this paper.

## 2 Theoretical predictions in a simultaneous fit

In this section we outline a general formalism that allows for the fast interface of theory predictions and to isolate their dependence on the physical parameters that we may want to fit simultaneously with the PDFs. We then focus on simplifying this approach through the use of multiplicative  $K$ -factors before discussing in the next section how we extend the NNPDF4.0 framework to fit PDFs along with other physical parameters.

### 2.1 Fast interface for theory predictions

A formalism that is widely adopted in order to obtain a fast interface of theory predictions in a PDF fit is the so-called Fast Kernel (FK) table approach. The latter was formulated in refs. [22–24] and it allows for the convolution of PDFs with the partonic cross section in an efficient way by reducing the non-trivial integral convolution to a tensor product using pre-computed integrals on interpolation grids. Moreover, encoded within the FK table, is the evolution kernel operator which performs the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution of the PDFs, allowing for the Bjorken- $x$  dependence of the PDF to be fitted at fixed scale,  $Q_0$ , before evolving the factorization scale to the relevant kinematic scale,  $Q$ , in a fast way. Using the notation employed in [23], we can write the theoretical prediction for any hadronic cross section as

$$T_I^{\text{hh}} = \sum_{i,j=1}^{N_{\text{pdf}}} \sum_{\alpha,\beta=1}^{N_x} \Sigma_{\alpha\beta ij}^I N_{\alpha i}^0 N_{\beta j}^0 \equiv N^0 \cdot \Sigma_I \cdot N^0, \tag{2.1}$$

where  $I$  indicates a specific hadronic observable included in a PDF fit;  $\alpha, \beta$  are the indices of the interpolation grids in  $x$ -space for the first and second parton respectively;  $i, j$  are the indices of the PDFs of the initial-state partons that contribute to the observable  $I$  and  $N^0$  are the (neural-network) parametrisation of the independent PDFs at the initial scale  $Q_0$ . The computation of the hadronic observables is reduced to a bilinear product over an interpolation grid in the  $x_{1,2}$  space and the basis of the input PDFs for a given process. The quantity  $\Sigma$  is the FK-table, which incorporates both the evolution of the PDFs from the initial scale to the scale of the measured observable and the partonic cross sections associated to each of the partonic channels that enter the computation of the hadronic cross section. For processes involving only one hadron and one lepton, like DIS, the expression is even simpler and reads:

$$T_I^{\text{hl}} = \sum_{i=1}^{N_{\text{pdf}}} \sum_{\alpha=1}^{N_x} \Sigma_{\alpha i}^I N_{\alpha i}^0 \equiv \Sigma_I \cdot N^0. \tag{2.2}$$

We can collectively refer to the theory prediction for a generic observable — whether it involves one or two hadrons in the initial states — as

$$T_I = \Sigma_I \cdot L^0, \tag{2.3}$$

where  $L^0$  indicates either the parametrisation of one independent PDF at the initial scale or the product of two of them.

If we now consider how the theoretical predictions  $T_I$  for each observable  $I$  explicitly depend upon a single parameter  $c$ , that could be for example the strong coupling constant,  $\alpha_s$ , or the top mass,  $m_t$ , the entire dependence upon this parameter is contained in the FK-tables  $\Sigma$ , since  $L^0$  only captures the initial scale parametrisation of PDFs, which is fitted from the data. For clarity of notation we focus here on the case where only one single parameter is fitted alongside PDFs, but the generalisation to more parameters is a straightforward extension of the following argument, simply requiring the use of multivariate Taylor expansions. Schematically we can write the FK-tables as

$$\Sigma_I(c) = [\hat{\sigma}(c) \otimes \Gamma(c)]_I \tag{2.4}$$

where  $\hat{\sigma}$  is the partonic cross section and  $\Gamma$  are the evolution kernels that evolve the PDFs from the initial scale to the scale of  $T_I$ . The shorthand  $\otimes$  denotes the usual convolution, where for sufficiently well behaved functions  $f, g$  we have

$$(f \otimes g)(x) = \int_x^1 \frac{dy}{y} f(y) g\left(\frac{x}{y}\right). \tag{2.5}$$

For a standard PDF-only determination, the parameter  $c$  is typically fixed to a certain value during the computation of the FK table. In general both  $\hat{\sigma}$  and  $\Gamma$  depend on the given parameter (the strong coupling being one such example), whilst in some other cases only the partonic cross sections,  $\hat{\sigma}$ , depends on the parameter under consideration (the Wilson coefficients of the SMEFT expansion being one such example).

If we now want to fit the parameter  $c$  alongside the PDFs, we need a fast interface to the dependence of each  $\Sigma_I(c)$  upon the parameter  $c$ .<sup>1</sup> One possible way to achieve such a fast interface is to assume that both the evolution kernel and the partonic cross section are suitably analytic such that they can be accurately described by their Taylor expansion about some point  $c^*$ . The closer the point  $c^*$  is to the actual parameter, the greater the validity of truncating the Taylor series. Dropping from now on the observable index  $I$ , we can write

$$\begin{aligned} \Sigma(c) &= \sum_{p,q} \frac{(c - c^*)^{p+q}}{p!q!} \frac{\partial^p \hat{\sigma}(c^*)}{\partial c^p} \otimes \frac{\partial^q \Gamma(c^*)}{\partial c^q} \\ &= \sum_k (c - c^*)^k \sum_{\substack{p,q: \\ p+q=k}} \frac{1}{p!q!} \frac{\partial^p \hat{\sigma}(c^*)}{\partial c^p} \otimes \frac{\partial^q \Gamma(c^*)}{\partial c^q} \\ &= \sum_k (c - c^*)^k \Sigma_k(c^*), \end{aligned} \tag{2.6}$$

---

<sup>1</sup>In section 3 we will see that the reason for this is that gradient descent will assess the optimal value of  $c$  at every step during learning by repeatedly evaluating theory predictions with different  $c$  values.

whereby for each power of  $c$  we have an order-by-order FK table,  $\Sigma_k(c^*)$ , that can be pre-computed before the fit, with the dependence on the parameter  $c$  being isolated from the FK table. In this way the task of querying the FK table for various values of  $c$  has been reduced to computing individual FK tables for each order and taking a weighted sum of these tensors which is a computationally trivial and fast operation.<sup>2</sup>

To make the above discussion more concrete, in appendix A we explicitly consider the case whereby we wish to isolate the FK table dependence on  $\alpha_s$ . The strong coupling constant  $\alpha_s(\mu_r)$  and its evolution with the scale from the reference value  $\alpha_s(m_Z)$  to the renormalization scale associated with each observable  $\mu_r \sim Q$ , with  $Q$  being the scale of the hard process, appears both in DGLAP evolution kernel  $\Gamma$  and in the partonic cross sections  $\hat{\sigma}$ . In appendix A we show explicitly how the dependence on  $\alpha_s(m_Z)$  and its evolution from  $m_Z$  to  $\mu_r \sim Q$  are combined in the FK tables. Then, following our above notation, we identify  $c = \alpha_s(m_Z)$  and  $c^* = \alpha_s^{\text{PDG}}(m_Z) = 0.1179(10)$  is the PDG value for the strong coupling evaluated at the  $Z$ -boson mass [6] and explicitly derive the first two terms in the Taylor expansion of eq. (2.6) in the case of  $\alpha_s$ . In the same appendix we also mention the approach one would have to use to Taylor expand about the central values of the electroweak parameters that enter a PDF fit.

## 2.2 Observable dependence on the SMEFT Wilson coefficients

We now focus on the parameters that we are specifically considering in this work, to display the potential of our approach. Specifically, we are interested in fitting  $N$  parameters  $\{c_n\}$  with  $n = 1, \dots, N$ , each of these parameters associated to the Wilson coefficient of a given operator in the SMEFT expansion,  $\mathcal{O}_n$ . We adopt an operator normalisation such that

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \sum_{n=1}^N \frac{c_n}{v^2} \mathcal{O}_n, \tag{2.7}$$

with  $N$  indicating in this specific case the number of operators that contribute to a given observable and  $c_n$  being the (dimensionless) Wilson coefficient associated to  $\mathcal{O}_n$ . The quantity  $v$  refers to the energy scale of new physics, making the Wilson coefficients dimensionless.

To include the effects of the corrections coming from the dim-6 operators included in the SMEFT expansion in eq. (2.7) in the theoretical prediction  $T_I$  for any observable included in the fit, one has to augment the SM partonic cross sections with the effects of the relevant operators with the linear and quadratic modifications of the SM cross section that the operators induce. Note that here we do not include the RGE running of the Wilson coefficients, rather we treat them as fixed numbers. This is justified by the fact that the anomalous dimensions of the operators that we consider in this work are either Yukawa-suppressed or suppressed by NLO EW corrections [17]. To include the scale dependence of the Wilson coefficients [26–28], we would have to fit the scale dependent Wilson coefficients at some fixed scale and use the relevant  $\beta$ -functions to evolve the operator to the relevant

---

<sup>2</sup>Importantly, such an operation can be implemented using the purely `TensorFlow` functionality already present in the `n3fit` methodology, thus allowing for the gradients to be computed using automatic differentiation techniques [25].

scale using some pre-computed evolution kernel, which would be factored into the FK-tables, exactly as is done for the PDF evolution. Furthermore, given that the SMEFT operators that we consider here do not modify the PDF DGLAP evolution, the corrections will only appear in the partonic cross sections  $\hat{\sigma}$  and, unlike in the case of the Taylor expansion around the PDG value of  $\alpha_s$  discussed in appendix A, here the sum is exact (and not approximated) when the Taylor expansion is truncated at order  $k = 2$  (linear and quadratic corrections). The linear SMEFT corrections for each observable  $I$  (from now omitting the index  $I$ ) can be parameterised as

$$R_{\text{SMEFT}}^{(n)} \equiv \left( \mathcal{L}_{ij}^{\text{NNLO}} \otimes d\hat{\sigma}_{ij,\text{SMEFT}}^{(n)} \right) / \left( \mathcal{L}_{ij}^{\text{NNLO}} \otimes d\hat{\sigma}_{ij,\text{SM}} \right), \quad n = 1 \dots, N, \quad (2.8)$$

with  $\mathcal{L}_{ij}^{\text{NNLO}}$  being the partonic luminosity evaluated at NNLO QCD

$$\mathcal{L}_{ij}(\tau, M_X) = \int_{\tau}^1 \frac{dx}{x} f_i(x, M_X) f_j(\tau/x, M_X), \quad (2.9)$$

$d\hat{\sigma}_{ij,\text{SM}}$  the bin-by-bin partonic SM cross section, and  $d\hat{\sigma}_{ij,\text{SMEFT}}^{(n)}$  the corresponding partonic cross section associated to the interference between  $\mathcal{O}_n$  and the SM amplitude  $\mathcal{A}_{\text{SM}}$  when setting  $c_n = 1$ . Likewise, the ratio encapsulating the quadratic effects is defined as

$$R_{\text{SMEFT}}^{(n,m)} \equiv \left( \mathcal{L}_{ij}^{\text{NNLO}} \otimes d\hat{\sigma}_{ij,\text{SMEFT}}^{(n,m)} \right) / \left( \mathcal{L}_{ij}^{\text{NNLO}} \otimes d\hat{\sigma}_{ij,\text{SM}} \right), \quad n, m = 1 \dots, N, \quad (2.10)$$

with the bin-by-bin partonic cross section  $d\hat{\sigma}_{ij,\text{SMEFT}}^{(n,m)}$  now being evaluated from the squared amplitude  $\mathcal{A}_n \mathcal{A}_m$  associated to the operators  $\mathcal{O}_n$  and  $\mathcal{O}_m$  when  $c_n = c_m = 1$ . The partonic cross sections in these ratios are computed at LO in QCD.<sup>3</sup> In terms of eqs. (2.8) and (2.10), we can define the EFT  $K$ -factors as

$$K(\{c_n\}) = 1 + \sum_{n=1}^N c_n R_{\text{SMEFT}}^{(n)} + \sum_{n,m=1}^N c_n c_m R_{\text{SMEFT}}^{(n,m)}, \quad (2.11)$$

which allow us to express a general cross section accounting for the dim-6 operators in eq. (2.7) as

$$T = T^{\text{SM}} \times K(\{c_n\}) \quad (2.12)$$

where the  $T$  is the SMEFT-modified theoretical prediction,  $T^{\text{SM}}$  is the state-of-the-art SM theoretical prediction including NNLO QCD and NLO EW corrections and  $K(\{c_n\})$  are the SMEFT  $K$ -factors defined in eq. (2.11). The coefficients associated with the linear (quadratic) corrections  $R_{\text{SMEFT}}^{(n)}$  ( $R_{\text{SMEFT}}^{(n,m)}$ ) in eq. (2.11) can be precomputed before the fit using a reference PDF set and then kept fixed<sup>4</sup> and the coefficients  $\{c_n\}$  can be fitted alongside the PDF parameters by multiplying the FK-tables,  $\Sigma$ , for each measurement

<sup>3</sup>Notice that, given that the SMEFT corrections are then multiplied by the SM predictions in eq. (2.11), which do include NNLO QCD and NLO EW corrections, the SMEFT  $K$ -factors inherit factorisable higher-order radiative corrections [29, 30]. To account for the full NLO QCD effects in the SMEFT one should include the relevant diagrams in the partonic cross section computation.

<sup>4</sup>The effect of varying the input NNLO PDF in eqs. (2.8) and (2.10) was quantitatively assessed in ref. [17] and it was found to be at most at the percent level.



included in the fit by the  $K$ -factors encapsulating the whole dependence on the parameters that we want to fit alongside the PDFs. Schematically this reads

$$\Sigma(\{c_n\}) = [\hat{\sigma} \otimes \Gamma] \times K(\{c_n\}) = \Sigma^{\text{SM}} \times K(\{c_n\}), \quad (2.13)$$

where the first factor does not depend on the Wilson coefficients and it is given by the same FK-tables that one computes for the standard NNPDF4.0 fits [2], which implicitly assume the SM to be valid at all scales at which observables are measured. The possibility of factoring the whole dependence upon the SMEFT parameters in a multiplicative factor simplifies the procedure highlighted in eq. (2.6) and it thus simplifies the way in which the dependence upon the parameters  $\{c_n\}$  is fitted within SIMUnet, as it will be outlined in section 3.2.

### 3 Methodology

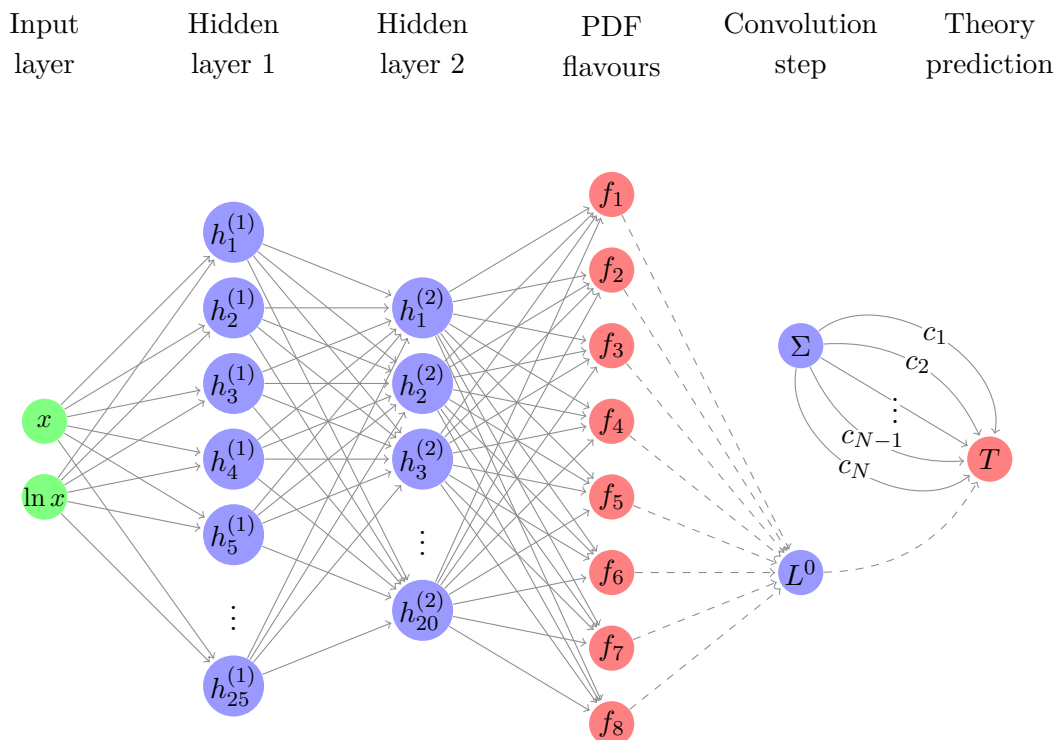
In this section we discuss the details of the new SIMUnet methodology which allows us to extend the deep-learning model of the n3fit approach [2, 21, 31] to optimize a set of parameters of the theory and determine their best fit values to the Monte Carlo replica representation of the experimental data. We also highlight various points of the n3fit methodology that remain pertinent to our study, in particular the hyperparameter selection as well as the cross-validation techniques employed to avoid overfitting.

#### 3.1 Neural network design

The recent NNPDF4.0 fit [2] that our methodology is built upon, shares the features of the previous NNPDF releases [23, 32], specifically the use of a Monte Carlo representation of PDF uncertainties and correlations, and the use of neural networks as basic interpolating functions. However all the details of the fitting methodology, such as the choice of neural network architecture and the minimization algorithm, are now selected through an automated hyperoptimization procedure [31]. As such, our methodology employs state-of-the-art deep-learning techniques through publicly available and highly optimised Machine Learning libraries such as TensorFlow [33] and Keras [34]. As a result, it boasts both performance and improved fit quality using the cutting edge in optimiser technology. Moreover, through the use of TensorFlow graph based execution, the code enjoys the readability Python is famed for, whilst still maintaining the performance of more traditional, statically typed, compiled languages. Additionally, we still retain parallel based execution capabilities, allowing for the ability to fit many replicas in a scalable way, whether on a local machine (using central or graphical processing units) or on a cluster.

The key feature of the SIMUnet methodology is the use of a custom *combination layer*, which captures the dependence of the theoretical predictions upon the external parameters  $\{c_n\}$ , with  $n = 1, \dots, N$ , that we fit alongside the PDFs. The edges of the combination layer are fitted simultaneously with the weights and biases associated with the parametrization of the PDFs at the initial scale  $Q_0$ .

The approach is represented schematically in figure 1, whereby the theory prediction,  $T$ , for each experimental observable included in the fit depends on a dynamical choice of



**Figure 1.** Schematic depiction of the SIMUnet methodology. The input nodes (shown in green) are Bjorken- $x$  and its logarithm. The forward pass through the deep hidden layers (blue) are performed as in the NNPDF4.0 methodology [2] to yield the output PDFs at the initial scale (red). The initial scale PDFs are then combined in the initial scale luminosity  $L^0$ , defined in eq. (2.3). The initial scale luminosity is then convoluted with the pre-computed FK-tables  $\Sigma$  (shown in blue) to obtain the theoretical prediction  $T$  (shown in red), which enters the figure of merit (3.1), which is minimised in the fit. The  $\Sigma$  dependence on the parameters  $\{c_n\}$  is fed into theoretical prediction  $T$  via the trainable edges of the combination layer. All trainable edges are shown by solid edges and are thus learned parameters determined through gradient descent, while dashed edges are non-trainable.

$\{c_n\}$ . The values of  $\{c_n\}$  are associated with the weights of the trainable edges which determine the FK table,  $\Sigma$ , as in eq. (2.6). Such dependence enters the theoretical prediction  $T$  via the bilinear produce between  $\Sigma(\{c_n\})$  and the initial scale PDFs, which in eq. (2.3) we refer to as  $L^0$ , where  $L^0$  indicates either the parametrization of one independent PDF at the initial scale or the product of two of them.

Letting  $\theta$  denote the set of trainable neural network parameters (the weights and biases) that parameterize the PDFs and  $\{c_n\}$  the parameters that we fit alongside the PDFs, SIMUnet fits the joint  $\hat{\theta} = \theta \cup \{c_n\}$  parameter set, by letting gradient descent determine their optimum value in order to minimize the figure of merit used in the fit, which is defined as

$$\chi^2(\hat{\theta}) = \frac{1}{N_{\text{dat}}} (\mathbf{D} - \mathbf{T}(\hat{\theta}))^T (\mathbf{cov})^{-1} (\mathbf{D} - \mathbf{T}(\hat{\theta})), \quad (3.1)$$

with  $\mathbf{D}$  being the vector of experimental central values,  $\mathbf{T}$  the vector of theoretical predictions and  $\mathbf{cov}$  the covariance matrix encapsulating the experimental uncertainties and the correlations therein.

Note that the covariance matrix that we use here is the one obtained by using the so-called  $t_0$  prescription [35], where multiplicative uncertainties are multiplied by theoretical predictions to avoid the D’Agostini bias [36]. If we wanted to include also correlated sources of theoretical uncertainties, such as those associated with missing higher order uncertainties in the theory predictions, we could include them using the method outlined in refs. [37–39]. We leave this endeavour to a future analysis, once the theory covariance matrix for missing higher orders will be available at NNLO.

### 3.2 Parameter fitting using linearisation

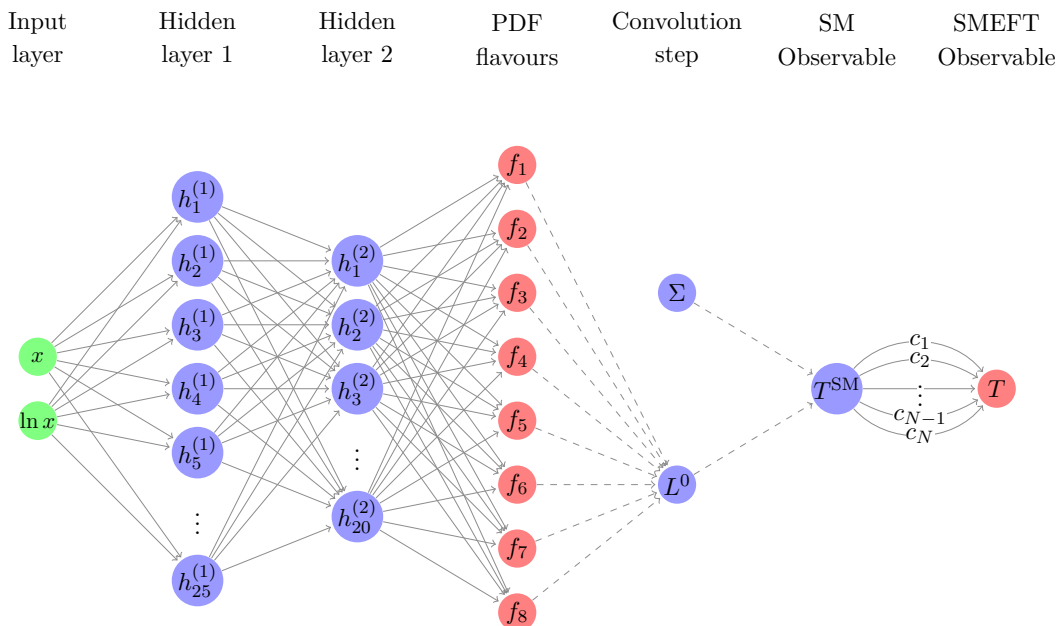
In the case of dim-6 operators, discussed in section 2.2, including only the interference of the SMEFT corrections with the SM diagrams is trivial, as we only add a linear dependence upon the Wilson coefficients. Indeed, the identity of eq. (2.12) allows us to write the theoretical predictions by `SIMUnet` at a particular configuration  $\hat{\theta}$  as

$$T(\hat{\theta}) = \Sigma(\{c_n\}) \cdot L^0(\theta) = T^{\text{SM}}(\theta) \cdot \left( 1 + \sum_{n=1}^N c_n R_{\text{SMEFT}}^{(n)} \right), \quad (3.2)$$

where  $T^{\text{SM}}(\theta) = \Sigma^{\text{SM}} \cdot L^0(\theta)$  is the SM theoretical prediction for each observable, that corresponds to  $\{c_n = 0\}$ .

Given that in this case the dependence upon the parameters  $\{c_n\}$  is factored out of the FK-tables into a multiplicative factor, we can visualize the dependence upon the parameter in the simplified schematic representation given in figure 2. Thanks to linearisation, the bracketed term is implemented using a `Keras` custom layer, which takes the usual SM observable predicted by the network at some interim configuration and maps it to a SMEFT modified observable with the strength of the new physics interaction being determined by the weights of the combination layer, which in this case is simply an extra sequential layer that maps  $T^{\text{SM}}$  into the theoretical prediction  $T$  which enters the figure of merit defined in eq. (3.1).

The SMEFT-modified theory prediction,  $T$ , is then split into disjoint training and validation splits. For each split we compute a training and validation  $\chi^2$ . Gradient descent attempts to minimize the training  $\chi^2$  by descending the loss surface in  $\hat{\theta}$ -space while the validation  $\chi^2$  is monitored to assess the network’s out of sample performance. The validation  $\chi^2$  decreases initially (since the network is learning the shared underlying laws that are common to both sets), but upon the onset of overfitting, the out of sample performance begins to deteriorate as the network fits the noise of the training data. This particular point in the training process corresponds to the validation  $\chi^2$  ceasing to decrease and instead beginning to increase. Upon reaching this regime, training is halted and various checks, such as positivity and integrability of the resulting PDF are assessed. It is at this point that the best fit values of the Wilson coefficients  $\{c_n\}$  are obtained, since gradient descent has modified the combination layer’s weights such that they best fit the input data. Using this cross-validation procedure [40–42] we ensure the fitting of statistical fluctuations are avoided as much as possible and it is the shared underlying physical laws that are being fitted instead. At the same time, in order to ensure that the  $\{c_n\}$  are not overfitted, we



**Figure 2.** Schematic representation of the architecture used by SIMUnet in the case of the fit of SMEFT coefficients, in which the dependence of the theoretical prediction  $T$  upon the parameters  $\{c_n\}$  can be factored into a multiplicative K-factor, as in eq. (2.13). The scheme is the same as the one of figure 1, however the initial scale PDFs  $L^0$  are first convoluted with the relevant SM FK tables to obtain the Standard Model theory prediction ( $T^{\text{SM}}$ ). The SM predictions are then incremented by the addition of the linear SMEFT corrections via a final linear *combination layer*. All solid edges are trainable and thus modified during gradient descent. The full structure of the hidden layers are suppressed for clarity. The precise nature of the manipulation performed by the final layer is outlined in eq. (3.2).

set the datasets that will be modified by these parameters to have a good representation both within the training set and validation set. Thus in this study we split such datasets to have their training and validation fractions equal:  $f_{\text{tr}} = f_{\text{val}} = 0.5$ , while keeping all other training and validations fractions as in the fits of ref. [17].

As with all deep learning studies, the user has freedom to choose the various hyperparameters of their model at their own discretion. Such parameters include the architecture of the network, the particular choice of initializer that sets the initial values of network parameters  $\{\theta\}$ , or the choice of minimizer (and the specific settings therein) that tunes these parameters to their optimal values such that the performance metric is minimized. Techniques exist to automate this process such as the `hyperopt` library [43] employing Bayesian methods to perform this optimization. Indeed this is employed by NNPDF4.0 and we choose to use the same settings that were found to be optimal there [2, 21]. In this study, the hyperoptimization procedure has not been performed to reassess the optimality of the various neural network settings. Indeed, this is justified, since we expect, *a priori* that the PDF modifications due the presence of the aforementioned SMEFT operators will be moderate, and thus the hyperparameters selected assuming a SM-only scenario will remain adequate. However, this assumption does not hold anymore in the event that a

more marked PDF modification is expected or obtained *a posteriori*. This point is also worth considering if one is to introduce a large number (relative to NNPDF4.0) of new measurements or any measurements which introduce tensions with other datasets. Should this be the case, hyperoptimization should be performed again on the layers preceding the combination layer, and this can be done straightforwardly using the standard procedure described in ref. [2].

The connections to the combination layer (one for each of the  $c_n$ ) are initialized to zero, since we assume *a priori* the Wilson coefficients will be small; although we observe that if they are initialized according to a normal distribution, then virtually identical results are obtained. These connections are then modified during back propagation. It is worth mentioning that for small values of Wilson coefficients (such as those in this study), it is highly desirable to scale the units such that the Wilson coefficients are  $\mathcal{O}(1)$  and thus comparable to the learning rate. This will assist the optimiser during gradient descent to converge upon the minimum in a timely fashion, since the step size will be more suited to the characteristic scale of the Wilson coefficients. In practice, the appropriate scaling is usually determined *a posteriori*, where one can analyse the typical values for the Wilson coefficients and refit with the normalization set accordingly.


### 3.3 Incorporating non-linear effects

The effect of including the SMEFT self-interference diagrams can often introduce a marked effect on the bounds obtained for a given SMEFT scenario [17, 44]. Moreover, with the impetus to produce high precision theoretical predictions for the LHC era, the inclusion of higher order corrections in the SMEFT scattering graphs are becoming particularly pertinent [45–48]. Such considerations introduce a non-linear dependence on the Wilson coefficients in the space of observables: quadratic in the former and quadratic in the highest order in the QCD expansion in the latter. This point serves as a major advantage of our methodology which can accommodate these effects during the fit by the simple addition of non-trainable edges.

When computing the amplitude of a Feynman diagram related to some process one has the schematic form  $\mathcal{A} = \mathcal{A}^{\text{SM}} + \sum \mathcal{A}_i$  where  $\mathcal{A}_i$  is the amplitude corresponding to the operator  $\mathcal{O}^{(i)}$  computed to some order in perturbation theory. We assume here that it is LO in the Wilson coefficients, but need not be in general and the extension to higher orders in the Wilson coefficient expansion is discussed at the end of the present section. When computing the observable, the matrix element,  $|\mathcal{A}|^2$ , introduces terms of the form  $\mathcal{A}_i \mathcal{A}_j$ . Since these amplitudes are computed to LO in perturbation theory, we can rewrite eq. (3.2) as

$$T(\hat{\theta}) = T^{\text{SM}}(\theta) \cdot \left( 1 + \sum_{n=1}^N c_n R_{\text{SMEFT}}^{(n)} + \sum_{1 \leq n < m \leq N} c_n c_m R_{\text{SMEFT}}^{(n,m)} \right) \quad (3.3)$$

with  $R_{\text{SMEFT}}^{(n,m)}$  being defined as in eq. (2.10). Including this contribution in the simultaneous fit is a straightforward task and simply requires that the manipulation performed by the combination layer correspond to that of eq. (3.3) instead of eq. (3.2).



**Figure 3.** Schematic representation of how SIMUnet allows for the effects of the SMEFT self-interaction diagrams ( $\text{dim-6}$ )<sup>2</sup> to be included. We show here how the SM observable can be transformed to a SMEFT observable which includes these  $\mathcal{O}(1/v^4)$  terms. The preceding PDF layers are omitted for clarity. The linear contributions are included in the usual way, with the strength of the SMEFT couplings being determined by the trainable edges,  $c_n$ : shown by the solid lines. The SMEFT-SMEFT interference contributions are instead non-trainable edges with their value being fixed by the strength of the corresponding pair of trainable edges. These are shown by the dashed lines. There will in general be  $N$  trainable edges, and  $N(N + 1)/2$  non-trainable edges.

In practice, however, it is more convenient to rewrite the terms in the right most sum of eq. (3.3) by defining  $c_{nm} = c_n c_m$ :

$$T(\hat{\theta}) = T^{\text{SM}}(\theta) \cdot \left( 1 + \sum_{n=1}^N c_n R_{\text{SMEFT}}^{(n)} + \sum_{1 \leq n \leq m \leq N} c_{nm} R_{\text{SMEFT}}^{(n,m)} \right), \quad (3.4)$$

we see that both summations are of the same form and so the manipulation required to incorporate the quadratic effects of the squared SMEFT amplitude is reduced to introducing  $N(N + 1)/2$  additional edges to the combination layer, one for each  $c_{nm}$ . These additional connections, however, are not trainable, that is to say their value is not determined by gradient descent during learning; since  $c_{nm}$  is completely determined by the product of trainable edges  $n$  and  $m$ . This is shown schematically in figure 3, whereby the trainable edges determine the non-trainable edges that perform the manipulation corresponding to the right most term in the brackets of eq. (3.3). Since the observable is always polynomial in the Wilson coefficients, it is possible to include the effects of higher dimensional operators in the SMEFT expansion (such as adding dimension-8 operators [49, 50]) in this way.

Moreover, to incorporate quantum corrections arising from Next-to-Leading Order (NLO) terms in the QCD perturbative expansion of the SMEFT corrections is a similarly straightforward task, simply requiring the computation of the corresponding K-factor and including an additional edge in the combination layer whose value is pinned to be the value of the corresponding trainable edge raised to some power. If one wanted to include also the scale dependence of the Wilson coefficients [26–28], one would have to fit the scale dependent Wilson coefficients at some fixed scale,  $Q_0$ , and use the relevant  $\beta$ -functions to evolve the operator to the relevant scale using some pre-computed evolution kernel, which would be factored into the pre-computed FK-tables.

Finally, it is well-known that one can obtain prior knowledge on the Wilson coefficients by assuming a UV completion of the EFT exists that is local, unitary and causal. Such matching conditions can often impose positivity bounds [51] on functions,  $f_i$ , of the Wilson coefficients by leveraging these standard conditions on the UV theory. To incorporate such a prior within our framework is analogous to the way positivity on physical cross sections is achieved within the NNPDF methodology [22, 52]. The minimizer is informed of the prior by modifying the training loss function to penalize negative values of Wilson coefficients, with one such choice of penalty term being

$$\text{loss} = \text{loss} + \lambda \sum_i \Theta(-f_i(\{c\})), \quad (3.5)$$

with  $\Theta$  the usual Heaviside step function which takes unit value for positive arguments and vanishes otherwise. The parameter  $\lambda$  is a non-negative scalar which can be treated as a hyperparameter and encapsulates the degree of belief in the prior: larger values impose the positivity more strictly while smaller values allow for slight violations of the positivity constraints. In a similar way to NNPDF it is also possible to filter the replicas *a posteriori* by discarding those that are deemed to violate positivity constraints too strongly.

### 3.4 Fixed PDF analysis

A desirable feature of any simultaneous fitting methodology is to be able to benchmark the simultaneous extraction with the analogous fixed PDF analysis. Indeed, this is precisely what is done in ref. [17] and the broadening of the Wilson coefficients bounds there being the key message of the paper.

This too is achievable with our methodology and proceeds as follows. We begin by first freezing the combination layer parameters by making them non-trainable. Mathematically, the combination layer thus performs an identity transformation, effectively removing itself from the process. At this stage we are in effect performing a regular NNPDF4.0 PDF fit. Upon successfully achieving the stopping criterion for the PDF only portion of the fit, we then freeze the PDF sector of the neural network (again by making the weights and biases non-trainable) and reinstate the combination layer, allowing for gradient descent to optimize the external parameters,  $c_i$ , only. In this way we eliminate the cross-talk between the PDFs and the external parameters and so the external parameters fitted in this way are equivalent to an analogous study keeping the PDFs fixed to the baseline.

## 4 Results

In this section we present the results we obtain by applying the SIMUnet methodology to fit the Wilson coefficients alongside the PDFs in the two benchmark scenarios. We begin by discussing the specific scenarios explored in this work by motivating our particular selection of operators. We then outline our dataset selection, which is chosen to provide strong constraints not only on the PDF, but also to the Wilson coefficients under consideration in each of the two scenarios. We then turn to describe the results we obtain in the two scenarios, both in terms of the resulting PDFs and bounds on the Wilson coefficients. We

compare our findings to those obtained when PDFs are kept fixed to the SM baseline, rather than fitted simultaneously alongside the Wilson coefficients. Finally, we summarize our results, by exploring the correlations between PDFs and the Wilson coefficients that we consider in this study and quantifying the agreement with respect to the previous findings of ref. [17].

### 4.1 Benchmark scenarios

The two SMEFT scenarios we explore in this work are the same as those introduced in ref. [17]. In this way, the SIMUnet methodology can be benchmarked against the rather less efficient methodology based on Hessian minimization and on fitting the PDFs at fixed values of the Wilson coefficients that we used in our previous work. Also, the two scenarios highlight the interplay between the PDFs and the EFT dynamics, illustrating in particular how the former changes and how constraints to the latter are modified. Moreover, they allow us to test the inclusion of purely linear SMEFT corrections (as in Benchmark Scenario I) and of both linear and quadratic SMEFT corrections (as in Benchmark Scenario II).

In particular, the first Benchmark Scenario belongs to the class of electroweak precision tests and is sensitive to a broad range of UV-complete theories proposed in the literature. The oblique corrections, originally proposed in [53, 54], play a key role in testing BSM theories. They parametrise the self-energy of the electroweak gauge bosons. Out of the four operators which can be identified with dim-6 operators in the SMEFT, two, namely  $T$  and  $S$ , are well constrained by precision LEP measurements [55] and grow slowly with the energy, while  $W$  and  $Y$  scale as  $\mathcal{O}(q^4)$  and are particularly sensitive to the tails of the high mass Drell-Yan distributions at the LHC [56, 57]. Their inclusions implies the addition of the following terms to the SM Lagrangian

$$\text{Benchmark Scenario I:} \quad \mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} - \frac{W}{4m_W^2} (D_\rho W_{\mu\nu}^a)^2 - \frac{Y}{4m_W^2} (\partial_\rho B_{\mu\nu})^2 \quad (4.1)$$

with  $m_W$  the  $W$  boson mass and  $D$  the usual covariant derivative. In this scenario the quadratic (dim-6)<sup>2</sup> effects are much less relevant than the linear effects coming from the interference between the SM and the SMEFT amplitudes, thus we can define the SMEFT-modified theoretical predictions as in eq. (3.2).

The second benchmark represents a consistency check of the existing hints of lepton universality violation in rare  $B$ -meson decays reported by the LHCb collaboration [58–60]. In particular, we look at the muon-philic operator, in which the Lagrangian is defined as:

$$\text{Benchmark Scenario II:} \quad \mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \frac{\mathbf{C}_{33}^{D\mu}}{v^2} (d_L^3 \gamma_\mu d_L^3)(\bar{\mu}_L \gamma^\mu \mu_L), \quad (4.2)$$

where  $v$  is the Higgs vacuum expectation value,  $\mathbf{C}_{33}^{D\mu}$  is a scalar Wilson coefficient for this scenario, and  $d^3(\mu)$  is the bottom quark (muon) Dirac spinor. Noting that the inclusion of this operator modifies only the muon channel whereas the electron channel is still described by the SM: when studying this particular SMEFT scenario, we only modify the muon channel in the high-mass DY datasets. Moreover, in this scenario the quadratic corrections are important, so we do include the presence of quadratic Wilson coefficients in the matrix



element to constrain  $C_{33}^{D\mu}$  and define the SMEFT-modified theoretical predictions according to eq. (3.3).

## 4.2 Experimental data


The dataset chosen for this study comprises 4022 data points, spanning a broad range of processes including neutral and charged current DIS data (the vast majority of which is composed of the HERA combined dataset) to high mass Drell-Yan measurements from the LHC. The dataset selection is identical to that of ref. [17] which is in itself an extension of the strangeness study of [61] and NNPDF3.1 [32]. In particular, the neutral current Drell-Yan measurements from ATLAS at 7, and 8 TeV [62, 63] and CMS at 7, 8, and 13 TeV [64–66] are LHC measurements that are used to constrain, not only the PDF, but are also sufficiently sensitive to the BSM scenario considered in this work to also constrain the SMEFT operators. The kinematic coverage of the data points used in this study are shown in figure 4. The points are shown in  $(x, Q^2)$  space with the data points that are modified by the EFT operators highlighted with a border, such points thus also constrain the Wilson coefficients as well as the PDFs. We note that, although DIS theory predictions are modified by the operators we consider in the two benchmark scenarios, the change in the HERA DIS cross sections upon the variation of the Wilson coefficients upon consideration is minimal, as is explicitly assessed in ref. [17].

## 4.3 Results for Benchmark Scenario I


We first present the results obtained within the first Benchmark Scenario outlined in section 4.1, in which the linear effects are dominant as compared to the quadratic effect and, as a result, SMEFT-modified theoretical predictions are defined as in eq. (3.2). We start by employing `SIMUnet` to individually constrain the  $W$  and  $Y$  operators with the ATLAS and CMS neutral-current (NC) high mass DY data from Run I and Run II. In the next subsection we will be able to constrain them simultaneously thanks to the inclusion of charged-current (CC) HL-LHC projections.

All details about the fit quality, once PDFs are allowed to vary alongside the  $W$  or the  $Y$  parameters are given in appendix B. We observe that the added degrees of freedom result in a slightly better fit quality when compared to the NNPDF4.0-like SM baseline fit based on the reduced dataset described in section 4.2 and on SM theoretical predictions. In particular the high-mass DY datasets that are sensitive to SMEFT corrections display a more significant improvement in fit quality, largely driven by the CMS measurements at 7 TeV, owing to the large weight carried by this dataset, which forms just under half the entire high-mass DY data points hereby considered.

In figures 5 we show the distribution of the optimal values of  $W$  and  $Y$  determined for each of the 1000 replicas of the Monte Carlo representation of the experimental data that we obtain during gradient descent in the simultaneous fit of either  $W$  or  $Y$  and the PDFs. The best-fit values are normally distributed in keeping with the Monte-Carlo pseudodata replica generation that the neural network replicas are fit to. We compare the distribution (in green) with the one that we obtain by keeping the PDFs fixed to the SM baseline (in orange), the latter obtained by using the methodology outlined in section 3.4.



**Figure 4.** The kinematic coverage of the data points used in this study grouped according to their process. Points marked with a black edge are also used to constrain the SMEFT operators. Note that the HL-LHC projections used in this study are omitted from this figure, as only actual data are displayed.



**Figure 5.** Distribution of best fit  $W$  (left panel) and  $Y$  values (right panel) obtained across 1000 replicas by fitting each of the Wilson coefficients alongside PDFs with the **SIMUnet** methodology (green) compared to the distribution that one would obtain by keeping the PDFs fixed to the SM baseline (orange).

	SM PDFs	SMEFT PDFs	best-fit shift	broadening
$W \times 10^4$ (68% C.L.)	[-9.0, 2.1]	[-8.2, 3.2]	+1.0	+2.6%
$W \times 10^4$ (95% C.L.)	[-14.5, 7.6]	[-13.9, 8.9]	+1.0	+3.2%
$Y \times 10^4$ (68% C.L.)	[-13.8, 5.7]	[-12.2, 7.9]	+1.9	+3.1%
$Y \times 10^4$ (95% C.L.)	[-23.5, 15.5]	[-22.2, 18.0]	+1.9	+3.1%

**Table 1.** The 68% and 95% C.L. bounds on the  $W$  and  $Y$  parameters obtained either for a fit in which PDFs are kept fixed (SM PDFs) or in a fit in which PDFs are fitted simultaneously with either  $W$  or  $Y$  (SMEFT PDFs). The fourth and fifth column indicate the absolute shift in best-fit values, eq. (4.3), and the percentage broadening of the SMEFT bounds, eq. (4.4), when the PDFs are allowed to change alongside the Wilson coefficients.

Both distributions are centred around zero, illustrating that the high-mass DY datasets are compatible with the SM-only hypothesis, but admit non-zero values, with  $Y$  being less constrained by the data than  $W$ . The distribution of best fits obtained in a simultaneous fit looks similar to the distribution of best fits obtained by keeping the PDFs fixed to the SM baseline.


For a more quantitative comparison, in table 1 we compare the bounds for the individual  $W$  and  $Y$  fits obtained in the simultaneous fits to those obtained by keeping the PDFs fixed to the SM baseline. We indicate the shift and the broadening of the bounds that are obtained once the Wilson coefficients are fitted alongside PDFs, by defining them as

$$\text{best fit shift} \equiv \left( \langle W \rangle \Big|_{\text{SMEFT PDFs}} - \langle W \rangle \Big|_{\text{SM PDFs}} \right), \quad (4.3)$$

$$\text{broadening} \equiv \left( \Delta_W \Big|_{\text{SMEFT PDFs}} - \Delta_W \Big|_{\text{SM PDFs}} \right) / \Delta_W \Big|_{\text{SM PDFs}}, \quad (4.4)$$

where  $\Delta$  is the size of either the 68% C.L. or the 95% C.L., depending on the bounds we consider. The same definitions apply for the  $Y$  parameter. As in [17], the effect of fitting the Wilson coefficients simultaneously with the effects that non-zero coefficients have on PDFs changes the interpretation of high-mass DY constraints in a moderate way, with the bounds of the simultaneous fit being within the uncertainty of the current analysis and only slightly looser (by a factor around 3%) than those obtained by keeping the PDFs fixed.

In figures 6 we show the quark-antiquark channel luminosity plots defined in eq. (2.9) with the error bands showing 68% C.L., normalised to the baseline SM PDFs. We notice that, while in the previous analysis of ref. [17], we could only produce sets of PDFs obtained at fixed values of  $W$  and  $Y$  (corresponding to the Benchmark Points that were taken under consideration), here we do really produce a set of PDFs obtained out of a simultaneous fit alongside the Wilson coefficients. We see the luminosity modification due to the simultaneous fit remains moderate, with a slightly larger deviation found at higher values of the produced lepton invariant mass, where the PDFs exhibit slightly larger uncertainties. This is indeed a result consistent with the indications given in ref. [17]. Indeed, the dominant luminosity channel for NC DY is  $u\bar{u}$  and  $d\bar{d}$  with the valence quark distribution being strongly constrained by DIS and the SMEFT modification being small relative to the ex-




**Figure 6.** The one dimensional luminosity in the  $q\bar{q}$  channel for a PDF fitted in the presence of a non-zero  $W$  operator (left panel) or of a non-zero  $Y$  operator (right panel) shown in orange, normalized to the SM baseline PDF shown in green.

perimental uncertainties of the highest invariant mass bins probed by current experimental data. Again, this finding shows that the interplay between EFT effects and PDFs remains moderate when one performs a truly simultaneous determination of both.

It is interesting to observe that, if we try to fit  $W$  and  $Y$  at the same time using only the current NC DY data, we identify both the flat direction and the strong anti-correlation between  $W$  and  $Y$ , which are known to exist in this case [56]. Results are shown in figure 7. Both the flat direction and the anti-correlation can be retrieved within the framework of the SIMUnet methodology, without the user having to be aware of the existence or particular nature of any flat directions which may exist. Indeed, the optimizer cannot preferentially differentiate one point from another within this landscape valley and so the points are distributed tightly along the flat direction. The preference for the upper left quadrant in the  $W$ - $Y$  plane is consistent with the findings of ref. [56] with the failure to capture the origin possibly due to the lack of inclusion of  $\mathcal{O}(1/\Lambda^4)$  terms in the partonic cross section which is known to ease such tensions when only  $\mathcal{O}(1/\Lambda^2)$  terms are present, as it was pointed out in ref. [44], where the effect of dim-8 operator is carefully assessed.

#### 4.4 Inclusion of the HL-LHC projections


The flat direction illustrated in figure 7 can be eliminated with the inclusion of Charged Current (CC) DY data [17]. No unfolded measurements of the high-mass transverse mass  $m_T$  distribution have been yet released at 13 TeV, thus we will base our analysis on the High Luminosity LHC (HL-LHC) high-mass Drell-Yan projections that we produced in ref. [17], inspired by the HL-LHC projections studied in ref. [67]. The invariant mass distribution projections are generated at  $\sqrt{s} = 14$  TeV, assuming an integrated luminosity of  $\mathcal{L} = 6 \text{ ab}^{-1}$  ( $3 \text{ ab}^{-1}$  collected by ATLAS and  $3 \text{ ab}^{-1}$  by CMS). Both in the case of NC and CC Drell-Yan cross sections, the pseudodata were generated using the MadGraph5\_aMCatNLO NLO Monte Carlo event generator [68] with additional  $K$ -factors to include the NNLO QCD and NLO EW corrections. The pseudodata consist of four datasets (associated with NC/CC distributions with muons/electrons in the final state), each comprising 16 bins in the  $m_{ll}$



**Figure 7.** Scatter plot for the best fit values of  $(W, Y)$  per replica using ATLAS and CMS high mass Drell-Yan data, which are exclusively NC observables. The upper and rightmost panels are histograms in their respective directions. A clear flat direction has been detected along with a strong anti-correlation.

invariant mass distribution or transverse mass  $m_T$  distributions with both  $m_{ll}$  and  $m_T$  greater than 500 GeV, with the highest energy bins reaching  $m_{ll} = 4$  TeV ( $m_T = 3.5$  TeV) for NC (CC) data. The rationale behind the choice of number of bins and of the width of each bin was outlined in ref. [17], and stemmed from the requirement that the expected number of events per bin was big enough to ensure the applicability of Gaussian statistics. The choice of binning for the  $m_{ll}$  ( $m_T$ ) distribution at the HL-LHC is displayed in figure 5.1 of ref. [17].

We performed a simultaneous fit of the PDFs and the two  $(W, Y)$  parameters by adding two trainable edges in the combination layer displayed in figure 2, and appending the aforementioned HL-LHC projected data to the data discussed in section 4.2. The best fit values of  $(W, Y)$  obtained for each of the 1000 replicas of the Monte Carlo ensemble are plotted in figure 8. We see that not only is the flat direction of figure 7 broken, but the HL-LHC projections heavily favour vanishing Wilson coefficients. Indeed, the origin is now covered in both the  $W$  and  $Y$  axes and much more heavily constrained, enjoying roughly two orders of magnitude tighter constraints in both directions. We notice a slightly favorable pull towards the upper-left quadrant in the  $W$ - $Y$  plane, that the ATLAS and CMS datasets seem to prefer. Comparing the best-fit distribution that we get in a simultaneous fit (displayed in green) to the one we get in a fit in which PDFs are kept fixed to the SM baseline (displayed in orange) we can see that the bounds are visibly tighter once



**Figure 8.** Scatter plot for best fit tuples of  $(W, Y)$  for each replica obtained in the simultaneous fit (green) compared to those obtained when PDFs are kept fixed to the SM baseline (orange). The upper and rightmost panels are histograms in their respective directions.

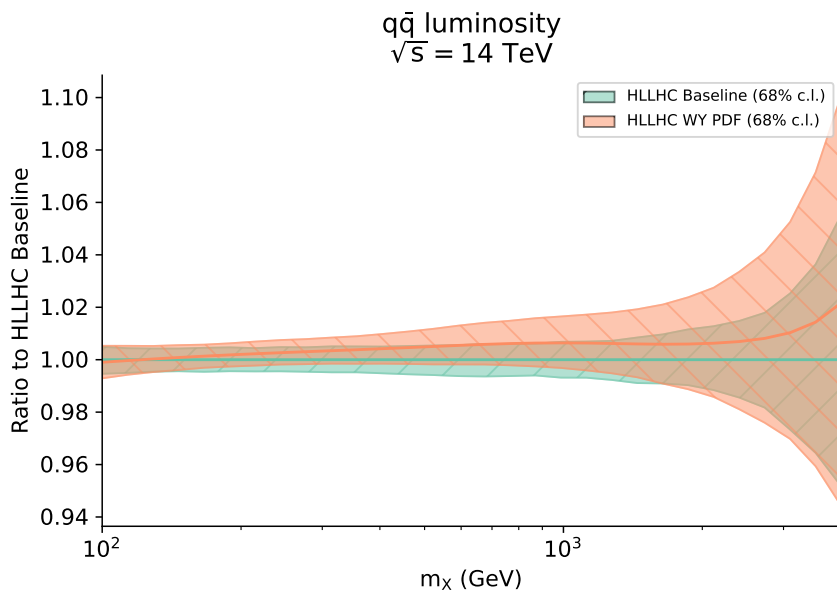
	SM PDFs	SMEFT PDFs	best-fit shift	broadening
$W \times 10^5$ (68% C.L.)	$[-1.1, 0.5]$	$[-2.4, 1.5]$	-0.2	+144%
$W \times 10^5$ (95% C.L.)	$[-2.0, 1.4]$	$[-4.3, 3.4]$	-0.2	+126%
$Y \times 10^5$ (68% C.L.)	$[-0.4, 5.2]$	$[0.6, 8.0]$	+1.9	+32%
$Y \times 10^5$ (95% C.L.)	$[-3.2, 8.1]$	$[-3.1, 11.7]$	+1.9	+31%

**Table 2.** Same as table 1 for the 68% C.L. and 95% C.L. marginalised bounds on the  $W$  and  $Y$  parameters obtained from the two-dimensional  $(W, Y)$  fits that include the HL-LHC pseudo-data for NC and CC Drell-Yan distributions.

PDFs are kept fixed to the SM baseline and are not allowed to consistently vary alongside the Wilson coefficients.

In table 2 we compare the bounds for the individual  $W$  and  $Y$  fits obtained in the simultaneous fits to those obtained by keeping the PDFs fixed to the SM baseline. We indicate the shift and the broadening of the bounds according to the definitions given in eqs. (4.3)–(4.4). Consistently to what was found in our previous study of ref. [17] we observe that including high-mass data at the LHC both in a fit of PDFs and in a fit of SMEFT coefficients and neglecting the interplay between them could result in a significant underestimate of the uncertainties associated to the SMEFT parameters. Indeed, the marginalised bounds on the  $W(Y)$  parameter increase by about 150% (30%) once a simultaneous fit of the PDFs and the  $(W, Y)$  parameters is performed. The broadening is smaller than observed in ref. [17] but it is still extremely significant. A detailed comparison is given in section 4.6.

As far as the quark-antiquark luminosity is concerned, we can see in figure 9 that once the PDFs are fitted simultaneously with the  $(W, Y)$  parameters two things happen. First

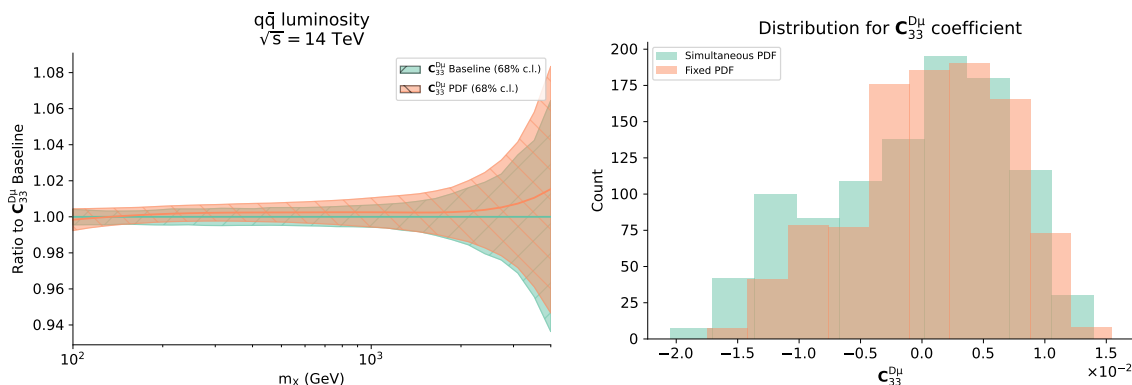


**Figure 9.**  $q\bar{q}$  luminosity channel of the PDF fitted in the presence of  $W$  and  $Y$  parameters fitted to the ATLAS and CMS high mass Drell-Yan data and the NC and CC DY HL-LHC projections, normalized to the appropriate baseline SM PDF.

of all the central values of the luminosity shift upwards for large values of the invariant mass ( $M_X \gtrsim 1$  TeV) towards the edge of the 68% C.L. error band. Second the error band significantly increases. The shift in the central value is compatible to what we observed before, namely that the luminosity plots, once PDFs are fitted at some representative values of the  $W$  and  $Y$  parameters, do change significantly, well outside the  $1\sigma$  error band of the SM PDFs, while the PDF uncertainties themselves are unchanged. However, in this case the PDF uncertainty does increase because here we are actually performing a simultaneous fit and, as a result, the PDF error band increases proportionally to the width of the range of  $W$  and  $Y$  that the data allow. This is a very interesting result and shows that PDF error bands at large- $x$  inherently have an extra source of theory uncertainty related with possible BSM effects that the data do not exclude.

#### 4.5 Results for Benchmark Scenario II

In this section we employ the left-handed muon-philic operator  $C_{33}^{D\mu}$  to showcase our methodology’s ability to constrain Wilson coefficients whilst accounting for the effect of quadratic dim-6 effects using the approach discussed in section 3.3. The fact that this second Benchmark SMEFT scenario is effectively unconstrained [17] at the linear level serves to act as the ideal setting to assess the ability to fit EFT operators whilst simultaneously accounting for their quadratic contributions, as in eq. (3.3). Furthermore, since the  $C_{33}^{D\mu}$  operator affects only muon final state observables, while electron final states are described by the SM, the combination layer uses only those datasets that have a muon in the final state to constrain  $C_{33}^{D\mu}$ . In particular, the CMS high-mass measurements at 13 TeV, which up until now have been for the combined decay channel, is now separated into the electron



**Figure 10.** Left:  $q\bar{q}$  luminosity channel of the PDF fitted in the presence of the  $C_{33}^{D\mu}$  parameter normalized to the baseline SM PDF. Right: histogram plot for best fit values for  $C_{33}^{D\mu}$  for each replica. The best-fit distribution over 1000 replicas obtained out of a simultaneous fit (green) is compared to the one obtained by keeping PDFs fixed to the SM baseline (orange).

	SM PDFs	SMEFT PDFs	best-fit shift	broadening
$C_{33}^{D\mu} \times 10^3$ (68% C.L.)	[-5.6, 6.9]	[-8.0, 6.7]	-0.9	+18%
$C_{33}^{D\mu} \times 10^3$ (95% C.L.)	[-11.9, 13.1]	[-15.3, 14.0]	-0.9	+17%

**Table 3.** Same as table 1 for the 68% C.L. and 95% C.L. bounds on the  $C_{33}^{D\mu}$  Wilson Coefficient, including both linear and quadratic terms in the SMEFT expansion.

and muon channels. As a result of splitting this particular dataset into separate channels, we have accordingly generated a new baseline PDF used in the comparisons.

The best-fit values of  $C_{33}^{D\mu}$  across the 1000 replicas in the fit are shown on the right-hand panel of figure 10. We compare the distribution obtained out of a simultaneous fit (in green) with the one obtained once PDFs are kept fixed to the SM baseline (in orange). We see that the distribution of best fit values is centred at the origin, although it is skewed towards  $C_{33}^{D\mu} \approx 5 \cdot 10^{-3}$ . The shape of the distribution is what we expect once quadratic terms are allowed in the fit of the Wilson coefficients by keeping PDF fixed [7], and it is interesting to see this feature not only holds but it is actually enhanced once the Wilson coefficient is fitted alongside PDFs (green histogram).

For a quantitative comparison of the bounds obtained in the simultaneous fit to those obtained in a Wilson coefficient-only fit, in table 3 we show the bounds that we obtain in the two cases. The interplay between PDFs and SMEFT coefficients is quite moderate in this particular scenario. In contrast with the marked effects in Benchmark Scenario I, in Benchmark Scenario II the obtained bounds on this Wilson coefficient would loosen by around 20%. The origin of this rather different behaviour can be traced back to the fact that in this scenario the electron channel data do not receive EFT corrections, and hence all the information that they provide makes it possible to exclusively constrain the PDFs. The muon channel distributions then determine the allowed range for the  $C_{33}^{D\mu}$ , restricted by the well-constrained large- $x$  quarks and antiquark PDFs from the electron data. This



claim is backed by the luminosity comparison, displayed on the left-hand panel of figure 10: the shift and the increase in the PDF uncertainty of the  $q\bar{q}$  luminosity are visible, but less enhanced than in the first Benchmark Scenario.


#### 4.6 Results overview

To conclude this section, we present an overview of our results. We have seen in section 4.3 that the current high-mass Drell-Yan data from LHC Run I and Run II do not allow to simultaneously fit  $W$  and  $Y$  in Benchmark Scenario I and loosely constrain  $\mathbf{C}_{33}^{D\mu}$  in Benchmark Scenario II. This is mostly due to the lack of unfolded CC Drell-Yan data, which would remove the flat direction that our algorithm is able to detect (see figure 7). Moreover, our analysis confirms what was outlined in ref. [17], namely that the interplay between the individual fits of  $W$  and  $Y$  and the fit of the large- $x$  quark distributions is mild at the level of current DY data. However, once the high-statistics data projections from the HL-LHC are included, the flat direction disappears and one is able to obtain strong constraints both on the  $(W, Y)$  plane [17, 56] and on the individual  $\mathbf{C}_{33}^{D\mu}$  coefficient [17, 29]. From the point of view of showcasing our methodology, the two scenarios are interesting, as in the Benchmark Scenario I, once the HL-LHC projections are included, we can simultaneously fit both  $W$  and  $Y$  alongside the PDFs, including only the linear SMEFT corrections while in the Benchmark Scenario II we can fit individually  $\mathbf{C}_{33}^{D\mu}$  alongside the PDFs, including both the linear and the quadratic SMEFT corrections. In the first scenario, we observe that there is a strong interplay between SMEFT and PDF fits, as the bounds for the SMEFT coefficients significantly broaden once PDFs are allowed to vary alongside  $W$  and  $Y$  (see figure 8) and the PDFs themselves display a sizeable shift (see figure 9). The interplay is more moderate in the second scenario, given that BSM effects only affect the data with muons in the final states, while the data with electrons in the final states constrain the large- $x$  quark distributions.

In this section, we focus on the results obtained including the HL-LHC projections. We first explore the correlation patterns between the PDFs and the SMEFT coefficients in both Benchmark Scenarios. These correlation coefficients can be evaluated as in ref. [18]. For example in the case of the gluon and the  $W$  Wilson coefficient, the correlation coefficient is defined as follows

$$\rho(W, g(x, Q)) = \frac{\langle W^{(\text{best-fit})} g(x, Q) \rangle - \langle W^{(\text{best-fit})} \rangle \langle g(x, Q) \rangle}{\sqrt{\langle W^{(\text{best-fit})2} \rangle - \langle W^{(\text{best-fit})} \rangle^2} \sqrt{\langle g(x, Q)^2 \rangle - \langle g(x, Q) \rangle^2}}, \quad (4.5)$$

where  $W^{(\text{best-fit})}$  is the best-fit of the  $W$  coefficient for each replica in the simultaneous fit, in which PDFs are allowed to vary alongside  $W$  and  $Y$ . In eq. (4.5), averages are computed over the  $N_{\text{rep}} = 1000$  replicas. This correlation coefficient provides a measure of how the variations in the PDFs translate into modifications of the best fit value of the Wilson coefficients. In figure 11 we show these correlation coefficients between the up valence quark and the gluon PDFs at  $Q = 100$  GeV. Each of the curves corresponds to one of Wilson coefficients considered in the HL-LHC analysis, the  $W$  and  $Y$  being fitted simultaneously alongside the PDFs (hence the “WY PDF” label) and the  $\mathbf{C}_{33}^{D\mu}$  being fitted



**Figure 11.** Correlation coefficients, defined in eq. (4.5), between the  $W$  (green),  $Y$  (orange),  $C_{33}^{D\mu}$  (blue) and the up quark PDF (left panel) and gluon PDF (right panel) computed at  $Q = 100$  GeV as a function of the momentum fraction  $x$ . The PDFs used to compute the correlation coefficients are those obtained in the simultaneous fit of  $(W, Y)$  and the PDFs described in section 4.4 (in the case of the green and orange curves) and those obtained in the simultaneous fit of  $C_{33}^{D\mu}$  and the PDFs described in section 4.5 (in the case of the blue curve).

individually alongside the PDFs (hence the “ $C_{33}^{D\mu}$  PDF” label). Although correlations are moderate, it is interesting to observe that the correlation/anticorrelation between  $W$  and the PDFs is much stronger than the correlation with the other Wilson coefficients. This explains why the broadening of the bounds in the  $W$  direction are more marked than those in the  $Y$  directions. Moreover the correlation is stronger in the large- $x$  region and for the gluon, which implies an anticorrelation with the singlet in the same region, which is where we observe a broadening of PDF uncertainties.

Throughout the paper we found that the results obtained with the **SIMUnet** methodology are in line with those presented in ref. [17]. In table 4 we make this comparison more quantitative, by focussing on the results in which the effects of the interplay between the SMEFT coefficients and the PDFs are more visible, namely the fits obtained by using the NC and CC projections from the HL-LHC. The results are comparable, although the effect of fitting the Wilson coefficients along with the PDFs is more moderate. This is not surprising, as there are two crucial differences in the two analyses. On the one hand the PDF set that we use here in the fixed SM PDF case is different compared to the one we used in the previous analysis, being based on the same dataset as before but on the NNPDF4.0 methodology rather than the NNPDF3.1 methodology. Secondly, because the previous methodology was based on the use of Benchmark Points in the Wilson coefficients parameter space, we determined the bounds on the parameters by using the partial  $\chi^2$  including only the data affected by the SMEFT corrections, rather than the global  $\chi^2$ . This approximation was forced because the statistical fluctuations of the global  $\chi^2$  were found to be significantly larger than those of the partial  $\chi^2$  and could only be tamed by running a very large batch of replicas for each Benchmark Point and by increasing the density of Benchmark Points in the region that is explored. This approximation is no longer necessary within **SIMUnet**, because we no longer rely on the interpolation over Benchmark Points, rather we perform a truly simultaneous fit of the PDFs and the Wilson coefficients based on the global  $\chi^2$ .

	SM PDFs	SMEFT PDFs	best-fit shift	broadening
$W \times 10^5$ (this work)	[-2.0, 1.4]	[-4.3, 3.4]	-0.2	+126%
$W \times 10^5$ [17]	[-1.4, 1.2]	[-8.1, 10.6]	-1.4	+620%
$Y \times 10^5$ (this work)	[-3.2, 8.1]	[-3.1, 11.7]	+1.9	+31%
$Y \times 10^5$ [17]	[-5.3, 6.3]	[-11.1, 12.6]	+0.3	+110%
$C_{33}^{D\mu} \times 10^3$ (this work)	[-11.9, 13.1]	[-15.3, 14.0]	-0.9	+17%
$C_{33}^{D\mu} \times 10^3$ [17]	[-10.4, 12.3]	[-12.5, 14.6]	-0.6	+18%

**Table 4.** 95% C.L. bounds on the simultaneous fit of the  $W$  and  $Y$  Wilson coefficients in Benchmark Scenario I and of the individual fit of the  $C_{33}^{D\mu}$  Wilson Coefficient in Benchmark Scenario II, based on a fit including the HL-LHC projections, compared to those obtained in the previous analysis presented in ref. [17]. The fourth and fifth column indicate the absolute shift in best-fit values, eq. (4.3) and the percentage broadening of the SMEFT bounds, eq. (4.4), when the PDFs are allowed to change alongside the Wilson coefficients.


Additionally, the minimizer used in this study employs a momentum driven stochastic gradient descent based algorithm (Nesterov-accelerated adaptive moment estimation [69, 70] available from the *Keras* library), while in ref. [17] we employed a more traditional genetic algorithm approach using legacy in-house implementations. As such, one in general expects to achieve an improved fit quality with our methodology. Finally, the bounds quoted in our study are defined using Monte Carlo based statistical estimators, whereas in the latter approach, the geometry of the  $\chi^2$  profile is used to define bounds on the Wilson coefficients. Notwithstanding, we observe that the results and the trends are consistent with each other, although the use of the partial  $\chi^2$  over-emphasises the broadening of the bounds.

Results in the HL-LHC scenario are displayed in figure 12. Shown are the results obtained including the current high-mass Drell-Yan data and the projected HL-LHC NC and CC Drell-Yan pseudo-data. In Scenario I,  $W$  and  $Y$  are fitted simultaneously by keeping only the linear term in the SMEFT expansion, while in Scenario II the  $C_{33}^{D\mu}$  coefficients is fitted individually including the  $\mathcal{O}(1/\Lambda^4)$  quadratic terms in the SMEFT expansion. The results from both studies are compatible, although the bounds obtained from our study are slightly more conservative with differences being explained by the differences between the methodologies employed that are outlined at in this section.

## 5 Methodology validation and closure testing

In this section, we assess the robustness of our approach through the closure testing framework defined in refs. [71, 72]. Here we do not address the robustness of the PDF part of the fit, which in a sense comes from the results presented in the NNPDF4.0 paper [2] and the following dedicated study of ref. [72], but rather focus on the robustness of the fit of the Wilson coefficients and of the PDFs in a simultaneous fit.

Crucially, the central values of the SMEFT-sensitive datasets are modified by artificially contaminating them with pre-chosen, extreme, values of Wilson coefficients. This



**Figure 12.** Comparison of 95% C.L. bounds between this study (in green) and those determined in by Greljo et al. in ref. [17]. The  $W$  and  $Y$  bounds are determined simultaneously in the HL-LHC scenario,  $C_{33}^{D\mu}$  is fitted individually including SMEFT quadratic terms in the HL-LHC scenario. The bounds obtained by keeping the PDFs fixed to the SM baseline (dashed lines) are compared to those obtained in a simultaneous fit of PDFs and Wilson coefficients (solid lines).


will emulate a situation where a dataset will favour a non-vanishing EFT operator. In the first part of this section we show that our methodology is sufficiently flexible to recover these chosen values. In the second part of the section we consider the case where the input datasets are replaced by theory predictions from a known underlying PDF set, rather than experimental central values. We show that even in the case of fixing both the underlying PDF and Wilson coefficients to some *a priori* values, the methodology is sufficiently robust that it can simultaneously reproduce both.

### 5.1 Closure test results on the Wilson coefficients

The PDF side of `SIMUnet` is of course identical to `NMPDF4.0` which has been demonstrated to be able to replicate an underlying law using the closure testing framework [2, 72]. In this section we will fix the value of the Wilson coefficients *a priori* and study how effectively we can retrieve these values using our methodology.

To perform this kind of closure tests, we artificially choose extreme values of the Wilson coefficients considered in Benchmark Scenario I, where the interplay between PDF and SMEFT fits is more marked, in two separate analyses:

1. The individual fit of the  $W$  Wilson coefficient including the current high-mass DY data from LHC Run I and Run II along with all other datasets listed in section 4.2.



**Figure 13.** Result of the closure testing framework for our methodology. Left: histogram for the distribution of the  $W$  parameter when the input data has been modified by setting  $W = 1 \times 10^{-3}$ . Right: distribution of  $(W, Y)$  when fitting to data that has been modified by setting  $(W, Y) = (1, -1) \times 10^{-4}$ . The upper and right panels show the histograms for the distribution of the best fit values in their respective directions.

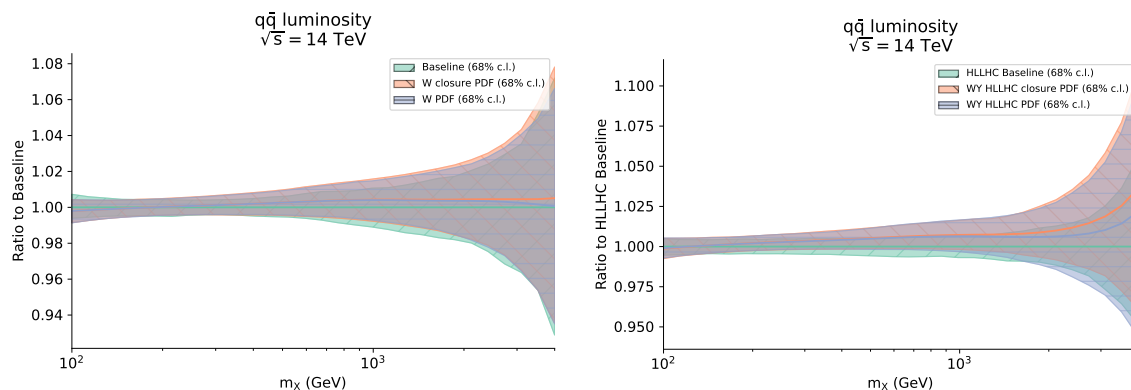
Specifically we set  $W = 1 \times 10^{-3}$ , which is outside the 95% C.L. bounds that are displayed in figure 5

2. The combined fit of the  $(W, Y)$  parameters including the HL-LHC projections. Specifically we set  $(W, Y) = (1, -1) \times 10^{-4}$ , which is also far outside the 95% C.L. contours displayed in figure 8.

The way we input these non-zero values of the Wilson coefficients in the underlying law is by multiplying the Monte Carlo pseudodata central values by the SMEFT  $K$ -factors obtained by setting the Wilson coefficient(s) to the aforementioned values. The fitting methodology proceeds as before. Importantly, the entire tool chain has no knowledge of what the value of the Wilson coefficient it is looking for are set to.

The result of the closure test are displayed in figure 13. On the left-hand panel we see that the distribution of best fit values of  $W$  across 1000 replicas, that were previously centred at the origin, now moves considerably towards the *a priori* value of  $W = 10^{-3}$ . Similarly, in the second analysis displayed on the right-hand panel, we see that the presence of HL-LHC data continues to eliminate the flat direction, with the distribution of best fit Wilson coefficients resembling that of figure 8, but with the best-fits values of the  $(W, Y)$  parameters consistently pushed towards the preselected *a priori* values of  $(W, Y) = (1, -1) \times 10^{-4}$ . We see that, despite an extremal choice of the injected values of the Wilson coefficients, the methodology is sufficiently flexible and robust to recover them.

The PDFs generated in the context of the closure tests are displayed in figure 14. On the left panel, we plot the quark-antiquark luminosity obtained when the experimental central values are modified by inputting  $W = 10^{-3}$  (called “W closure PDF”) and compare it to the ones that we obtain in the simultaneously fit of the PDFs and  $W$  presented in section 4.3 (called “W PDF”), both normalised to the SM baseline. We observe that the PDFs generated with our closure test for  $W = 10^{-3}$  are similar to those that we obtain in the simultaneous fit, despite the fact that the training dataset has been heavily




**Figure 14.** The  $q\bar{q}$  luminosity obtained from the closure tests when the experimental central values are modified to encode a specific SMEFT benchmarking scenario (“closure PDF” in orange), compared to the results of the simultaneous fits (“PDF” in blue) normalized to the corresponding baseline (in green). Shown in the left panel is the  $W = 1 \times 10^{-3}$  scenario while the  $(W, Y) = (1, -1) \times 10^{-4}$  scenario is displayed in the right panel.

modified by the extreme *a priori* choice of  $W$ . This is to be expected, since one can view the combination layer as capturing the data’s dependence on the Wilson coefficients, whilst the complementary PDF sector of the network architecture captures the data’s dependence on the underlying PDF, which of course remains unchanged. The combination layer, in effect, subtracts off the EFT dependence, leaving behind the pure SM contribution for the PDF sector to parameterise. This is even more evident in the fit obtained in the presence of the HL-LHC projections, by setting  $(W, Y) = (1, -1) \times 10^{-4}$  displayed on the right-hand panel of the figure. There we can observe that the central value of the “WY HLLHC closure PDF” luminosity is pushed slightly upwards and the uncertainty is rather smaller once the closure test is performed as compared to the simultaneous fit “WY HLLHC PDF”. This is once again to be expected, given that the range of Wilson Coefficients allowed in the simultaneous fit is larger as compared to the range that we obtain from a simultaneous fit.

## 5.2 Closure test results on the simultaneous fit

The natural extension of the closure test described in the previous subsection is to assess the degree to which SIMUnet is able to replicate, not only fixed Wilson coefficients, but also a known underlying PDF. For this scenario we employ the NNPDF level 2 [52] closure test strategy. In the context of a simultaneous fitting methodology this amounts to generating Standard Model predictions using a known PDF set (referred to henceforth as the underlying law) and mapping these to SMEFT observables by multiplying the SM theory predictions with SMEFT K-factors scaled by a previously determined choice of Wilson coefficients. These SMEFT observables, generated by the underlying law, replace the usual MC pseudodata replicas, and are used to train the neural network in the usual way.

Through this approach we are able to assess the degree to which the parameterisation is able to capture not only an underlying choice of Wilson coefficients, but ensures it is sufficiently flexible to adequately replicate a known PDF. For such a closure test, we use



**Figure 15.** The gluon (left) and up quark (right) PDFs obtained from the closure test framework in which both the underlying PDF set and Wilson coefficients are known. Shown in green is the PDF replica used as the underlying law which generates the fake data used to train our model. The resulting PDFs are shown in orange along with their 68% confidence level bands. The fake data generated by the underlying law is subsequently modified so as to encode the  $(W, Y) = (1, -1) \times 10^{-4}$  condition.


the HL-LHC baseline used in this study as the underlying law with the SMEFT scenario being again the simultaneous  $(W, Y)$  determination, with the input data being adjusted to have  $(W, Y) = (1, -1) \times 10^{-4}$  encoded within it.

The PDFs generated in this way are shown in figure 15 for representative choices of parton flavour: we display both the resulting PDFs as well as the underlying law. We see that, despite modifying the training data with an extreme choice of SMEFT benchmark, the methodology is sufficiently robust so as to be able to recover the true PDF set with good precision. Indeed, the distribution of best fit  $(W, Y)$  values plotted in figure 16 shows that not only did the methodology retrieve the underlying law, but also managed to recover the chosen SMEFT scenario: being able to correctly determine the  $(W, Y) = (1, -1) \times 10^{-4}$  condition. Such behaviour is reminiscent of the above closure test whereby the combination layer is able to parameterise the BSM dynamics while the preceding layers of the model are left to parameterise the PDFs. This conclusion thus illustrates how our methodology is able to correctly disentangle the interplay between PDFs and BSM dynamics.

## 6 Conclusions and outlook

In this work we have presented a novel methodology, dubbed **SIMUnet**, which employs the latest in deep learning techniques, to achieve a first truly simultaneous fit of PDFs alongside any physical parameter that determines theoretical predictions at the LHC. As a showcase we fit the Wilson coefficients in the SMEFT and the PDFs in two motivated Benchmark Scenarios.

In particular we employ our methodology to constrain each of  $W$  and  $Y$  individually using Run I and Run II ATLAS and CMS neutral current (NC) high mass DY data and show that the interplay between SMEFT effects and the fit of large- $x$  PDFs remain moderate. Interestingly we show how a scenario like this, displaying a flat direction once  $W$  and  $Y$  are



**Figure 16.** Distribution of best fit  $(W, Y)$  parameters for each Monte Carlo replica. The data used for the fit was generated by a pre-selected PDF before the SM predictions were transformed to a SMEFT observable with  $(W, Y) = (1, -1) \times 10^{-4}$ . The upper panel is the histogram of best fit values in the  $W$  axis while the right panel is the histogram for the  $Y$  axis.

fitted simultaneously, are uncovered naturally by our approach, without the need for any user intervention or prior knowledge, and how the framework comes to learn that the flat direction is lifted by the inclusion of such data that eliminates this redundancy. Indeed, once the HL-LHC projections are added to the fit, such projections including charged current (CC) DY data, the flat direction is broken and much tighter bounds are obtained.

Most importantly, we confirm the results already shown in our previous analysis of ref. [17], namely that the inclusion of high-mass DY data both in a fit of PDFs and in a fit of SMEFT coefficients by neglecting the interplay between them, could result both in a significant underestimate of the uncertainties associated to the SMEFT parameters (by a factor of 2 in the case of the  $W$  coefficient) and of the large- $x$  PDF uncertainties, that increase by about 70% in the high-mass luminosity regions once the SMEFT affects are accounted for.

We also showcase the ability of our framework to fit SMEFT parameters while including the quadratic effect of SMEFT-SMEFT diagrams. Such a scenario is known [17] to be essentially unconstrained at linear EFT level and we illustrate how our methodology replicates bounds for  $C_{33}^{D\mu}$  presented in previous analyses.

We additionally stress test the methodology by providing it artificially contaminated data possessing data shifted by an *a priori* choice of Wilson coefficients. We show that our methodology is sufficiently flexible to recover these chosen values. Moreover, we consider the case where the input datasets are replaced by theory predictions from a known underlying PDF set, rather than experimental central values. We show that even in the case of fixing both the underlying PDF and Wilson coefficients to known values, the methodology is sufficiently robust that it can simultaneously reproduce both.

The new methodology that we present here provides a crucial step towards the interpretation of indirect searches under a unified framework. We can assess the impact of new



datasets not only on the PDFs, but now on the couplings of an EFT expansion, or on any other physical parameter.

The next steps of our work will include more SMEFT parameters to the combination layer, which is conceptually straightforward. We can add the effects of dim-8 operators or consider the effect of quadratic dim-6 operator effects. Both these scenarios fit very naturally within our approach. Furthermore, one could assess the impact of other high energy datasets from the LHC, exploiting different processes to better constrain other subsets of the Warsaw basis. Finally, an interesting prospect would be to explore simultaneously fitting precision SM parameters, such as the strong coupling or electroweak parameters. Indeed, our framework extends naturally, not only to BSM studies, but to any parameter which may modify the SM prediction through the use of K-factors or other kind of interpolation.

## Acknowledgments

We thank Richard Ball, Luigi Del Debbio, Juan Cruz-Martinez, Ben Day, Stefano Forte, James Moore, Juan Rojo, and Zahari Kassabov for particularly useful discussion. We thank the NNPDF collaboration for the public release of the code. We thank the PBSP collaboration for providing the relevant SMEFT  $K$ -factors for the various datasets studied in this work. M.U. is supported by the European Research Council under the European Union’s Horizon 2020 research and innovation Programme (grant agreement n. 950246). M.U. and S.I. are supported by the Royal Society grant RGF/EA/180148. The work of M.U. is also funded by the Royal Society grant DH150088. M.U. and S.I. are partially supported by STFC consolidated grant ST/T000694/1.

## A FK table dependence on the strong coupling

In this appendix we consider the case whereby we wish to isolate the FK table dependence on the strong coupling constant at the  $Z$ -pole, such that, following our notation in section 2, we have  $c = \alpha_s(M_Z)$ . We take the PDG value for the strong coupling evaluated at the  $Z$ -boson mass  $c^* = \alpha_s(M_Z) = 0.1179(10)$  [6] to be the point about which we perform the Taylor expansion. The partonic cross sections are given as a power expansion in  $\alpha_s$  using perturbation theory, which allows us to write the exact expression

$$\hat{\sigma}(c) = \sum_p (c - c^*)^p \hat{\sigma}_p(c^*) \tag{A.1}$$

where we have simply centred the order by order sum around  $c^*$ . For illustrative purposes, we restrict ourselves to the case where the process depends solely on the non-singlet quark distribution. The evolution kernel in Mellin space to leading order in the anomalous dimension (the Mellin transform of the splitting function)  $\gamma_0$ , is then given by [42]:

$$\Gamma(N, \alpha_s(Q^2), \alpha_s(Q_0^2)) = \left( \frac{\alpha_s(Q^2)}{\alpha_s(Q_0^2)} \right)^{-\gamma_0(N)/\beta_0} . \tag{A.2}$$

If we consider the leading-log evolution of  $\alpha_s$  via renormalisation group equation from  $M_Z^2$  to  $Q^2$  or  $Q_0^2$  and set  $c = \alpha_s(M_Z^2)$ , we get

$$\alpha_s(Q^2) = \frac{c}{1 + \beta_0 c \ln \frac{Q^2}{M_Z^2}} \quad \alpha_s(Q_0^2) = \frac{c}{1 + \beta_0 c \ln \frac{Q_0^2}{M_Z^2}} \quad (\text{A.3})$$

where  $\beta_0 = 11 - \frac{2}{3}N_f$ . Expanding the evolution kernel  $\Gamma$  according to eq. (2.6), taking  $c^* = \alpha_s^{\text{PDG}}(M_Z) = 0.1179$ , it is easy to see that we get

$$\Gamma(c) = \Gamma_0 + (c - c^*)\Gamma_1 + \dots, \quad (\text{A.4})$$

where

$$\Gamma_0 = \frac{1 + c^* \beta_0 \ln \frac{Q_0^2}{M_Z^2}}{1 + c^* \beta_0 \ln \frac{Q^2}{M_Z^2}} \quad (\text{A.5})$$

$$\Gamma_1 = -\frac{\gamma_0(N)}{\beta_0} \Gamma_0^{-\frac{\gamma_0(N)}{\beta_0} - 1} \left( \frac{\beta_0 \ln \frac{Q_0^2}{M_Z^2}}{1 + c^* \beta_0 \ln \frac{Q^2}{M_Z^2}} - \frac{\left(1 + c^* \beta_0 \ln \frac{Q_0^2}{M_Z^2}\right) \beta_0 \ln \frac{Q^2}{M_Z^2}}{\left(1 + c^* \beta_0 \ln \frac{Q^2}{M_Z^2}\right)^2} \right). \quad (\text{A.6})$$

Finally, once eq. (A.4) is combined with eq. (A.1) we obtain the order by order expansion for the FK tables:

$$K = \hat{\sigma}_0 \otimes \Gamma_0 + (c - c^*)(\hat{\sigma}_1 \otimes \Gamma_0 + \hat{\sigma}_0 \otimes \Gamma_1) + \dots \quad (\text{A.7})$$

$$\equiv K_0 + (c - c^*)K_1 + \dots \quad (\text{A.8})$$

which can, in principle, be computed and stored permanently in the usual way. In this way we can capture the FK table dependence on the strong coupling parameter in the neighborhood of some prior choice  $c^*$ .

So far, we have limited ourselves to the leading order evolution of the PDFs and the strong coupling constant, however a similar expression can be straightforwardly obtained at next-to-leading (NLO) and next-to-next-to-leading order (NNLO). The only caveat is that, from NLO onwards terms of the form  $\ln \left(1 + \beta_0 c \ln \frac{Q^2}{M_Z^2}\right)$  arise which spoil the validity of the Taylor expansion. The term  $\ln \frac{Q^2}{M_Z^2}$  can in principle be made arbitrarily large thus exiting the unit disc which is the region of analyticity of  $\ln 1 + x$ . This problem can be circumvented, however, by noting:

$$\ln \left(1 + \beta_0 c \ln \frac{Q^2}{M_Z^2}\right) = \ln \left(1 + \beta_0 c^* \ln \frac{Q^2}{M_Z^2}\right) + \ln \left(1 + \frac{\beta_0 (c - c^*) \ln \frac{Q^2}{M_Z^2}}{1 + \beta_0 c^* \ln \frac{Q^2}{M_Z^2}}\right) \quad (\text{A.9})$$

where the rightmost term can now be Taylor expanded since  $c - c^*$  can be made arbitrarily small so as to suppress the large logarithm. Indeed, note that in the worst case:

$$\lim_{Q^2 \rightarrow \infty} \left( (c - c^*) \frac{\beta_0 \ln \frac{Q^2}{M_Z^2}}{1 + \beta_0 c^* \frac{Q^2}{M_Z^2}} \right) = \frac{c - c^*}{c^*} \leq 1 \quad (\text{A.10})$$

$$\implies c \leq 2c^* \quad (\text{A.11})$$

which can be implemented within the optimizer to restrict it from venturing to values greater than  $2c^*$ .

A viable alternative to the Taylor expansion approach defined in section 2 and detailed here for the simultaneous extraction of  $\alpha_s$  and the PDFs is to compute various FK tables using preset values of the strong coupling and then perform an element-wise interpolation between these tensors. In this way one can avoid having to implement the Taylor series expansion, while still accurately replicating the linear behaviour of the FK table. Such a task is reserved for future work [73].

Finally if one was to include electroweak corrections to the DGLAP evolution equation, as is for example done in APFEL [74], then electroweak parameters will in general also be present in the combined QCD and QED evolution operator. For such a scenario, the prescription outlined in this appendix must then be followed. However, the corrections to the pure QCD splitting functions introduced by electroweak considerations have no dependence on the CKM matrix elements, the weak mixing angle,  $\theta_W$ , or gauge boson masses, amongst others. Such quantities manifest solely in the partonic cross section and in general yield a prescription much simpler than that outlined here.

## B Fit quality

For each MC replica there is the pair  $(\mathbf{f}_i, \mathbf{c}_i)$  which are the best fit PDFs and Wilson coefficients respectively. Together they can be used to generate the corresponding theory predictions for each dataset, noting that any datasets which were not modified by the SMEFT operators will effectively have  $\mathbf{c} = \mathbf{0}$ . We thus have an ensemble of  $N_{\text{rep}}$  vectors of theory predictions,  $\mathbf{T}$ , with the central theory prediction being given by the average across replicas:

$$\langle \mathbf{T} \rangle = \frac{1}{N_{\text{rep}}} \sum_{i=1}^{N_{\text{rep}}} \mathbf{T}(\mathbf{f}_i, \mathbf{c}_i). \tag{B.1}$$

The central  $\chi^2$  per data point is then computed in the usual way

$$\chi^2 = \frac{1}{N_{\text{dat}}} (\mathbf{d} - \langle \mathbf{T} \rangle)^T V^{-1} (\mathbf{d} - \langle \mathbf{T} \rangle) \tag{B.2}$$

with  $\mathbf{d}$  being the vector of experimental central values and  $V$  the covariance matrix encapsulating the experimental uncertainties and the correlations therein. These values are tabulated in table 5 for each of the various SMEFT scenarios considered in this work and in table 6 for the high mass Drell-Yan measurements in the muonphilic operator case. We also tabulate the  $\chi^2$  for various groupings of these datasets, such as DIS only, including or excluding the high-mass DY measurements etc. For these particular entries various correlated systematics may exist between datasets, such as the uncertainty in beam luminosity, which introduces off-diagonal entries in the covariance matrix; as such the grouped  $\chi^2$  is not necessarily equal to the weighted average of the individual  $\chi^2$  values constituting the grouping. This effect is particularly marked for the HL-LHC entries.

In all scenarios considered, the added degrees of freedom result in the  $\chi^2$  per data point to drop when a simultaneous determination is performed when contrasted to the


Dataset	$n_{\text{data}}$	$\chi^2/n_{\text{data}}$				
		SM Baseline	$W$	$Y$	HL-LHC SM Baseline	HL-LHC ( $W, Y$ )
SLAC	67	0.866	0.855	0.854	0.837	0.850
BCDMS	581	1.285	1.265	1.265	1.294	1.266
NMC	325	1.320	1.318	1.318	1.333	1.320
CHORUS	832	1.208	1.209	1.210	1.206	1.208
NuTeV	76	0.444	0.486	0.487	0.472	0.497
HERA inclusive	1145	1.188	1.183	1.183	1.190	1.184
HERA charm	37	1.435	1.384	1.382	1.463	1.391
HERA bottom	29	1.113	1.110	1.110	1.118	1.110
<b>Total DIS</b>	<b>3092</b>	<b>1.202</b>	<b>1.197</b>	<b>1.197</b>	<b>1.206</b>	<b>1.198</b>
E886 $\sigma_{DY}^d/\sigma_{DY}^p$	15	0.669	0.609	0.600	0.974	0.679
E886 $\sigma_{DY}^p$	89	1.572	1.593	1.604	1.566	1.631
E605 $\sigma_{DY}^p$	85	1.197	1.200	1.205	1.197	1.213
CDF $d\sigma_Z/dy_Z$	29	1.613	1.546	1.548	1.623	1.563
D0 $d\sigma_Z/dy_Z$	28	0.612	0.610	0.610	0.613	0.613
D0 $W \rightarrow \mu\nu$ asy.	9	1.845	1.510	1.503	2.052	1.587
ATLAS $W, Z$ 2010	30	1.021	1.014	1.014	1.017	1.017
ATLAS low-mass $Z \rightarrow ee$	6	0.923	0.921	0.921	0.923	0.921
ATLAS $W, Z$ 2011 CC	46	2.095	2.005	2.006	2.091	2.010
ATLAS $W, Z$ 2011 CF	15	1.062	1.072	1.073	1.066	1.071
ATLAS $W + c$ rapidity	22	0.453	0.463	0.463	0.447	0.457
ATLAS $Z_{pT}$	92	0.959	0.938	0.936	0.939	0.928
ATLAS $W_{pT}$ jets	32	1.685	1.672	1.670	1.676	1.665
CMS $W e$ asy.	11	0.785	0.790	0.789	0.815	0.804
CMS $W \mu$ asy.	11	1.767	1.732	1.733	1.765	1.732
CMS $\sigma_{W+c}$ 7 TeV	5	0.513	0.518	0.516	0.502	0.504
CMS $\sigma_{W+c}/\sigma_{W-c}$ 7 TeV	5	1.822	1.791	1.796	1.884	1.848
CMS $Z_{pT}$	28	1.303	1.312	1.311	1.287	1.306
CMS $W \rightarrow \mu\nu$ rapidity	22	1.472	1.337	1.340	1.422	1.310
CMS $W + c$ rapidity 13 TeV	5	0.719	0.722	0.721	0.712	0.711
LHCb $Z \rightarrow \mu\mu$	9	1.503	1.545	1.550	1.506	1.549
LHCb $W, Z \rightarrow \mu$ 7 TeV	29	2.043	1.973	1.977	2.066	2.005
LHCb $W, Z \rightarrow ee$	17	1.249	1.236	1.236	1.220	1.231
LHCb $W, Z \rightarrow \mu$ 8 TeV	30	1.621	1.497	1.502	1.615	1.543
<b>Total DY (excl. HM)</b>	<b>670</b>	<b>1.302</b>	<b>1.274</b>	<b>1.276</b>	<b>1.307</b>	<b>1.286</b>
ATLAS DY high-mass 7 TeV	13	1.680	1.575	1.609	1.654	1.626
ATLAS DY high-mass 8 TeV	46	1.174	1.177	1.171	1.175	1.171
CMS DY high-mass 7 TeV	117	1.694	1.671	1.676	1.677	1.669
CMS DY high-mass 8 TeV	41	0.923	0.944	0.941	0.893	0.914
CMS DY high-mass 13 TeV	43	2.003	2.064	2.037	2.000	2.005
<b>Total DY (HM only)</b>	<b>260</b>	<b>1.531</b>	<b>1.529</b>	<b>1.527</b>	<b>1.517</b>	<b>1.515</b>
<b>Total (excl. HL-LHC)</b>	<b>4022</b>	<b>1.245</b>	<b>1.236</b>	<b>1.237</b>	<b>1.248</b>	<b>1.238</b>
HL-LHC CC $e$	16	<i>1.119</i>	<i>119.9</i>	<i>0.922</i>	0.588	0.544
HL-LHC CC $\mu$	16	<i>1.414</i>	<i>112.8</i>	<i>1.162</i>	0.894	0.803
HL-LHC NC $e$	12	<i>1.164</i>	<i>17.65</i>	<i>7.495</i>	1.104	0.961
HL-LHC NC $\mu$	12	<i>1.041</i>	<i>13.32</i>	<i>5.048</i>	0.964	1.071
<b>Total HL-LHC only</b>	<b>56</b>	<b>1.298</b>	<b>72.88</b>	<b>5.546</b>	<b>0.894</b>	<b>0.836</b>
<b>Total</b>	<b>4078</b>	<b>1.246</b>	<b>2.220</b>	<b>1.296</b>	<b>1.243</b>	<b>1.232</b>

**Table 5.** Values of the  $\chi^2$  per data point across all datasets used in this study. We tabulate values for the baseline PDF set as well as those obtained in the various SMEFT scenarios. Shown also is the  $\chi^2$  for the HL-LHC scenario. The rows indicating total  $\chi^2$  values are computed accounting for any relevant correlated systematic errors. Values in italics indicate the dataset was not used in the corresponding fit.

Dataset	$n_{\text{data}}$	SM Baseline	$C_{33}^{D\mu}$
ATLAS DY high-mass 7 TeV	13	1.631	1.625
ATLAS DY high-mass 8 TeV	46	1.179	1.168
CMS DY high-mass 7 TeV	117	1.666	1.664
CMS DY high-mass 8 TeV	41	0.858	0.893
CMS DY high-mass 13 TeV $ee^*$	43	2.579	2.574
CMS DY high-mass 13 TeV $\mu\mu$	43	0.836	0.837
<b>Total DY (HM only)</b>	<b>303</b>	1.501	1.503
HL-LHC CC $e^*$	16	0.437	0.424
HL-LHC CC $\mu^*$	16	0.841	0.752
HL-LHC NC $e^*$	12	1.108	1.130
HL-LHC NC $\mu$	12	0.943	0.969
<b>Total HL-LHC only</b>	<b>56</b>	0.830	0.826

**Table 6.** Values of the  $\chi^2$  per data point across all high mass Drell-Yan measurements. We tabulate values for the baseline PDF set as well as those obtained in the EFT scenario II. Shown also is the  $\chi^2$  for the HL-LHC projections. Quadratic terms in the EFT parameter are used in the  $\chi^2$  calculation. The rows indicating total  $\chi^2$  values are computed accounting for any relevant correlated systematic errors. Datasets marked with an asterisk indicate that they were not used to constrain  $C_{33}^{D\mu}$ .

purely Standard Model fits. The SM  $\chi^2$  in this sense serves as an upper bound, since the optimizer is free to determine  $\mathbf{c} = 0$  and so the goodness of fit can be no worse than the SM fit. In particular we see the SMEFT sensitive high-mass DY datasets experience a large overall improvement in fit quality, largely driven by the CMS measurements at 7 TeV owing to the large weight carried by this dataset: forming just under half the entire high-mass DY data points considered in this study. Moreover, the DIS only grouping experiences a marked improvement in fit quality with the  $\chi^2$  per data point dropping by 0.014 across 3092 data points in the case of the simultaneous  $(W, Y)$  determination. The reader is again reminded that the HERA combined dataset, which forms the majority of the DIS data points, is included in the set of datasets that are modified by the  $W$  and  $Y$  operators. It is interesting to observe that the improvement fit quality is propagated down to those datasets, such as the low-mass measurements, that are not used to explicitly constrain the Wilson coefficients. Such datasets, however, are correlated through shared sources of systematic errors, such as the luminosity uncertainty or detector effects, thus improving the fit to one such dataset necessarily affects others. Also tabulated are the fit quality values for the HL-LHC projections, even for those fits which do not incorporate these projections in their training data. These entries illustrate the pull the HL-LHC projections have on the Wilson coefficients. We see that not including these data points in the fit renders the fit virtually useless in the context of the projected data. This is indeed reflected in the




**Figure 17.** The  $q\bar{q}$  luminosity of the 100 replica (low statistics) fits normalized to the 1000 replica (high statistics) fits for various SMEFT scenarios considered in this work. Shown in the top left (top right) is the individual  $W$  ( $Y$ ) scenario, while the lower panel is the combined ( $W, Y$ ) scenario with HL-LHC projections being used in the fit.

values of the K-factors used for these projections, reaching  $K \simeq 5$  for the highest transverse mass bins.

### C Stability on replica number

When performing purely PDF fits to experimental data, it is often the case one will have to compute  $\sim 100$  Monte Carlo PDF replicas in order to achieve a percent level, faithful, uncertainty estimation [22]. As this study acts as a proof-of-concept for our methodology, we perform high statistic fits composed of 1000 MC replicas. However, with each replica requiring approximately 3 hours of compute time, one necessarily requires access to a cluster of nodes in order to asynchronously compute the roughly 3000 hours of total wall clock time that is needed for each fit. It is possible, however, to reduce this time by an order of magnitude by instead computing  $\sim 100$  MC replicas per fit, and in this appendix we show that doing so poses little risk in underestimating the statistics when compared to a high replica fit.

From each of our high replica fits we randomly sample, without replacement, a set of 100 replicas thereby effectively emulating the scenario whereby the user would have performed a low statistics fit. In figure 17 we plot the  $q\bar{q}$  luminosity for the various SMEFT



**Figure 18.** Distributions of the best fit  $W$  (left) and  $Y$  (right) parameters using both high ( $N_{\text{rep}} = 1000$ ) and low ( $N_{\text{rep}} = 100$ ) statistics fits. We bring the readers attention to the different  $y$  axes for the histogram overlays.

scenarios considered in this work. We plot the low replica luminosity normalized to the analogous high replica set. We see that the luminosities remain virtually identical, with no discernible difference at the ensemble level. Such behaviour is inherited directly from NNPDF4.0 where a typical fit will typically only possess 100 MC replicas after the post-fit selection has filtered poorly performing replicas.

In figure 18 we present the distributions of best-fit  $W$  and  $Y$  values using both the high replica and reduced sets. We see that the distribution of best fit Wilson coefficients are accurately reproduced by the low statistics set: implying that, had we stopped the fitting process with 100 MC replicas, the additional 900 replicas would have changed the ensemble statistics such as the mean, standard deviation, or bounds very little. This is a typical pattern in the Monte Carlo approach to PDF fitting, whereby one quickly reaches saturation after  $\sim 100$  MC replicas and further replicas only serve to accurately reproduce the experimental correlations.

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

## References

- [1] D. Britzger, *The Linear Template Fit*, [arXiv:2112.01548](https://arxiv.org/abs/2112.01548) [[INSPIRE](#)].
- [2] R.D. Ball et al., *The Path to Proton Structure at One-Percent Accuracy*, [arXiv:2109.02653](https://arxiv.org/abs/2109.02653) [[INSPIRE](#)].
- [3] S. Bailey, T. Cridge, L.A. Harland-Lang, A.D. Martin and R.S. Thorne, *Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs*, *Eur. Phys. J. C* **81** (2021) 341 [[arXiv:2012.04684](https://arxiv.org/abs/2012.04684)] [[INSPIRE](#)].
- [4] T.-J. Hou et al., *New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC*, *Phys. Rev. D* **103** (2021) 014013 [[arXiv:1912.10053](https://arxiv.org/abs/1912.10053)] [[INSPIRE](#)].
- [5] S. Alekhin, J. Blümlein, S. Moch and R. Placakyte, *Parton distribution functions,  $\alpha_s$ , and heavy-quark masses for LHC Run II*, *Phys. Rev. D* **96** (2017) 014011 [[arXiv:1701.05838](https://arxiv.org/abs/1701.05838)] [[INSPIRE](#)].

- [6] PARTICLE DATA collaboration, *Review of Particle Physics*, *Prog. Theor. Exp. Phys.* **2020** (2020) 083C01 [INSPIRE].
- [7] SMEFT collaboration, *Combined SMEFT interpretation of Higgs, diboson, and top quark data from the LHC*, *JHEP* **11** (2021) 089 [arXiv:2105.00006] [INSPIRE].
- [8] J. Ellis, M. Madigan, K. Mimasu, V. Sanz and T. You, *Top, Higgs, Diboson and Electroweak Fit to the Standard Model Effective Field Theory*, *JHEP* **04** (2021) 279 [arXiv:2012.02779] [INSPIRE].
- [9] S. Forte and Z. Kassabov, *Why  $\alpha_s$  cannot be determined from hadronic processes without simultaneously determining the parton distributions*, *Eur. Phys. J. C* **80** (2020) 182 [arXiv:2001.04986] [INSPIRE].
- [10] S. Alekhin, S. Moch and S. Thier, *Determination of the top-quark mass from hadro-production of single top-quarks*, *Phys. Lett. B* **763** (2016) 341 [arXiv:1608.05212] [INSPIRE].
- [11] E. Bagnaschi and A. Vicini, *Parton Density Uncertainties and the Determination of Electroweak Parameters at Hadron Colliders*, *Phys. Rev. Lett.* **126** (2021) 041801 [arXiv:1910.04726] [INSPIRE].
- [12] H1 collaboration, *Determination of electroweak parameters in polarised deep-inelastic scattering at HERA*, *Eur. Phys. J. C* **78** (2018) 777 [arXiv:1806.01176] [INSPIRE].
- [13] J.J. Ethier and E.R. Nocera, *Parton Distributions in Nucleons and Nuclei*, *Ann. Rev. Nucl. Part. Sci.* **70** (2020) 43 [arXiv:2001.07722] [INSPIRE].
- [14] R.A. Khalek, J.J. Ethier, E.R. Nocera and J. Rojo, *Self-consistent determination of proton and nuclear PDFs at the Electron Ion Collider*, *Phys. Rev. D* **103** (2021) 096005 [arXiv:2102.00018] [INSPIRE].
- [15] NNPDF collaboration, *Charged hadron fragmentation functions from collider data*, *Eur. Phys. J. C* **78** (2018) 651 [arXiv:1807.03310] [INSPIRE].
- [16] I. Brivio and M. Trott, *The Standard Model as an Effective Field Theory*, *Phys. Rept.* **793** (2019) 1 [arXiv:1706.08945] [INSPIRE].
- [17] A. Greljo et al., *Parton distributions in the SMEFT from high-energy Drell-Yan tails*, *JHEP* **07** (2021) 122 [arXiv:2104.02723] [INSPIRE].
- [18] S. Carrazza, C. Degrande, S. Iranipour, J. Rojo and M. Ubiali, *Can New Physics hide inside the proton?*, *Phys. Rev. Lett.* **123** (2019) 132001 [arXiv:1905.05215] [INSPIRE].
- [19] ZEUS collaboration, *Limits on contact interactions and leptoquarks at HERA*, *Phys. Rev. D* **99** (2019) 092006 [arXiv:1902.03048] [INSPIRE].
- [20] CMS collaboration, *Measurement and QCD analysis of double-differential inclusive jet cross sections in proton-proton collisions at  $\sqrt{s} = 13$  TeV*, *JHEP* **02** (2022) 142 [arXiv:2111.10431] [INSPIRE].
- [21] NNPDF collaboration, *An open-source machine learning framework for global analyses of parton distributions*, *Eur. Phys. J. C* **81** (2021) 958 [arXiv:2109.02671] [INSPIRE].
- [22] R.D. Ball et al., *A first unbiased global NLO determination of parton distributions and their uncertainties*, *Nucl. Phys. B* **838** (2010) 136 [arXiv:1002.4407] [INSPIRE].
- [23] R.D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244 [arXiv:1207.1303] [INSPIRE].



- [24] V. Bertone, S. Carrazza and N.P. Hartland, *APFELgrid: a high performance tool for parton density determinations*, *Comput. Phys. Commun.* **212** (2017) 205 [[arXiv:1605.02070](#)] [[INSPIRE](#)].
- [25] A.G. Baydin, B.A. Pearlmutter, A.A. Radul and J.M. Siskind, *Automatic differentiation in machine learning: a survey*, [arXiv:1502.05767](#).
- [26] M. González-Alonso, J.M. Camalich and K. Mimouni, *Renormalization-group evolution of new physics contributions to (semi)leptonic meson decays*, *Phys. Lett. B* **772** (2017) 777 [[arXiv:1706.00410](#)] [[INSPIRE](#)].
- [27] E.E. Jenkins, A.V. Manohar and M. Trott, *Renormalization Group Evolution of the Standard Model Dimension Six Operators. Part II. Yukawa Dependence*, *JHEP* **01** (2014) 035 [[arXiv:1310.4838](#)] [[INSPIRE](#)].
- [28] E.E. Jenkins, A.V. Manohar and M. Trott, *Renormalization Group Evolution of the Standard Model Dimension Six Operators. Part I. Formalism and lambda Dependence*, *JHEP* **10** (2013) 087 [[arXiv:1308.2627](#)] [[INSPIRE](#)].
- [29] A. Greljo and D. Marzocca, *High- $p_T$  dilepton tails and flavor physics*, *Eur. Phys. J. C* **77** (2017) 548 [[arXiv:1704.09015](#)] [[INSPIRE](#)].
- [30] R. Torre, L. Ricci and A. Wulzer, *On the  $W\mathcal{B}Y$  interpretation of high-energy Drell-Yan measurements*, *JHEP* **02** (2021) 144 [[arXiv:2008.12978](#)] [[INSPIRE](#)].
- [31] S. Carrazza and J. Cruz-Martinez, *Towards a new generation of parton densities with deep learning models*, *Eur. Phys. J. C* **79** (2019) 676 [[arXiv:1907.05075](#)] [[INSPIRE](#)].
- [32] NNPDF collaboration, *Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 663 [[arXiv:1706.00428](#)] [[INSPIRE](#)].
- [33] M. Abadi et al., *Tensorflow: A system for large-scale machine learning*, in proceedings of the *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, Savannah, GA, U.S.A., 2–4 November 2016, USENIX Association, Berkeley, CA, U.S.A. (2016), pp. 265–283 [[arXiv:1605.08695](#)].
- [34] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins and D.D. Cox, *Hyperopt: a Python library for model selection and hyperparameter optimization*, *Comput. Sci. Discov.* **8** (2015) 014008.
- [35] NNPDF collaboration, *Fitting Parton Distribution Data with Multiplicative Normalization Uncertainties*, *JHEP* **05** (2010) 075 [[arXiv:0912.2276](#)] [[INSPIRE](#)].
- [36] G. D’Agostini, *On the use of the covariance matrix to fit correlated data*, *Nucl. Instrum. Meth. A* **346** (1994) 306 [[INSPIRE](#)].
- [37] NNPDF collaboration, *Parton Distributions with Theory Uncertainties: General Formalism and First Phenomenological Studies*, *Eur. Phys. J. C* **79** (2019) 931 [[arXiv:1906.10698](#)] [[INSPIRE](#)].
- [38] NNPDF collaboration, *A first determination of parton distributions with theoretical uncertainties*, *Eur. Phys. J. C* **79** (2019) 838 [[arXiv:1905.04311](#)] [[INSPIRE](#)].
- [39] R.D. Ball, E.R. Nocera and R.L. Pearson, *Deuteron Uncertainties in the Determination of Proton PDFs*, *Eur. Phys. J. C* **81** (2021) 37 [[arXiv:2011.00009](#)] [[INSPIRE](#)].
- [40] C.M. Bishop et al., *Neural networks for pattern recognition*, Oxford University Press, Oxford, U.K. (1995).
- [41] NNPDF collaboration, *Neural network determination of parton distributions: The Nonsinglet case*, *JHEP* **03** (2007) 039 [[hep-ph/0701127](#)] [[INSPIRE](#)].

- [42] NNPDF collaboration, *A Determination of parton distributions with faithful uncertainty estimation*, *Nucl. Phys. B* **809** (2009) 1 [Erratum *ibid.* **816** (2009) 293] [[arXiv:0808.1231](#)] [[INSPIRE](#)].
- [43] J. Bergstra, D. Yamins and D. Cox, *Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures*, in proceedings of the *30th International Conference on Machine Learning*, Atlanta, GA, U.S.A., 17–19 June 2013, *Proceedings of Machine Learning Research* **28**, S. Dasgupta and D. McAllester eds., JMLR, Cambridge, MA, U.S.A. (2013), pp. 115–123.
- [44] R. Boughezal, E. Mereghetti and F. Petriello, *Dilepton production in the SMEFT at  $\mathcal{O}(1/\Lambda^4)$* , *Phys. Rev. D* **104** (2021) 095022 [[arXiv:2106.05337](#)] [[INSPIRE](#)].
- [45] C. Grojean, E.E. Jenkins, A.V. Manohar and M. Trott, *Renormalization Group Scaling of Higgs Operators and  $h \rightarrow \gamma\gamma$  decay*, *JHEP* **04** (2013) 016 [[arXiv:1301.2588](#)] [[INSPIRE](#)].
- [46] J. Baglio, S. Dawson and S. Homiller, *QCD corrections in Standard Model EFT fits to WZ and WW production*, *Phys. Rev. D* **100** (2019) 113010 [[arXiv:1909.11576](#)] [[INSPIRE](#)].
- [47] C. Degrande, G. Durieux, F. Maltoni, K. Mimasu, E. Vryonidou and C. Zhang, *Automated one-loop computations in the standard model effective field theory*, *Phys. Rev. D* **103** (2021) 096024 [[arXiv:2008.11743](#)] [[INSPIRE](#)].
- [48] S. Dawson and P.P. Giardino, *New physics through Drell-Yan standard model EFT measurements at NLO*, *Phys. Rev. D* **104** (2021) 073004 [[arXiv:2105.05852](#)] [[INSPIRE](#)].
- [49] H.-L. Li, Z. Ren, J. Shu, M.-L. Xiao, J.-H. Yu and Y.-H. Zheng, *Complete set of dimension-eight operators in the standard model effective field theory*, *Phys. Rev. D* **104** (2021) 015026 [[arXiv:2005.00008](#)] [[INSPIRE](#)].
- [50] C.W. Murphy, *Dimension-8 operators in the Standard Model Effective Field Theory*, *JHEP* **10** (2020) 174 [[arXiv:2005.00059](#)] [[INSPIRE](#)].
- [51] J. Davighi, S. Melville and T. You, *Natural selection rules: new positivity bounds for massive spinning particles*, *JHEP* **02** (2022) 167 [[arXiv:2108.06334](#)] [[INSPIRE](#)].
- [52] NNPDF collaboration, *Parton distributions for the LHC Run II*, *JHEP* **04** (2015) 040 [[arXiv:1410.8849](#)] [[INSPIRE](#)].
- [53] M.E. Peskin and T. Takeuchi, *Estimation of oblique electroweak corrections*, *Phys. Rev. D* **46** (1992) 381 [[INSPIRE](#)].
- [54] G. Altarelli, R. Barbieri and S. Jadach, *Toward a model independent analysis of electroweak data*, *Nucl. Phys. B* **369** (1992) 3 [Erratum *ibid.* **376** (1992) 444] [[INSPIRE](#)].
- [55] R. Barbieri, A. Pomarol, R. Rattazzi and A. Strumia, *Electroweak symmetry breaking after LEP-1 and LEP-2*, *Nucl. Phys. B* **703** (2004) 127 [[hep-ph/0405040](#)] [[INSPIRE](#)].
- [56] M. Farina, G. Panico, D. Pappadopulo, J.T. Ruderman, R. Torre and A. Wulzer, *Energy helps accuracy: electroweak precision tests at hadron colliders*, *Phys. Lett. B* **772** (2017) 210 [[arXiv:1609.08157](#)] [[INSPIRE](#)].
- [57] G. Panico, L. Ricci and A. Wulzer, *High-energy EFT probes with fully differential Drell-Yan measurements*, *JHEP* **07** (2021) 086 [[arXiv:2103.10532](#)] [[INSPIRE](#)].
- [58] LHCb collaboration, *Test of lepton universality using  $B^+ \rightarrow K^+ \ell^+ \ell^-$  decays*, *Phys. Rev. Lett.* **113** (2014) 151601 [[arXiv:1406.6482](#)] [[INSPIRE](#)].
- [59] LHCb collaboration, *Test of lepton universality with  $B^0 \rightarrow K^{*0} \ell^+ \ell^-$  decays*, *JHEP* **08** (2017) 055 [[arXiv:1705.05802](#)] [[INSPIRE](#)].

- [60] LHCb collaboration, *Search for lepton-universality violation in  $B^+ \rightarrow K^+ \ell^+ \ell^-$  decays*, *Phys. Rev. Lett.* **122** (2019) 191801 [[arXiv:1903.09252](#)] [[INSPIRE](#)].
- [61] F. Faura, S. Iranipour, E.R. Nocera, J. Rojo and M. Ubiali, *The Strangest Proton?*, *Eur. Phys. J. C* **80** (2020) 1168 [[arXiv:2009.00014](#)] [[INSPIRE](#)].
- [62] ATLAS collaboration, *Measurement of the high-mass Drell-Yan differential cross-section in  $pp$  collisions at  $\sqrt{s} = 7$  TeV with the ATLAS detector*, *Phys. Lett. B* **725** (2013) 223 [[arXiv:1305.4192](#)] [[INSPIRE](#)].
- [63] ATLAS collaboration, *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](#)] [[INSPIRE](#)].
- [64] CMS collaboration, *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](#)] [[INSPIRE](#)].
- [65] CMS collaboration, *Measurements of differential and double-differential Drell-Yan cross sections in proton-proton collisions at 8 TeV*, *Eur. Phys. J. C* **75** (2015) 147 [[arXiv:1412.1115](#)] [[INSPIRE](#)].
- [66] CMS collaboration, *Measurement of the differential Drell-Yan cross section in proton-proton collisions at  $\sqrt{s} = 13$  TeV*, *JHEP* **12** (2019) 059 [[arXiv:1812.10529](#)] [[INSPIRE](#)].
- [67] R.A. Khalek, S. Bailey, J. Gao, L. Harland-Lang and J. Rojo, *Towards Ultimate Parton Distributions at the High-Luminosity LHC*, *Eur. Phys. J. C* **78** (2018) 962 [[arXiv:1810.03639](#)] [[INSPIRE](#)].
- [68] R. Frederix, S. Frixione, V. Hirschi, D. Pagani, H.S. Shao and M. Zaro, *The automation of next-to-leading order electroweak calculations*, *JHEP* **07** (2018) 185 [Erratum *JHEP* **11** (2021) 085] [[arXiv:1804.10017](#)] [[INSPIRE](#)].
- [69] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [70] Y. Nesterov, *A method for solving the convex programming problem with convergence rate  $\mathcal{O}(1/k^2)$* , *Proc. USSR Acad. Sci.* **269** (1983) 543.
- [71] N.P. Hartland et al., *A Monte Carlo global analysis of the Standard Model Effective Field Theory: the top quark sector*, *JHEP* **04** (2019) 100 [[arXiv:1901.05965](#)] [[INSPIRE](#)].
- [72] L. Del Debbio, T. Giani and M. Wilson, *Bayesian approach to inverse problems: an application to NNPDF closure testing*, *Eur. Phys. J. C* **82** (2022) 330 [[arXiv:2111.05787](#)] [[INSPIRE](#)].
- [73] S. Iranipour et al., *A simultaneous determination of PDFs and  $\alpha_s$  using deep-learning, in preparation.*
- [74] V. Bertone, S. Carrazza and J. Rojo, *APFEL: A PDF Evolution Library with QED corrections*, *Comput. Phys. Commun.* **185** (2014) 1647 [[arXiv:1310.1394](#)] [[INSPIRE](#)].