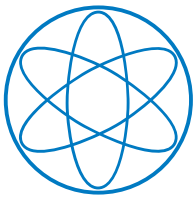

Technische Universität Müchen
TUM School of Natural Sciences

NNLO QCD calculations matched to parton shower for diboson processes with MiNNLO$_{\mathrm{PS}}$

Daniele Lombardi

Vollständiger Abdruck der von der TUM School of Natural Sciences der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz:

- Prof. Dr. Lukas Heinrich

Prüfende der Dissertation:

1. Prof. Dr. Andreas Weiler
2. Prof. Dr. Giulia Zanderighi

Die Dissertation wurde am 01.08.2022 bei der Technischen Universität München eingereicht und durch die TUM School of Natural Sciences am 21.10.2022 angenommen.

To my parents,
Gianni and Rosa,
and to their inspiring
perseverance in love.

Abstract

Increasing the accuracy of theory simulations for high-energy particle collisions is indispensable to take full advantage of the vast amount of data that has been and will continue to be collected at the LHC. Any deviation emerging from the comparison of theoretical predictions based on the SM of particle physics and experimental measurements might potentially be a signal of new and unexplored physics.

A simulation of a LHC collision event relies on many theoretical tools, among others fixed-order calculations and parton showers. The former is particularly suitable for the description of the hard core of the events, while the latter is more appropriate when multiple soft and/or collinear emissions from hard down to non-perturbative scales are involved. The complementary nature of the two methods makes their combination extremely important to obtain realistic theoretical predictions. Nowadays, the frontier of simulation accuracy is represented by the possibility to interface NNLO QCD predictions to parton showers.

We present the MiNNLO$_{\text{PS}}$ method, an extension of the MiNLO$'$ merging algorithm, which allows to reach NNLO QCD accuracy for physics observables inclusive over QCD radiation. Moreover, NLO and LO accuracy can also be claimed for observables involving one or two hard jets, respectively. All that is done without spoiling the logarithmic accuracy of the parton shower, to which the MiNNLO$_{\text{PS}}$ predictions are interfaced, thanks to the Powheg matching approach. In particular, we describe the extension of the MiNNLO$_{\text{PS}}$ framework to general colour-singlet production.

We apply our algorithm to diboson production processes, which are of fundamental importance in the LHC research programme. Predictions using the MiNNLO$_{\text{PS}}$ method are reported for $Z\gamma$ production, also including the effect of anomalous triple gauge couplings. For the relevant class of massive diboson processes, we show results for W^+W^- and ZZ production, where in the latter case the gluon-gluon loop-induced channel is also included at NLO accuracy and properly interfaced to the parton shower.

Zusammenfassung

Um die enormen Datenmengen, die am LHC bisher gesammelt wurden und weiterhin gesammelt werden, voll auszunutzen, ist eine höhere Genauigkeit bei der Simulation von hochenergetischen Teilchenkollisionen unentbehrlich. Selbst kleinste Abweichungen, welche aus dem aus dem Vergleich zwischen den theoretischen Vorhersagen des Standardmodells (SM) der Teilchenphysik und den experimentellen Messungen hervorgehen, könnte ein Anzeichen für neue und unerforschte Physik liefern.

Die Simulation von Kollisionsereignissen am LHC beruht auf verschiedensten theoretischen Werkzeugen, dazu gehören die Berechnungen zu fester Ordnung in der Störungstheorie und die Modellierung durch sogenannte Partonenschauer. Erstere sind insbesondere dafür geeignet den harten Kern der LHC Ereignisse zu beschreiben, letztere sind dagegen notwendig, wenn mehrere energiearme und/oder kollineare Abstrahlungen von harten bis runter zu nicht-perturbativen Energieskalen berücksichtigt werden. Die komplementäre Natur dieser beiden Methoden macht ihre Kombination besonders wichtig um realistische theoretische Vorhersagen zu erhalten. Das derzeitig höchste Ziel um die Genauigkeit von LHC Simulationen zu verbessern liegt darin NNLO Vorhersagen in der perturbativen Quantenchromodynamik mit Partonschauern zu verbinden.

In dieser Arbeit, wird die sogenannte MiNNLO$_{\text{PS}}$ Methode präsentiert, welche eine Erweiterung des "MiNLO′ merging" Algorithmus darstellt. Diese Methode erlaubt es NNLO QCD Genauigkeit für physikalische Observablen zu erreichen, die inklusiv bezüglich zusätzlicher QCD Strahlung sind. Gleichzeitig, wird NLO bzw. LO Genauigkeit erlangt für Observablen, die auf jeweils ein bzw. zwei harten Jets beruhen. Dank der zugrundeliegenden Powheg Methodik zur Verbindung mit dem Partonenschauer, erfolgt dies ohne die logarithmische Genauigkeit des Partonenschauers zu beeinträchtigen, mit dem die MiNNLO$_{\text{PS}}$ Vorhersagen verbunden ist. Im Rahmen dieser Arbeit, werden wir insbesondere darauf eingehen wie sich MiNNLO$_{\text{PS}}$ Methode für die Produktion von generellen farbneutralen Endzuständen verallgemeinern lässt.

Unser Algorithmus wird auf die Vektorbosonpaarproduktion angewendet, welche von fundamentaler Bedeutung für das LHC-Forschungsprogramm ist. Es werden Vorhersagen mittels der MiNNLO$_{\text{PS}}$ Methode für die Produktion von $Z\gamma$-Endzuständen berechnet sowohl im SM als auch unter Berücksichtigung von Effekten durch anomalen Vektorbosonselbstkopplungen. Bezüglich der Produktion von zwei massiven Vektorbosonen zeigen wir Ergebnisse für die W^+W^--Produktion und die ZZ-Produktion, wobei im letzteren Fall wir ebenfalls den loop-induzierten Beitrag durch Gluonenfusion auf NLO-QCD-Genauigkeit mitnehmen und konsistent mit dem Partonschauer verbinden.

Acronyms

AP Altarelli-Parisi

aTGC Anomalous triple gauge coupling

BSM Beyond the Standard Model

CERN Conseil Européen pour la Recherche Nucléaire

CKM Cabibbo-Kobayashi-Maskawa

CM Center of mass

CMW Catani-Marchesini-Webber

CP Charge-parity

CS Catani-Seymour

DESY Deutsches Elektronen-Synchrotron

DIS Deep Inelastic Scattering

EFT Effective field theory

EW Electroweak

FKS Frixione-Kunst-Signer

FO Fixed-order

GPMC General purpose Monte Carlo

IR Infrared

LEP Large Electron–Positron Collider

LHC Large Hadron Collider

LHE Les Houches Event

LL Leading logarithmic

LO Leading order

MC Monte Carlo

MPI Multi-particle interactions

NLL Next-to-leading logarithmic

NLO Next-to-leading order

NNLL Next-to-next-to-leading logarithmic

NNLO Next-to-next-to-leading order

PDF Parton distribution function

PDG Particle Data Group

PS Parton shower

QCD Quantum chromodynamics

QED Quantum electrodynamics

QFT Quantum field theory

RGE Renormalization group equation

SLAC Stanford Linear Accelerator Center

SM Standard Model

UV Ultraviolet

VEV Vacuum expectation value

Contents

List of publications

The bulk of this work is based on the following list of publications by the author, where the results discussed in Chapter 4 of this Thesis were originally presented.

- D. Lombardi, M. Wiesemann and G. Zanderighi, *Anomalous couplings in $Z\gamma$ events at NNLO+PS and improving $\nu\bar{\nu}\gamma$ backgrounds in dark-matter searches* Phys.Lett.B 824 (2022) [arXiv:2108.11315], 25 August 2021.

- L. Buonocore, G. Koole, D. Lombardi, L. Rottoli, M. Wiesemann and G. Zanderighi, *ZZ production at nNNLO+PS with MiNNLO$_{PS}$* JHEP 01 (2022) 072 [arXiv:2108.05337], 11 August 2021.

- D. Lombardi, M. Wiesemann and G. Zanderighi, *W^+W^- production at NNLO+PS with MiNNLO$_{PS}$* JHEP 11 (2021) 230 [arXiv:2103.12077], 22 March 2021.

- D. Lombardi, M. Wiesemann and G. Zanderighi, *Advancing MiNNLO$_{PS}$ to diboson processes: $Z\gamma$ production at NNLO+PS* JHEP 06 (2021) 188 [arXiv:2010.10478], 20 October 2020.

1 Introduction

Developed in the second-half of the twentieth century through the cooperative effort of many physicists all over the world, the Standard Model (SM) of particle physics is one of the most comprehensive and predictive theory that has ever been conceived. Since it was finalized in the mid-1970s, the SM has received continuous confirmations from many experiments and no compelling deviation from its predictions has been found so far.

1.1 Basic facts of the Standard Model

The SM describes the elementary constituents of nature and their interactions in a Quantum field theory (QFT) language. Particles comprising the entire matter of the known universe are associated to quanta of fermionic (spin 1/2) relativistic quantum fields, while carriers of the three fundamental forces accounted for by the SM (the electromagnetic, weak and strong ones) are related to bosonic (spin 1) fields. Finally, the Higgs boson, a scalar (spin 0) field named after one of the theoretical physicists who first proposed its existence, plays the special role of providing masses to all SM massive particles via a spontaneous symmetry breaking mechanism. These elementary degrees of freedom enter one Lorentz invariant, renormalizable [1] Lagrangian which is symmetric under the $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge group. This Lagrangian, together with ninteeen experimental inputs, which are free parameters of the theory, can be used to describe with astounishing accuracy a vast amount of the physics we observe.

Each of the force-carrier fields belongs to the adjoint representation of one of the three simple, compact subgroups of the full SM gauge group. The massless gluon field is the vector field which mediates the strong force and belongs to the adjoint representation of $SU(3)_c$, the group of colour symmetry. We denote this field by G^μ_A, where μ and $A = \{1, \ldots, 8\}$ are a Lorentz index and the colour index of the adjoint representation, respectively. All fermionic matter fields which transform under $SU(3)_c$ belong to its fundamental representation and are endowed with a colour quantum number: they are spinor fields named quarks q_α (where $\alpha = \{1, 2, 3\}$ is the colour index of the fundamental representation). Quantum chromodynamics (QCD) is that branch of the SM describing strong interactions according to the QCD Lagrangian [1] (we adopt the choice of natural

[1] For a brief recap of the concept of renormalizability in QFT, see discussion in Section 2.1.1 of this work.

units $\hbar = c = 1$ and the Einstein convention for repeated indices) [2]:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G^A_{\mu\nu} G^{\mu\nu}_A + \sum_q \bar{q}_\alpha i\gamma_\mu \mathcal{D}^\mu_{\alpha\beta} q_\beta + \frac{g_S^2}{32\pi^2} \theta_{\text{QCD}} \tilde{G}^B_{\mu\nu} G^{\mu\nu}_B \,. \tag{1.1}$$

The first term of Eq. (1.1) describes the propagation and self-interactions of the eight gluon fields via the gluon field strength tensor:

$$G^A_{\mu\nu} = \partial_\mu G^A_\nu - \partial_\nu G^A_\mu - g_S f^{ABC} G^B_\mu G^C_\nu \,, \tag{1.2}$$

where g_S (often appearing in the combination $\alpha_s = g_S^2/(4\pi)$) is the strong coupling constant. If t^A are the generators of the $SU(3)_c$, f^{ABC} are the structure constants of the associated Lie algebra, i.e $[t^A, t^B] = i f^{ABC} t^C$. The second term of Eq. (1.1) accounts for the propagation and interactions of the quark fields, where γ_μ are the four Dirac gamma matrices, $\bar{q} = q^\dagger \gamma^0$ is the adjoint Dirac field and $\mathcal{D}^\mu_{\alpha\beta}$ the covariant derivative, defined as:

$$\mathcal{D}^\mu_{\alpha\beta} = \partial^\mu \delta_{\alpha\beta} + i g_S t^A_{\alpha\beta} G^A_\mu \,. \tag{1.3}$$

In the second term of the QCD Lagrangian, the sum is carried out over the six quark species of the SM, each of them coming in three colours. Specifically, we have $q = u, d, s, c, b, t$, denoting the up, down, strange, charm, bottom and top quarks, respectively. The generators t^A obey three relevant colour-algebra relations: $t^A_{\alpha\gamma} t^A_{\gamma\beta} = C_F \delta_{\alpha\beta}$, $f_{ECD} f_{BCD} = C_A \delta_{EB}$ and $t^A_{\alpha\beta} t^B_{\alpha\beta} = T_R \delta_{AB}$, where $C_F = 4/3$, $C_A = 3$ and $T_R = 1/2$ are the colour-factors respectively associated with the emission of a gluon from a quark/antiquark, the emission of a gluon from another gluon and the splitting of a gluon into a quark-antiquark pair [2]. Finally, the third term of Eq. (1.1), also dubbed the θ-term $\mathcal{L}_\theta$, is another gauge invariant renormalizable contribution that can be added to the QCD Lagrangian, where $\tilde{G}^A_{\mu\nu} = \epsilon^{\mu\nu\rho\delta} G^A_{\rho\delta}/2$ is the dual gluon field strength tensor and $\epsilon^{\mu\nu\rho\delta}$ the completely antisymmetric Levi-Civita tensor. Since it can be written in terms of a total derivative, it has no impact on theoretical calculations based on a perturbative approach (see Section 2.1), but it does have non-vanishing physical effects at the non-perturbative level, due to the non-trivial topological structure of the QCD vacuum. One of these physical effects would be the violation of the charge and parity (CP) discrete symmetry in strong interactions by an amount dictated by the value of θ_{QCD}. Anyway, experimental measurements on the neutron electric dipole moment constrain its value to be surprisingly small ($\theta_{\text{QCD}} < 10^{-10}$ [3]), which poses serious fine-tuning issues (something usually referred to as *strong CP problem*, discussed at length in the literature and more recently in Ref. [4]).

Overall, the QCD Lagrangian of the SM comes with two free parameters, whose values are not predicted by the theory: g_S and θ_{QCD}.

[2] The QCD Lagrangian in Eq. (1.1) should be completed with a mass term for the quarks. Since all particle masses originate from the Higgs sector within the SM, we omit such a term here and we account for this contribution in Eq. (1.15), presented later in our discussion. In this brief summary we also refrain from explicitly including *gauge-fixing* or *ghost* terms, which are indispensable in QCD perturbation theory and would deserve a dedicated discussion on their own.

Three massless vector bosons W_μ^I (with $I = \{1, 2, 3\}$) belong to the adjont representation of weak isospin $SU(2)_L$ and are the force-carriers of the weak force, while a single massless vector field B_μ is associated to the hypercharge symmetry group $U(1)_Y$. The $SU(2)_L$ group is chiral in the SM, since the left- and right-handed components of the Dirac matter fields transform according to different group representations. The quark fields are both colourful and weakly charged, and their left-handed components are grouped in three generations which transform like $SU(2)_L$ doublets:

$$Q = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}. \tag{1.4}$$

The six remaining right-handed fields q_R are instead $SU(2)_L$ singlets. Together with quarks, six lepton fields ℓ contribute to the matter content of the universe, where $\ell = e, \mu, \tau, \nu_e, \nu_\mu, \nu_\tau$ stand for the electron, the muon, the tau and the three corresponding neutrinos, respectively. Differently from quarks, leptons are $SU(3)_c$ singlets, but weakly charged. As with the quark fields, their left-handed components come in three $SU(2)_L$ doublets:

$$L = \begin{pmatrix} \nu_{e,\,L} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,\,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,\,L} \\ \tau_L \end{pmatrix}, \tag{1.5}$$

and the right-handed ones ℓ_R are weak-isospin singlets. The propagation of these degrees of freedom and their weak/hypercharge interactions are described by the electroweak (EW) sector of the SM Lagrangian [1]:

$$\begin{aligned} \mathcal{L}_{\text{EW}} = &-\frac{1}{4} W^I_{\mu\nu} W_I^{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + \sum_L \bar{L}_i i\gamma_\mu (\partial^\mu \delta_{ij} + i g_W \tau^I_{ij} W_I^\mu + i g' \delta_{ij} Y_L^\ell B^\mu) L_j \\ &+ \sum_\ell \bar{\ell}_R i\gamma_\mu (\partial^\mu + i g' Y_R^\ell B^\mu) \ell_R + \sum_Q \bar{Q}_i i\gamma_\mu (i g_W \tau^I_{ij} W_I^\mu + i g' \delta_{ij} Y_L^q B^\mu) Q_j \\ &+ \sum_q \bar{q}_R i\gamma_\mu (i g' Y_R^q B^\mu) q_R \,. \end{aligned} \tag{1.6}$$

The first two terms of the previous equation are written using the field strength tensor $W^I_{\mu\nu}$ for the weak isospin group and $B_{\mu\nu}$ for the hypercharge group. They are defined as:

$$W^I_{\mu\nu} = \partial_\mu W^I_\nu - \partial_\nu W^I_\mu - g_W \epsilon^{IJK} W^J_\mu W^K_\nu \qquad \text{and} \qquad B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \,, \tag{1.7}$$

where g_W is the EW gauge coupling and ϵ^{IJK} are the structure constants of the $SU(2)_L$ Lie algebra with generators τ^I, i.e $[\tau^I, \tau^J] = \epsilon^{IJK} \tau^K$. We notice that, being $U(1)_Y$ an abelian group, no self-interaction terms are generated by $B_{\mu\nu}$. The remaining terms of Eq. (1.6) account for the lepton and quark electroweak interactions: weak interactions are proportional to g_W, while hypercharge ones are constructed using the $U(1)_Y$ coupling g' weighted by the different matter field hypercharges $Y = \{Y_L^\ell, Y_R^\ell, Y_L^q, Y_R^q\}$. Clearly, the indices i, j vary over the two components of the $SU(2)_L$ fundamental representation. We stress that we have not reported the derivative term for the quark fields in Eq. (1.6),

since it was already included in the QCD part of the Lagrangian in Eq. (1.1), so that, by adding up the two equations the full covariant derivative for the quark fields is correctly recovered (being the quark EW interactions colour diagonal).

Therefore, the electroweak Lagragian adds to the theory two free parameters, namely g_W and g'.

The SM Lagrangian is then completed by adding the terms involving the Higgs field, which is given by a complex $SU(2)_L$ doublet:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} , \tag{1.8}$$

and enters the Higgs sector of the SM Lagrangian [1]:

$$\begin{aligned} \mathcal{L}_{\text{Higgs}} = &\left(\partial^\mu \phi_i^\dagger + i g_W \tau^I_{ki} W^\mu_I \phi^\dagger_k + i \frac{g'}{2} B^\mu \phi_i^\dagger \right) \left(\partial_\mu \phi_i - i g_W \tau^I_{ij} W^I_\mu \phi_j - i \frac{g'}{2} B_\mu \phi_i \right) \\ &- \mathcal{V}(\phi^\dagger \phi) + \mathcal{L}_{\text{yuk}} , \end{aligned} \tag{1.9}$$

where the first line describes the propagation and the electroweak interactions of the Higgs field. The term $\mathcal{V}(\phi^\dagger\phi)$ denotes the Higgs potential, which is quadratic in the field ϕ and plays a crucial role in the spontaneous symmetry breaking mechanism [5]:

$$\mathcal{V}(\phi^\dagger \phi) = \lambda (\phi^\dagger \phi)^2 - \mu^2 \phi^\dagger \phi . \tag{1.10}$$

Since λ and μ^2 are two positive definite parameters, the function in Eq. (1.10) has a set of non-zero degenerate minima at classical level which are located on the circle $|\phi|^2 \equiv v^2 = \mu^2/\lambda$, where v is the Higgs vacuum expectation value (VEV). Even if $\mathcal{L}_{\text{Higgs}}$ is $SU(2)_L \times U(1)_Y$ invariant, the VEV of the theory:

$$\langle \phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \tag{1.11}$$

is not, and it breaks $SU(2)_L \times U(1)_Y$ to the residual electromagnetic symmetry group $U(1)_{\text{em}}$, associated to the only unbroken generator $Q_{ij} \equiv \tau_{3,ij} + Y \delta_{ij}$ (*Gell-Mann–Nishijima formula*). This generator is identified with the electric charge operator: the hypercharge values Y are indeed chosen in such a way to reproduce the measured electric charges of the SM particles (eigenvalues of Q_{ij}), once the transformation under the weak isospin gauge group is fixed ($\tau_{3,ij}$ has zero eigenvalues on right-handed matter fields). We can then parametrize the Higgs field fluctuations around the VEV of Eq. (1.11) as:

$$\phi = e^{i \tau_k \xi_k / v} \begin{pmatrix} 0 \\ (v + H)/\sqrt{2} \end{pmatrix} \tag{1.12}$$

in terms of three massless Goldstone modes ξ_k, which correspond to the broken generators of $SU(2)_L \times U(1)_Y$ according to the Goldstone theorem [6], and a massive field H. The latter has been identified with the neutral scalar field, first observed at the Large Hadron

Collider (LHC) experiment in 2012 [7, 8]. By making use of the gauge invariance of the SM Lagrangian, we can choose a specific gauge (unitary gauge) where the three Goldstone modes disappear from the Lagrangian, but effectively provide longitudinal components (and therefore masses) to three weak bosons: the two gauge bosons $W_\mu^{1/2}$ (or, better said, to the electric charge eigenstates $W_\mu^\pm = (W_\mu^1 \mp i W_\mu^2)/\sqrt{2}$, named $W^\pm$ bosons) and the gauge boson Z_μ (so called Z^0 or simply Z boson). The neutral gauge boson Z is associated to the third broken generator of $SU(2)_L \times U(1)_Y$ and is obtained from the transformation:

$$\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} . \tag{1.13}$$

Performing this rotation, the mass eigenstates of the theory are correctly defined. In particular, A_μ is the massless vector field of the $U(1)_{\rm em}$ group: its quanta are identified with photons, whose propagation and interactions are described by Quantum electrodynamics (QED). The angle θ_W (Weinberg angle) is fixed by the theory and given by $\sin^2\theta_W = g'^2/(g_W^2 + g'^2)$. Using Eq. (1.12) and Eq. (1.13) in Eq. (1.9) with a unitary gauge, the Higgs, $W^\pm$ and Z boson masses are predicted by the SM to be [3]:

$$m_H = \sqrt{2\lambda} v\,, \quad m_W = \frac{v}{2} g_W \quad \text{and} \quad m_Z = \frac{v}{2}\sqrt{g_W^2 + g'^2}\,. \tag{1.14}$$

Moreover, the Higgs trilinear and quartic self-couplings, together with its interactions with massive gauge bosons, are also fixed. Therefore, the first two terms of Eq. (1.9) discussed so far just require the inclusion of two more free parameters, which can be chosen to be, for instance, λ and v.

The last part of Eq. (1.9), i.e $\mathcal{L}_{\rm yuk}$, contains the interactions of the Higgs field ϕ with the matter fields:

$$\mathcal{L}_{\rm yuk} = -\,Y^e_{mn}\bar{L}^m_i \phi_i e^n_R - Y^d_{mn}\bar{Q}^m_i \phi_i d^n_R - Y^u_{mn}\bar{Q}^m_i (i\sigma_{2,ij}\phi^\star_j) u^n_R + h.c. \tag{1.15}$$

where σ_2 is the second Pauli matrix and *h.c.* stands for the complex conjugate of the previous part of the equation, required to have a real Lagrangian. In Eq. (1.15), we have also introduced the notation $e^n_R \in \{e_R, \mu_R, \tau_R\}$ (right-handed charged leptons), $u^n_R \in \{u_R, c_R, t_R\}$ (right-handed up-type quarks) and $d^n_R \in \{d_R, s_R, b_R\}$ (right-handed down-type quarks). Finally, Y^e_{mn}, Y^u_{mn} and Y^d_{mn} are complex 3×3 matrices (named Yukawa matrices). We notice that right-handed neutrinos have no interaction with SM particles, so that in practise they do not enter the SM theory. Once the Higgs field assumes a non-zero VEV, the first term of Eq. (1.15) can be diagonalized without any consequence for the other terms of the SM Lagrangian. That provides mass terms for the charged leptons, which depend on the three eigenvalues of Y^e_{mn} (Yukawa couplings), i.e y_e, y_μ, y_τ, not predicted by the theory. Neutrinos are massless in the SM. The last

[3] These mass values are clearly the tree-level ones: more accurate predictions for these parameters can be obtained by including perturbative corrections. In this respect, the instability under perturbative corrections of the SM prediction for the Higgs mass m_H is at the root of the well-known *hierarchy problem* (see Section 1.3).

two terms of $\mathcal{L}_{\text{yuk}}$ give masses to the six quark fields, where two terms are needed to separately provide masses to the down and up components. One needs unitary matrices acting differently on the right- and left-handed fields and on the up- and down-type components to move from the flavour to the mass basis of the quark fields (namely, the complex matrices Y^d_{mn} and Y^u_{mn} are diagonalized by two bi-unitary transformations). The six Yukawa eigenvalues of Y^u_{mn} and Y^d_{mn} resulting from this procedure are additional free parameters which define the quark masses. But since the $W^{\pm}$ boson interactions (see Eq. (1.6)) mix the up and down components, these interactions written in terms of massive quark eigenstates will be modified by a 3×3 matrix V_{CKM}, named the Cabibbo-Kobayashi-Maskawa (CKM) matrix [9]: its effect is to allow for interactions that mix quark generations of Eq. (1.4) with each others. As mentioned above, a similar term is missing for the lepton sector, since neutrinos are massless and one has some residual freedom to rotate a generation-mixing matrix away. Since V_{CKM} can be parametrized by three real angles θ_{12}, θ_{13}, θ_{23} and a complex phase δ, that brings the number of SM free parameters to ninteen.

1.2 Experimental evidences of particle physics

The path that brought to the formulation of the SM as it is commonly accepted nowadays is one of the most exciting example of how theory and experimental measurements can reinforce and support each other. As it is clear from a quantum field theory perspective, higher and higher energies need to be probed to explore the elementary constituents of nature and, to this end, particle accelerators turned out to be fundamental tools.

Before 1974, the three fundamental forces described by the SM were postulated in order to account for observed phenomena, such as electromagnetism, the radioactive beta decay (first explained by Fermi in 1933 in terms of a new 4-fermion contact interaction [10]) and the existence of nuclei, i.e bound states of neutrons and protons. But at that time, only the photon was known to exist. Among leptons, the electron and the muon (the latter discovered by Carl D. Anderson and Seth Neddermeyer in 1936 studying cosmic radiation [11]) were known, together with neutrinos, whose existence was first postulated by Pauli in 1930 [12], still in the context of beta decays. The idea of quark was first introduced by Murray Gell-Mann and George Zweig in 1964 [13] to explain the approximate $SU(3)$ flavour symmetry in the spectrum of low-mass hadrons (mesons and baryons). It soon became clear that three different elementary particles (the up, down and strange quarks) were needed, having spin 1/2 to accomodate the observed meson and baryon spins. Moreover, quarks had to be endowed with a new quantum number, named colour, with three possible values to reconcile Fermi-Dirac statistics with the existence of fully symmetric baryon states like the Δ^{++} baryon. The introduction of the colour hypothesis had to be supplemented by the *confinement* assumption, according to which only colour-singlet states can exist as asymptotic states, to justify the limited amount of measured hadron species.

Results from electron-proton Deep Inelastic Scattering (DIS) experiments carried out in California at the Stanford Linear Accelerator Center (SLAC) in 1968-69 first

supported the idea that protons comprise of point-like constituents (as confirmed by the scaling behaviour of proton structure functions) of half-integer spin (suggested by the experimental verification of the Callan-Gross relation [14]). An important step sheding further light on experimental observations was the theoretical discovery of asymptotic freedom in 1973 by David Gross, Frank Wilczek [15] and David Politzer [16]: that allowed to explain why, despite the confinement hypothesis, quarks within the proton behaved as free particles at very high energy. The running behaviour of the strong coupling constant, connecting the confined and the asymptotically free regimes of QCD, was then confirmed by many data collected at different energy scales by numerous collaborations over the years. Once the quark masses are neglected, α_s is the only free parameter of QCD and its accurate determination is still a crucial task (see Ref. [17] for a recent review). Its value is generally extracted by a proper combination of measurements from different fields, such as hadronic τ and bottonium decay measurements, together with studies on various processes recorded at DIS experiments, electron-positron and hadron colliders. Hints of the underlying $SU(3)_c$ symmetry of QCD, with three as a number of colours, came from different measurements, such as the one of the neutral pion decay rate into two photons or of the ratio R, i.e the ratio of the e^+e^- total hadronic cross section to the cross section for $\mu^+\mu^-$ production. The latter originates from data collected over the years in electron-positron colliders from various experiments operating at different energy scales [18]. Specifically, R measurements in low energy regions were performed from 1977 to 1999 by the DASP and PLUTO experiments of DORIS accelerator at the Deutsches Elektronen-Synchrotron (DESY) in Hamburg, at DCI in Orsay, at Adone in Frascati, with the SPEAR storage ring at SLAC laboratories, at Novosibirsk and at Beijing; data within the charm and bottom thesholds were collected again at DORIS (also by the LENA detector), by the Crystal Ball detector at SPEAR, by Mark I (also known as SLAC-LBL magnetic detector) and CLEO and CUSB detectors at Cornell Electron Storage Ring (CESR); above the bottom threshold measurements were done by three more e^+e^- colliders, namely the Positron-Electron Project (PEP) at SLAC, the Positron-Elektron-Tandem-Ring-Anlage (PETRA) at DESY, and the Large Electron–Positron Collider (LEP) at CERN in Geneva. Further evidences that QCD was a meaningful description of nature came from the three-jet event measurement [19], first achieved by the four different collaborations at the PETRA e^+e^- accelerator (1978-86), which is customarily regarded as an experimental proof of the gluon existence. Event shape studies carried out at LEP confirmed the vectorial nature of gluons, while the non-abelian structure of QCD resulted from evidences of the triple-gluon coupling, highlighted by angular correlation measurements on four-jet events (for instance the ones carried out by the L3 collaboration at LEP on the famous Bengtsson-Zerwas $\chi_{\rm BS}$ angle [20]): these studies were also extremely useful to extract the QCD colour factors C_A and C_F.

Alongside with QCD researches, the indipendent consolidation of the theory of electroweak interactions completed the path towards the structuring of the SM. A decisive step was done in 1956 by Chien-Shiung Wu with her experiment establishing the parity-violating nature of weak interactions [21]: that made clear that weak interactions had to be described by a chiral theory. Theoretical studies carried out from 1961 to 1967 by Sheldon Glashow [22], Abdus Salam [23] and Stephen Weinberg [24] led to the

unified description of electromagnetic and weak interactions in the electroweak theory, whose renormalizability was proven in 1971 by Martinus J. G. Veltman and Gerard 't Hooft [25, 26]. Many observations started to shed light on the correctness of the theory, but two of them were absolutely crucial for this purpose. First, in 1973 neutral currents were discovered at the Gargamelle detector, located at CERN, which was the first indirect evidence of the existence of the neutral Z boson [27]. Then, in 1893 the UA1 and UA2 experiments at the CERN proton-antiproton collider ($Sp\bar{p}S$) detected the $W^{\pm}$ [28, 29] and Z [30] bosons. Between these two important discoveries, the particle content of the emerging SM kept on being enriched. In November 1974 (the so called *November revolution*) the Brookhaven National Laboratory [31], in Long Island, and SLAC SPEAR accelerator [32] discovered the J/ψ meson, the first evidence of the charm quark, whose existence was proposed in 1970 in the context of the S. Glashow, J. Iliopoulos and L. Maiani (GIM) mechanism [33] to explain why flavour-changing neutral currents (FCNC) are suppressed. The observation of CP violation in the kaon system [34, 35] also required the introduction of a third quark generation: this was understood in 1973 by M. Kobayashi and T. Maskawa as the only viable way for the CKM matrix to have a complex CP-violating phase [9]. Shortly after, the idea was corroborated by the discovery of the bottonium Υ ($\bar{b}b$ ground state) in 1977 at Fermilab in Illinois [36], precedeed by the τ lepton discovery in 1975 at SLAC [37]. In order to have further confirmations of CP violation, the scientific community had to wait 2001 with the B-meson decay observation by the Belle experiment at KEK in Japan [38] and the BaBar experiment at SLAC [39], which strongly supported the theory of Kobayashi and Maskawa. Once the bottom quark was found, many new electron-positron machines were designed to find the top quark, as well: PETRA at DESY, TRISTAN in Japan, PEP and SLC at SLAC, and LEP at CERN. But due to a mass higher than the energy reach of any of these accelerators, the top quark would have been discovered much later in 1994 at the Tevatron proton-antiproton collider of Fermilab [40] (even if already in 1993 some indirect determinations of the top mass were performed at LEP and at SLC). Despite that, electron-positron colliders provided a very clean environment where more and more accurate measurements of the electroweak parameters could be performed. The most important ones are predictions for the number of SM neutrinos (reinforced by the discovery of the tauonic neutrino in 2000 at Fermilab with the DONUT experiment [41]) and measurements for the $W^{\pm}$ and Z boson masses and decay width, together with the electroweak mixing angle $\sin^2\theta_W$.

The last missing pillar of the SM was the Higgs boson. Even though some constraints on its mass had been obtained by some e^+e^- collider analysis, none of these machines was able to really detect the Higgs boson. It was known that vector boson scatterings at high energy would violate unitarity without the Higgs field: that knowledge allowed to derive a bound (the Lee-Quigg-Thacker bound) on the Higgs mass, i.e $m_H \lesssim 1\,\text{TeV}$ [42]. Consequently, a machine capable of reaching such energy regimes would have either discovered the Higgs or shown that weak interactions become non-perturbative at these energy scales. This and other arguments strongly motivated the construction of one of the world's largest and most powerful particle accelerator: the LHC, built at CERN in the LEP tunnel after its shutdown. Inside a 27-kilometre ring structure two proton beams are accelerated and made collide at four different locations around the accelerator

ring, corresponding to the positions of four particle detectors: ATLAS, CMS, ALICE and LHCb. In 2009 a beam collision energy $\sqrt{s}$ of 900 GeV was reached, followed by an energy reach of $\sqrt{s} = 7$ TeV in 2010 and of $\sqrt{s} = 8$ TeV in 2012, the year when the Higgs boson discovery was announced. Its discovery [7, 8] was the result of a combined analysis of the $H \to ZZ^* \to 4l$, $H \to \gamma\gamma$ and $H \to WW^* \to e\nu_e\mu\nu_\mu$ decay channels. If LHC Run 1, from 2009 to 2012, essentially allowed to probe the Higgs couplings to gauge bosons, LHC Run 2, operating at $\sqrt{s} = 13$ TeV from 2015 to 2018, first made it possible to study the Higgs Yukawa couplings to the third generation of charged fermions, namely the top and bottom quarks and the τ lepton. Despite these important achievements, many aspects of the SM Higgs sector still need to be established, such as the Higgs couplings to the second and first fermion generations or the Higgs self-couplings, which would shed light on the structure of the Higgs potential. Moreover, many cross section and parameter measurements have been performed at the LHC, whose precision is getting more and more competitive to previous measurements carried out at electron-positron colliders: the most representative one is the $W^\pm$ mass measurement, whose accuracy has reached the 0.02% level [2]. All that definitely changed a long-standing paradigm, according to which e^+e^- colliders were considered precision machines, due to their clean environment, as opposed to hadron colliders, mainly seen as discovery machines where higher energies could be probed. At the LHC it soon became clear that precision physics could be carried out, as well.

1.3 Precise phenomenology: a path towards new physics

The great amount of data collected during LHC operational time so far strongly supported the SM field theory and did not show any significant deviation from its predictions. But some physics beyond the SM is expected to exist, since the SM itself, as it is, is not a complete theory. There are many phenomena that are not accounted for by the SM Lagrangian, among which the most striking of all is the gravitational force. Moreover, cosmological observations suggest that visible matter is only roughly 5% of the whole energy content of our universe. As also supported by many gravitational observations (among which one of the first dates back to Zwicky's work of 1933 [43]), a 26% of it should comprise of dark matter, but the SM does not contain any particle which would be a good dark matter candidate. The remaining 69% is expected to be dark energy, whose idea first appeared with Einstein's cosmological constant and then seriously entered cosmology with the first inflationary models [44]. But dark energy, which can be thought of as the vacuum energy of a quantum field, is again difficult to accomodate within the SM theory. Another problem that the SM can not fully explain is the matter-antimatter asymmetry of the universe: even though all of the three Sakharov conditions [45] for baryogenesis are satisfied by the SM, they are quantitatively insufficient.

Together with cosmological arguments, also experimental observations from particle physics soon started to point to the fact that the SM can not be our ultimate theory. Even if SM neutrinos are massless left-handed particles, the observation of neutrino oscillations in different contexts made it clear that neutrinos do have masses. However,

explaining the mechanism that provides masses to neutrinos is still an open problem. Moreover, among many other questions in neutrino physics which remain unsolved, a very fundamental one is whether neutrinos should be described in terms of a Dirac field or whether they are their own antiparticles, and so Majorana-like fields, as it would emerge from the observation of neutrinoless double-beta decays. For a review on neutrino physics, see for instance Ref. [46]. But hints of new physics also seem to come from esperiments on B meson decays, which might challenge the SM idea that lepton couplings to EW gauge bosons are universal: in 2012 the BaBar experiment at PEP-II collider (at SLAC) reported small deviations of 3.4σ [47] from SM expectations, materializing in a small excess of the $\bar{B} \to D^{(\star)}\tau^-\bar{\nu}_\tau$ decay rate. Measurements going in the same direction (even though with a smaller statistical significance) were also registered by LHCb [48] at CERN and by the Belle experiment at KEKB collider (in Japan) [49] in 2015. Last year (2021), some results from Fermilab on the anomalous magnetic dipole moment of the muon seemed to show deviations from SM predictions with a 4.2σ significance [50], while this year a CDF analysis based on Tevatron data reported a $W^\pm$ mass measurement exceeding the SM value with a surprisingly high 7σ significance [51]. If only some of these results were confirmed, that might disclose a path towards physics beyond the SM (BSM).

Another set of arguments providing reasons for the SM not to be the final theory of nature simply arises from theoretical considerations. For instance, the SM comes with 19 free parameters, which means it is unable to explain the origin of their values: the strong hierarchy among fermion masses or the stringent bounds on θ_{QCD} (the already mentioned *strong CP problem*) remain unexplained. Even if the anomaly cancellations, required not to break gauge symmetries at quantum level, force the electric charge to be quantized, no really deep explanation of why particle charges should only come in fractions of the elementary electron charge is given in the SM. Finally, even though the presence of the Higgs boson is a key ingredient of the SM, the Higgs mass is not stable under radiative corrections and the m_H^2 parameter is quadratically sensitive to scales of new physics. Since the measured Higgs mass is of the order of the EW scale, only large cancellations taking place with an unnatural fine-tuning of the Lagrangian parameters can solve this issue within the SM (the so called *hierarchy problem*).

This short and necessarily incomplete summary strongly motivates to look for BSM physics. As we have briefly shown, particle colliders are extremely precious tools to explore the elementary constituents of nature, and particularly nowadays the LHC still has a prominent role in this context. Despite the fact that, after the Higgs discovery, no new particle predicted by many theoretical extensions of the SM has been directly found, new physics can still be discovered in an indirect way: any deviation from the comparison of precise experimental measurements and accurate theoretical predictions based on the SM can be a signal of new physics. As we have already mentioned, LHC measurements are getting more and more accurate and they are bound to improve even further with the high-luminosity (HL) upgrade of the LHC, which is expected to be finalized towards the end of 2027. To avoid that the theory uncertainties become the limiting factor of this data-theory comparison, a similar effort in improving the accuracy of the theoretical predictions is demanded from the theory community. Unfortunately, that is by no means

a simple task, due to the complexity involved in theoretical simulations of LHC events. Indeed, even a single LHC collision comes with a huge amount of particles and necessarily requires to deal with many different energy scales at the same time. Furthermore, it is clear that the composite structure of the colliding objects, i.e the protons, and the non-abelian nature of QCD, which leads to colour confinement at scales of the order of $\Lambda_{\rm QCD} \sim 1\,\text{GeV}$, significantly complicate the pattern of a LHC event.

Due to efficiency reasons, event simulations are built around the hard event taking place at a high-energy scale $Q \gg \Lambda_{\rm QCD}$ characteristic of the process that one wants to accurately describe. Thanks to asymptotic freedom, at such high-energy scales α_s is small enough to legitimate a perturbative approach (see Section 2.1): the cross section of the process can be computed with a well-established machinery starting from the Feynman rules which can be read off the QCD Lagrangian of Eq. (1.1), supplemented with a proper gauge-fixing term and ghost contributions. These fixed-order calculations can predict with an increasing accuracy physical observables by progressively adding one more perturbative contribution, which is suppressed by an extra power of α_s with respect to the lower order one. At scales of the order of $\Lambda_{\rm QCD}$, the perturbative picture breaks down and hadronization models are needed to describe the formation of hadrons. To evolve the process from Q to $\Lambda_{\rm QCD}$ scales, parton showers are nowadays the most powerful tools. As described in Section 2.2, they can account for the increasing number of particle splittings and for the enhanced all-order contributions that occur when considering non-trivial observables or more realistic events involving experimental analysis cuts. After the formation of primary hadrons described with hadronization models, their decays to a set of hadrons which is stable enough to reach the detectors have to be correctly modelled, as well. Since protons and not partons are made collide at the LHC, we also need to parametrize the probability to resolve, inside the proton, a parton which can partecipate in the hard event: that is done in terms of a non-perturbative object, named Parton distribution function (PDF). But the remaining partons in the proton can also interact, contributing to extra QCD activity which accompanies the hard event, commonly referred to as *underlying events* (UE). In their turn, the two hadrons responsible for the hard-scattering process are part of colliding beams comprising of many hadrons, which are also expected to interact (*pile-up*), even though mainly at energies substantially lower than Q. Indeed, for a proper LHC event simulation, all possible classes of events have to be covered, in such a way that the total proton-proton cross section results from the sum of their different contributions. Events described so far typically generate final state particles covering the whole pseudorapidity [4] range: they are named non-diffractive inelastic contributions. But 40% of the total hadronic cross section at LHC comes from diffractive events, where a large pseudorapidity interval is devoid of hadronic activity. They originate from exchanges of soft gluons which can not resolve the internal partonic structure of the proton. These events are mainly described in terms of elastic $pp \to pp$ scatterings or of processes where one or both protons turn into a bunch of particles X with the same proton quantum numbers. The latter are further classified into single

[4] The pseudorapidity, usually denoted by η, is a kinematic quantity related to the polar angle θ between the beam and the detection directions, and defined as $\eta \equiv -\log\tan(\theta/2)$.

$pp \to p + X$ and double $pp \to X + Y$ dissociation processes, where here the + sign denotes the large pseudorapidity gap characteristic of diffractive events.

All of the ingredients entering a LHC event simulation are commonly available in a General purpose Monte Carlo (GPMC) event generator. Improving on the accuracy of theoretical predictions for collider physics translates into a cooperative effort to improve the many different components of the simulation itself. The fact that the various building blocks of the process evolution from high to low scales can be treated to a large extent indipendently is based on QCD factorization. According to the factorization assumption, physics occuring at low energy scales is not expected to drastically affect high-energy phenomena which take place at much earlier time in the evolution. Within this complex picture, the proper combination of fixed-order calculations, needed for the description of the hard event, and the parton shower, accounting for QCD radiation down to the hadronization scales, plays a crucial role. Such a combination, which is the main focus of this Thesis, allows to interface the best-available fixed-order results for a given process with the flexibility of GPMC generators, which is indispensable to obtain a realistic event description to be compared to LHC data. In Chapter 2, the basic ingredients of the fixed-order (FO) approach and of parton shower (PS) simulations are briefly introduced in sections 2.1 and 2.2, respectively. Available techniques to combine both strategies at different levels of accuracy are summarized in Section 2.3. After this general overview, in Chapter 3 we specialize the discussion to the theoretical framework that has been used and developed in this work. Therefore, in Section 3.1 the POWHEG method, empowered by the MiNLO′ approach (presented in Section 3.2), is discussed: its usage allows to *merge* different next-to-leading order (NLO) QCD calculations and *match* them with parton shower programs. Finally, the recently developed MiNNLO$_{\rm PS}$ method, to interface any colour-singlet process at next-to-next-to-leading order (NNLO) accuracy to parton showers (so reaching NNLO+PS accuracy), is discussed in Section 3.3: the extension of the MiNNLO$_{\rm PS}$ approach to general colour-singlet production is the main result of this work. In Chapter 4, this method is applied to an extremely relevant class of processes for LHC search, namely diboson processes. Here we present results for $Z\gamma$ at NNLO+PS in Section 4.1 and 4.2, but also for massive diboson production such as W^+W^- at NNLO+PS in Section 4.3 and ZZ at nNNLO+PS (once the loop-induced gluon-gluon channel is also properly included) in Section 4.4. The main achievements of this Thesis are finally summarized in the Conclusions of Chapter 5.

2 Theoretical tools for precise phenomenology

In the previous chapter we have described the main components of a GPMC event generator, which is indispensable for comparing theory results to hadron collider data. In view of high-precision phenomenology, in this chapter we focus on fixed-order calculations and parton shower techniques, which describe the evolution of the physics process within the perturbative regime, before hadronization has occurred. After a brief review of the two approaches, respectively in Section 2.1 and Section 2.2, in Section 2.3 we present some of the customary techniques to consistently combine them.

2.1 Fixed-order perturbation theory

In the context of QCD, theoretical predictions for the hadronic cross section of a process of interest rely on the parton model picture and the factorization ansatz. We can start considering the cross section $\sigma(h_1 h_2 \to X)$ for a hadron scattering process $h_1 h_2 \to X$, where h_1 and h_2 are the incoming hadrons and X denotes a generic final state. Then, $\sigma(h_1 h_2 \to X)$ can be written as a sum of convolutions of the partonic cross sections $\sigma_{ij,\,X}$ (where i and j label the incoming partons which contribute to the process), which enclose the short distance physics, and some non-perturbative universal functions, named parton distribution functions (PDFs). If $K_\oplus$ and $K_\ominus$ are the four-momenta of the hadrons h_1 and h_2 moving in opposite directions along the collision axis, and if X is a system of n particles with momenta $k_1, \ldots, k_n$, then four-momentum conservation reads:

$$x_\oplus K_\oplus + x_\ominus K_\ominus = \sum_{i=1}^{n} k_i \tag{2.1}$$

where $x_\oplus$, $x_\ominus \in [0, 1]$ are the momentum fractions of the hadron momenta which are shared with the colliding partons, whose momenta are $k_\oplus = x_\oplus K_\oplus$ and $k_\ominus = x_\ominus K_\ominus$. Hadron constituents are assumed to be massless, so that $k_\oplus^2 = k_\ominus^2 = 0$, while final-state momenta are considered to be on-shell. If we define the total hadronic energy $S = (K_\oplus + K_\ominus)^2$, the latter will be related to the partonic energy $Q^2 = (k_\oplus + k_\ominus)^2$ (often denoted also with the symbol of the Mandelstam invariant s) via the relation $Q^2 = x_\oplus x_\ominus S$. Then, in the collinear factorization ansatz [52, 53], which is the one we will refer to throughout this Thesis, the hadronic cross section can be written as:

$$\sigma(h_1 h_2 \to X) = \sum_{ij} \int dx_\oplus dx_\ominus f_{i/h_1}(x_\oplus,\, \mu_\mathrm{F}) f_{j/h_2}(x_\ominus,\, \mu_\mathrm{F}) \sigma_{ij,\,X}(Q^2, \mu_\mathrm{R}^2, \mu_\mathrm{F}^2) + \mathcal{O}\left(\frac{\Lambda_\mathrm{QCD}^2}{Q^2}\right) \tag{2.2}$$

where we stress how the characteristic scale of the hard process Q^2 in $\sigma_{ij,X}$ depends on $x_\oplus$ and $x_\ominus$, which are integrated over in the convolution with the PDFs of the incoming hadrons. The non-perturbative function $f_{i/h}(x, \mu_\mathrm{F})$ provides the probability to resolve at the energy scale μ_F a parton i with momentum fraction x inside the hadron h. The scale μ_F is an unphysical scale named *factorization scale*, which also enters the partonic cross section, together with a second unphysical scale μ_R, named *renormalization scale*. The need to introduce these two unphysical scales in the theory will be motivated later in due course. Finally, the last term in Eq. (2.2) reminds us that the factorization of long- and short-distance physics holds up to non-perturbative corrections. As clarified for instance in Ref. [2], all high-energy QCD processes come with power-correction terms proportional to the ratio of the non-perturbative over the perturbative scale $(\Lambda_\mathrm{QCD}/Q)^p$, with a power p which depends on the observable of interest. Since these corrections can range from the percent to the permille level to the power of p, especially linear power corrections ($p = 1$) can become comparable to the accuracy reached in the calculation of the first part of Eq. (2.2), based on perturbation theory, as we will discuss in the following. Therefore, gaining better control on non-perturbative power-corrections is also important. Some recent progresses in this direction for collider observables can be found for instance in Ref. [54, 55].

For the time being, we can just assume that the partonic cross section $\sigma_{ij,X}$ to produce a final state X of n particles coincides with the partonic cross section $\sigma_{ij,n}$ to generate a final state comprising of exactly n particles and nothing else. Then, $\sigma_{ij,n}$ can be derived from first principles and connected to the fundamental theory enclosed in the SM Lagrangian presented in Section 1.1 by means of the well-known *Fermi's golden rule*:

$$\sigma_{ij,n} = \frac{(2\pi)^4}{4k_\oplus \cdot k_\ominus} \int \prod_{m=1}^{n} \frac{d^3k_m}{(2\pi)^3 2k_m^0} \delta^4\left(k_\oplus + k_\ominus - \sum_{m=1}^{n} k_m\right) |\mathcal{M}_{ij,n}|^2 \,. \tag{2.3}$$

The kinematic ingredients of the formula comprise of a flux factor $4k_\oplus \cdot k_\ominus$ and the *n-particle phase space* $d\Phi_\mathrm{n}$:

$$d\Phi_\mathrm{n}(k_\oplus + k_\ominus;\, k_1, \ldots, k_n) = (2\pi)^4 \delta^4\left(k_\oplus + k_\ominus - \sum_{m=1}^{n} k_m\right) \prod_{m=1}^{n} \frac{d^3k_m}{(2\pi)^3 2k_m^0} \,. \tag{2.4}$$

If $\Phi_\mathrm{n} = \{k_1, \ldots, k_n\}$, it will turn out to be useful to define the set of variables $\mathbf{\Phi}_\mathrm{n} = \{x_\oplus, x_\ominus, k_1, \ldots, k_n\}$, so that $d\mathbf{\Phi}_\mathrm{n} = dx_\oplus dx_\ominus d\Phi_\mathrm{n}(k_\oplus + k_\ominus; k_1, \ldots, k_n)$. On the other hand, the theory model enters Eq. (2.3) only via the *squared matrix element* (or squared amplitude) $|\mathcal{M}_{ij,n}|^2$, where sums over colour and spin indices are understood. Quantum field theory provides a recipe to compute $|\mathcal{M}_{ij,n}|^2$, making use of the Wick's theorem [56] and the Lehman-Symanzik-Zimmermann (LSZ) reduction formula [57]. What different models have to provide is the correct set of interaction verteces (or Feynman rules) among the fields of the theory, which can be extracted from the Lagrangian (for QCD interactions, for instance, directly from Eq. (1.1), once gauge-fixing and ghost terms are also accounted for). Then, $\mathcal{M}_{ij,n}$ can be calculated by accounting for all Feynman diagrams contributing to the scattering process $ij \to X$ which can be built from this set of interactions. Each interaction vertex is proportional to the product of some coupling constants characteristic

of the theory. Relying on the smallness of some coupling constants, perturbation theory allows to organize the calculation as a perturbative expansion of $\mathcal{M}_{ij,\,n}$, where dominant contributions come from terms counting the lowest powers of these couplings. Within the SM, starting from the first term contributing to a given process, named the *leading order* (LO) term, different kinds of corrections can be added, depending on the coupling used to carried out the expansion: QCD, QED or EW corrections. When simultaneously considered, mixed terms will also appear. Anyway, QCD corrections are known to be the dominant ones: even if the strong coupling $\alpha_s \sim 0.1$ for energies within 100 GeV and 1 TeV, and so well within the perturbative regime, its value is still much larger than the QED and EW couplings in the same energy range. That is why in the bulk of this Thesis we will mainly focus on QCD types of corrections. By dropping the parton indices for ease of notation (which can easily be reintroduced back where needed), we can express the α_s-perturbative expansion of $\mathcal{M}_{ij,\,n}$ as:

$$\begin{aligned}
\mathcal{M}_n =& \left(\frac{g_S}{\sqrt{4\pi}}\right)^N \left(\mathcal{M}_n^{(0)} + \left(\frac{g_S}{\sqrt{4\pi}}\right)^2 \mathcal{M}_n^{(1)} + \left(\frac{g_S}{\sqrt{4\pi}}\right)^4 \mathcal{M}_n^{(2)} + \mathcal{O}\left(\frac{g_S^6}{(4\pi)^3}\right)\right) \\
\rightarrow \quad & |\mathcal{M}_n|^2 = \alpha_s^N \Big(|\mathcal{M}_n^{(0)}|^2 + 2\alpha_s Re[\mathcal{M}_n^{(0)\star} \cdot \mathcal{M}_n^{(1)}] + \alpha_s^2 |\mathcal{M}_n^{(1)}|^2 \\
& + 2\alpha_s^2 Re[\mathcal{M}_n^{(0)\star} \cdot \mathcal{M}_n^{(2)}] + \mathcal{O}(\alpha_s^3) \Big) . \qquad (2.5)
\end{aligned}$$

In the previous formula, we factored out N powers of α_s, which contribute to the LO squared amplitude, $|\mathcal{M}_n^{(0)}|^2$, usually referred to as *Born* term $\mathcal{B}$.

2.1.1 Ultraviolet divergences

Starting from the next-to-leading order (NLO) term of Eq. (2.5), which is suppressed by one extra power of α_s with respect to the Born term, Feynman diagrams contributing to the scattering process require the calculation of loop integrals, where the momentum ℓ associated to an inner line of the diagram and not constrained by momentum conservation has to be integrated over the whole phase-space volume. Depending on the number of internal lines (i.e particle propagators) forming the loop, these integrals might present non-integrable divergences for $\ell^2 \rightarrow \infty$, which are usually named *ultraviolet* (UV) divergences. This kind of divergences, that emerge when probing high energies, is handled via a standard machinery in QFT. First of all, divergences are isolated by means of a *regularization* procedure. A very common one is *dimensional regularization* [58], which has the nice virtue of preserving gauge invariance throughout the calculation. In this procedure, loop integrals are made finite by rising the phase-space dimensionality to $d = 4 - 2\epsilon$, for $\epsilon > 0$: by doing that, divergences are recovered through the limit $\epsilon \rightarrow 0$, and they appear as poles in the ϵ regulator. For a special class of theories, named *renormalizable*, these divergences can be systematically removed by introducing a finite set of counterterms, defined in such a way to compensate order by order in perturbation theory for the UV divergences emerging from loop calculations. These counterterms can be absorbed using proper *renormalization constants* into a redefinition of the original parameters and

fields of the Lagrangian, which then becomes a function of new *bare* quantities. The renormalizability condition is satisfied by any theory whose Lagrangian comprises of terms (to be interpreted as *operators* after quantization) having at most mass dimension equal to four: this requirement, met by the SM Lagrangian, ensures that all UV divergences can be removed by simply casting the parameters of the theory in terms of bare ones, without any need to introduce new kinds of interaction verteces. These theories can be thought of as being to a great extent insensitive to high-energy scales, which might be associated to new physics. Once the ϵ-poles are removed, there is still some arbitrariness in the definition of the finite part of the counterterms: different choices define different *renormalization schemes*. A scheme which is largely used in QCD, where confinement does not allow to use physical quark masses to resolve the ambiguity in the counterterm definition, is the $\overline{\text{MS}}$ scheme. Here counterterms just subtract the ϵ singularities, while their finite parts are set to zero, up to a $\log(4\pi) - \gamma_E$ term (where γ_E is the Eulero-Mascheroni constant), which is also subtracted for convenience. One extra arbitrariness in the dimensional regularization procedure resides in the possibility to analytically continue to d dimensions only the loop momentum ℓ or also other quantities, such us external momenta and helicity vectors: different choices define different *reduction schemes*, which introduce mismatches amounting to finite constants in intermediate results.

Anyway, all this machinery comes with a price: for consistency reasons, when analytically continuing to d dimensions, the dimensionless nature of the expansion parameter (i.e g_S in QCD) has to be preserved. That is done by introducing the unphysical renormalization scale μ_{R}, mentioned in the previous section. If Z_g is the renormalization constant associated to the strong coupling, the bare $g_{S,0}$ and physical g_S couplings will be related through:

$$g_{S,0} = \mu_{\text{R}}^{\epsilon} Z_g(\epsilon)\, g_S \,. \tag{2.6}$$

Since the bare Lagrangian quantities can not depend on μ_{R}, the μ_{R} derivative of Eq. (2.6) results in an equation, named *renormalization group equation* (RGE), which describes the μ_{R} dependence of g_S. In terms of α_s, the RGE for the strong coupling constant in $d = 4$ reads:

$$\mu_{\text{R}}^2 \frac{\partial \alpha_s(\mu_{\text{R}})}{\partial \mu_{\text{R}}^2} \equiv \beta(\alpha_s) = \lim_{\epsilon \to 0} \left[-\alpha_s(\mu_{\text{R}}) \cdot \mu_{\text{R}}^2 \frac{\partial}{\partial \mu_{\text{R}}^2} \log(\mu_{\text{R}}^{2\epsilon} Z_g^2(\epsilon)) \right] = -\alpha_s(\mu_{\text{R}}) \sum_{k=1}^{\infty} \beta_k \alpha_s^k(\mu_{\text{R}}) \,, \tag{2.7}$$

where $\beta(\alpha_s)$ is the famous 4-dimensional QCD β-function, which admits a perturbative expansion starting from order α_s^2, and whose first coefficients are given by:

$$\begin{aligned} \beta_0 &= \frac{11 C_A - 2 n_f}{12\pi} \,, \quad \beta_1 = \frac{17 C_A^2 - 5 C_A n_f - 3 C_F n_f}{24\pi^2} \,, \\ \beta_2 &= \frac{2857 C_A^3 + (54 C_F^2 - 615 C_F C_A - 1415 C_A^2) n_f + (66 C_F + 79 C_A) n_f^2}{3456 \pi^3} \,, \end{aligned} \tag{2.8}$$

with n_f the number of active light quark flavours [1]. Starting from β_2, the terms of the

[1] A quark flavour is referred to as *active* if the corresponding quark mass is lower than the probed energy scale.

β-function are renormalization scheme dependent: here we reported the result for β_2 in the $\overline{\text{MS}}$ scheme. Since Eq. (2.7) describes how α_s depends on an arbitrary energy scale, it can be used to connect the values of α_s at two different scales, for instance μ_R and Q, provided that both $\alpha_s(\mu_\text{R})$ and $\alpha_s(Q)$ are in the perturbative regime. Solving Eq. (2.7) at first order in α_s gives [59, 60]:

$$\alpha_s(Q) = \frac{\alpha_s(\mu_\text{R})}{1 + \alpha_s(\mu_\text{R})\beta_0 \log\left(\frac{Q^2}{\mu_\text{R}^2}\right)} \tag{2.9}$$

which encloses QCD asymptotic freedom, according to which $\alpha_s(Q) \to 0$ for $Q \to \infty$. The solution of Eq. (2.7) can also be expressed in terms of an integration constant Λ_QCD, which gives the location of the Landau pole, i.e the energy at which the strong coupling would diverge if we extrapolate the QCD RGE outside the perturbative regime. At first order in α_s we can write:

$$\alpha_s(Q) = \frac{1}{\beta_0 \log\left(\frac{Q^2}{\Lambda^2_\text{QCD}}\right)} . \tag{2.10}$$

Despite the fact that the exact value of Λ_QCD depends on the perturbative order at which Eq. (2.7) is solved, on the number of active quark flavours and on the renormalization scheme employed (see for instance the discussion in Ref. [1]), still its order of magnitude gives an indication of the energy regime at which non-perturbative dynamics starts to dominate.

Moving back to Eq. (2.5), it is now clear that α_s therein and the different coefficients of the expansions have a dependence on μ_R, which translates in a μ_R dependence of the partonic cross section $\sigma_{ij,n}$. But since the physical hadronic cross section can not depend on an unphysical scale, this μ_R dependence has to cancel out when contributions from all orders are summed up. By truncating the expansion of Eq. (2.5) at a given order k in α_s, a residual μ_R dependence survives for terms of order $k' > k$: the higher k the weaker the residual dependence on μ_R. Therefore, evaluating the cross section for values of μ_R around a central scale $\mu_{\text{R},0}$ can give an indication of the size of the missing higher order contributions that have been neglected by truncating the expansion. Moreover, in order to improve the convergence of the perturbative series, which might be jeopardized by potentially large logarithmic terms $\sim \log(Q/\mu_\text{R})$ entering Eq. (2.9), it is customary to fix the arbitrary scale $\mu_{\text{R},0}$ to energies characteristic of the hard event, i.e $\mu_{\text{R},0} \sim Q$. For non-trivial processes, multiple scales can characterize the hard event: as we will see for instance in Section 4.3.4, in such cases there is no correct choice a priori, and the dependence of the result on different μ_R choices can give a hint on the convergence of the perturbative series and on the importance of higher order terms.

To summarize, when considering the matrix element $\mathcal{M}_n$ entering the cross section for the production of exactly n final states making up the system X, then tree-level [2],

[2] To simplify this general discussion, here we explicitly referred to processes receiving their LO contribution from tree-level diagrams, which involve no loops. However, one should keep in mind that some processes are loop-induced, so that the calculation of one-loop diagrams is required already at the LO, of two-loop diagrams at the NLO and so on.

one-loop and two-loop diagrams contribute to its perturbative expansion at order g_S^0, g_S^2 and g_S^4, respectively. At the squared amplitude level, the LO contribution comes from the squared tree-level amplitude or Born term $\mathcal{B}(\mathbf{\Phi}_\mathrm{n})$, where the phase-space dependence has been highlighted. Once UV divergences have been subtracted, the cross section receives NLO contirbutions of order $\mathcal{O}(\alpha_s^{N+1})$ from the so called *bare virtual* term $\mathcal{V}_b(\mathbf{\Phi}_\mathrm{n})$, given by twice the real part of the tree-level and one-loop amplitude interference. Then, some next-to-next-to-leading order (NNLO) contributions of order $\mathcal{O}(\alpha_s^{N+2})$ come from the interference of the tree-level and two-loop amplitudes and from the one-loop squared term.

2.1.2 Infrared divergences

When performing loop integrals such as the ones introduced in the previous section, UV divergences are not the only kind of singularities that emerge. When massless particles enter the loops (as it can be if we have a gluon or a photon propagator, or when quark masses are neglected), the integral can diverge also in the limit where the loop momentum ℓ gets soft, i.e $\ell^2 \to 0$, and/or collinear to an external momentum k, i.e $\ell \cdot k \to 0$. Divergences arising from these limits are named *infrared* (IR) divergences. By making use of dimensional regularization, these integrals can be made finite by analytic continuation with a regulator $\epsilon < 0$, so that each of the two IR divergences is recovered in terms of an ϵ-pole singularity (which means at one-loop level we can have at most a $1/\epsilon^2$ singularity for a simultaneously soft and collinear divergence). But the handling of IR divergences is quite different from that of UV ones: the cancellation of IR divergences occurs when physically sound quantities are computed. Indeed, for theories like QED, we know that a propagating particle (for instance an electron) going through a detector can not be experimentally distinguished from another particle emitting a certain number of soft and/or highly collinear photons, due to the limited detector resolution. In QCD, we can not refer to an experimental resolution cutoff, since quarks and gluons are not asymptotic states, but a physical cutoff comes from the theory itself. Indeed, IR singularities are associated to phase-space regions where perturbation theory breaks down. Still, the possibility to compute physical quantites that are free of IR divergences signal that such quantities are insensitive to the non-perturbative regime: they are usually referred to as *infrared safe observables*.

According to the physical picture we have provided, it can be understood that the partonic cross section $\sigma_{ij,\,n}$ for the production of exactly n final-state partons is not an infrared safe observable. However, $\sigma_{ij,\,X}$, which is the *inclusive* cross section giving the probability of having n partons plus additional radiation, is infrared safe. For the calculation of $\sigma_{ij,\,X}$, amplitudes for different final-state multiplicities have to be included. In what follows, we specialize the discussion to the NLO QCD case. Up to order $\mathcal{O}(\alpha_s^{N+1})$, $\sigma_{ij,\,X}$ receives contributions from $\mathcal{B}(\mathbf{\Phi}_\mathrm{n}) = |\mathcal{M}_n^{(0)}|^2$, $\mathcal{V}_b(\mathbf{\Phi}_\mathrm{n}) = 2\,Re[\mathcal{M}_n^{(0)\star}\mathcal{M}_n^{(1)}]$ and the *real* term $\mathcal{R}(\mathbf{\Phi}_{\mathrm{n}+1}) = |\mathcal{M}_{n+1}^{(0)}|^2$, so that we can write our hadronic cross section of Eq. (2.2)

as:

$$\sigma(h_1h_2 \to X) = \int d\mathbf{\Phi}_{\rm n}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^N \left[\mathcal{B}(\mathbf{\Phi}_{\rm n}) + \alpha_s\, \mathcal{V}_b(\mathbf{\Phi}_{\rm n})\right] + \int d\mathbf{\Phi}_{\rm n+1}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^{N+1}\, \mathcal{R}(\mathbf{\Phi}_{\rm n+1}) + \mathcal{O}(\alpha_s^{N+2})\,, \tag{2.11}$$

where the $\mu_{\rm F}$ dependence has been momentarily dropped and the sum over different partonic channels is understood. Contrary to the loop amplitude $\mathcal{M}_n^{(1)}$, the tree-level one $\mathcal{M}_{n+1}^{(0)}$ is finite, but IR divergences that exactly cancel the IR loop ones appear when $|\mathcal{M}_{n+1}^{(0)}|^2$ is integrated over the full real phase space $\mathbf{\Phi}_{\rm n+1}$. That is due to the universal behaviour of $|\mathcal{M}_{n+1}^{(0)}|^2$ when one of the final-state partons become soft and/or collinear to another parton.

Specifically, if a pair of partons with momenta k_i and k_j gets collinear, then we have [61]:

$$|\mathcal{M}_{n+1}^{(0)}(k_1, ..., k_i, ..., k_j, ..., k_{n+1})|^2 \quad \xrightarrow{\vec{k}_i \parallel \vec{k}_j} \tag{2.12}$$

$$\frac{8\pi\alpha_s\mu^{2\epsilon}}{2k_i\cdot k_j}\mathcal{M}^{(0)}_{s_1,...,s_I,...,s_n;n}(k_1, ..., k_I, ..., k_n)\hat{P}_{j,I}^{s_I,s_I'}(z, k_\perp, \epsilon)\mathcal{M}^{(0)\star}_{s_1,...,s_I',...,s_n;n}(k_1, ..., k_I, ..., k_n)$$

where the final-state momentum dependence has been made explicit at the level of the matrix elements, together with the spin indices s. The momentum k_I is simply obtained by replacing parton momenta k_i and k_j with their sum, i.e $k_I = k_i + k_j$. The spin index s_I is the spin of the splitting parton I in the splitting process $I \to i + j$. The previous equation shows that in the collinear limit the real amplitude has a universal behaviour: a single pole in $k_i \cdot k_j$ multiplies the spin-correlated Born amplitude, where spin correlation is controlled by the universal d-dimensional spin-dependent *Altarelli-Parisi* (AP) *splitting functions* $\hat{P}_{j,I}^{s_I,s_I'}(z, k_\perp, \epsilon)$. They depend on the regulator ϵ, the energy fraction z that the parton j takes from the splitting parton I in the emission of i, and the transverse momentum $k_\perp$ of the splitting pair. Clearly, both z and $k_\perp$ have to be defined within a proper parametrization of the kinematics of the emitted parton. The exact form of such splitting functions can be found for instance in Ref. [61]. Here we report their spin-averaged form at leading order in the limit $\epsilon \to 0$, which are the ones we will mainly refer to throughout this work:

$$\begin{aligned} \hat{P}_{qq}^{(0)}(z) &= C_F\left[\frac{1+z^2}{1-z}\right] & \hat{P}_{gq}^{(0)}(z) &= C_F\left[\frac{1+(1-z)^2}{z}\right] \\ \hat{P}_{qg}^{(0)}(z) &= T_R\Big[1-2z(1-z)\Big] & \hat{P}_{gg}^{(0)}(z) &= 2C_A\left[\frac{1-z}{z}+\frac{z}{1-z}+z(1-z)\right]. \end{aligned} \tag{2.13}$$

We see that the splitting functions can be singular themselves in the limits $z \to 0$ or $z \to 1$, where one parton of the splitting pair also gets soft. Only the function $\hat{P}_{qg}(z)$ for a gluon splitting in a quark-antiquark pair does not come with soft singularities.

On the other hand, if we consider the limit of a gluon with momentum k_s getting soft, we can write [61]:

$$|\mathcal{M}^{(0)}_{n+1}(k_1, \ldots, k_s, \ldots, k_{n+1})|^2 \quad \xrightarrow{k_s \to 0} \tag{2.14}$$
$$-8\pi\alpha_s\mu^{2\epsilon} \sum_{i,j=1}^{n} S_{ij}(k_s)\mathcal{M}^{(0)}_{c_1,\ldots,c_i,\ldots,c_j,\ldots,c_n;n}(k_1, \ldots, k_n)T^a_{c_i,c'_i}T^a_{c_j,c'_j}\mathcal{M}^{(0)\star}_{c_1,\ldots,c'_i,\ldots,c'_j,\ldots,c_n;n}(k_1, \ldots, k_n)$$

where the labels c are colour indices and $T^a_{c,c'}$ are the $SU(3)_c$ generators either in the adjoint or fundamental representation according to whether the emitting parton is a gluon or a quark, respectively. These colour matrices enter the definition of the colour-correlated Born matrix element, multiplied by the universal soft function $S_{ij}(k_s)$, named *eikonal factor*, for the soft gluon with momentum k_s which is colour-connected to the partons i and j (clearly $i \neq j$):

$$S_{ij}(k_s) = \frac{k_i \cdot k_j}{2(k_i \cdot k_s)(k_j \cdot k_s)}\,, \tag{2.15}$$

which contains the singular terms of the soft kinematics.

Therefore, when integrating $|\mathcal{M}^{(0)}_{n+1}|^2$ over $\mathbf{\Phi}_{\mathrm{n+1}}$, divergences arise from phase-space regions where a parton enters the soft and/or collinear regime, where $|\mathcal{M}^{(0)}_{n+1}|^2$ exhibits the singular behaviours of Eq. (2.12) and Eq. (2.14). Such $n+1$ parton configurations are kinematically degenerate with a n-parton one and compensate for the IR virtual singularities, which indeed live in the $\mathbf{\Phi}_{\mathrm{n}}$ phase space. The fact that an exact cancellation between the two sources of IR divergences in $\mathcal{V}_b(\mathbf{\Phi}_{\mathrm{n}})$ and $\mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})$ takes place is guaranteed by the Kinoshita-Lee-Nauenberg (KLN) theorem [62, 63], which states that any unitary theory is free of IR divergences as soon as configurations involving all possible final and initial states are summed over.

We can generalize Eq. (2.11) so as to account for the expectation value of a general observable O at NLO accuracy:

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\mathrm{n}}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^N\, [\mathcal{B}(\mathbf{\Phi}_{\mathrm{n}}) + \alpha_s\, \mathcal{V}_b(\mathbf{\Phi}_{\mathrm{n}})]\, O_n(\mathbf{\Phi}_{\mathrm{n}}) \\ &+ \int d\mathbf{\Phi}_{\mathrm{n+1}}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^{N+1}\, \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}}) O_{n+1}(\mathbf{\Phi}_{\mathrm{n+1}}) + \mathcal{O}(\alpha_s^{N+2})\end{aligned} \tag{2.16}$$

where O_n and O_{n+1} are the expressions of the observable O in terms of n and $n+1$ final-state particles, respectively. We can simply recover $\sigma(h_1h_2 \to X)$ by setting $O_n(\mathbf{\Phi}_{\mathrm{n}}) = 1$ for all n. In the general case of Eq. (2.16), the cancellation of IR divergences between the virtual and the integrated real contributions can take place only if O is infrared safe, a condition that we can formulate as:

$$O_{n+1}(x_\oplus, x_\ominus; k_1, \ldots, k_i, \ldots, k_j, \ldots, k_{n+1}) \xrightarrow{\vec{k}_i \parallel \vec{k}_j} O_n(x_\oplus, x_\ominus; k_1, \ldots, k_i + k_j, \ldots, k_n) \tag{2.17}$$

$$O_{n+1}(x_\oplus, x_\ominus; k_1, \ldots, k_s, \ldots, k_{n+1}) \xrightarrow{\vec{k}_s \to 0} O_n(x_\oplus, x_\ominus; k_1, \ldots, k_{s-1}, k_{s+1}, \ldots, k_n) \tag{2.18}$$

which simply states that O is insensitive to any additional soft and/or collinear splitting.

By now, we have just explicitly referred to final-state momenta when considering the IR limits of $|\mathcal{M}^{(0)}_{n+1}|^2$ in Eq. (2.12) and Eq. (2.14). Similar formulae clearly hold also for initial-state momenta. Unfortunately, for lepton-hadron and hadron-hadron collisions initial-state collinear singularities (i.e when a final-state parton with momenum k_i gets collinear to $k_\oplus$ or $k_\ominus$) do not cancel out in Eq. (2.16), but, in a dimensional regularization picture, they survive as $1/\epsilon$ poles [61]. Indeed, the KLN theorem does not apply to processes involving initial-state hadrons, whose outgoing partons carry a well-defined momentum. Singularities associated to *identified partons* can be dealt with by treating the PDFs as bare divergent objects [64], $f_{0,\,h_1}(x_\oplus)$ and $f_{0,\,h_2}(x_\ominus)$, to be contrasted with the measurable ones. Measurable PDFs are the ones obtained via a renormalization procedure that, as with UV and final-state IR divergences, introduces an additional unphysical scale $\mu_{\rm F}$ in the theory, dubbed factorization scale. Since the form of initial-state collinear singularities is also universal, the collinear factorization theorem [3] guarantees that these singularities can be factored into the PDFs in a process independent way [64]. All over the calculation, the bare PDF $f_{0,\,h}(x)$ has to be replaced by the renormalized PDF $f_h(x,\,\mu_{\rm F})$, with the additional $\mu_{\rm F}^2$ dependence introduced by the renormalization procedure and the subtraction of singularities. This approach is actually equivalent to introducing two bare collinear counterterms $\mathcal{G}_{\oplus,b}$ and $\mathcal{G}_{\ominus,b}$ at the level of the cross section. These counterterms, constructed order by order in perturbation theory, are defined on the phase-space configurations:

$$\mathbf{\Phi}_{\rm n,\oplus} = \{x_\oplus,\, x_\ominus,\, z,\, k_1,\, \ldots,\, k_n\} \quad z x_\oplus K_\oplus + x_\ominus K_\ominus = \sum_{i=1}^{n} k_i\,, \tag{2.19}$$

$$\mathbf{\Phi}_{\rm n,\ominus} = \{x_\oplus,\, x_\ominus,\, z,\, k_1,\, \ldots,\, k_n\} \quad x_\oplus K_\oplus + z x_\ominus K_\ominus = \sum_{i=1}^{n} k_i\,, \tag{2.20}$$

where z is the parton energy fraction taken away from the parton momentum $k_\oplus$ or $k_\ominus$ in the initial splitting process. In the explicit form of $\mathcal{G}_{\oplus,b}$ and $\mathcal{G}_{\ominus,b}$, the freedom of subtracting some additional finite contributions defines a *factorization scheme* choice. It is worth remarking that, since PDFs are non-perturbative objects, they must be provided as an input to the theory. Even though some first attempts to derive PDFs from first principles in the context of lattice QCD have started to appear (for the most recent results see for instance Refs. [68–70] and references therein), they are usually extracted from fits of data coming mainly from DIS (more sensitive to PDFs, since they only involve one hadron), but also from many LHC processes. Depending on the accuracy of the matrix elements used to describe the data, LO, NLO or even NNLO PDF fits can be carried out, and the results are then stored into proper grids.

Despite their non-perturbative nature, the dependence of PDFs on the scale $\mu_{\rm F}$ can be described within perturbation theory. The requirement for measurable quantities (such

[3] In the collinear factorization of long- and short-distance physics, bare PDFs only depend on the parton energy fraction x: the parton coming out of the hadron and taking part in the hard process is assumed to be collinear to the original hadron direction. A more differential approach, which relaxes the collinear factorization assumption, is to consider transverse momentum dependent (TMD) PDFs, where the transverse momentum of the outgoing parton with respect to the hadron direction is left unintegrated [65–67].

as structure functions in DIS experiments) to be independent of the unphysical scale $\mu_{\rm F}$ leads to evolution equations for PDFs, known as Dokshitzer, Gribov, Lipatov, Altarelli and Parisi (DGLAP) equations [71–73]. In their most general form, DGLAP equations are formulated as $(2n_f+1)$-dimensional matrix element equations in the space of quarks, antiquarks and gluons [4]:

$$\mu_{\rm F}^2 \frac{\partial}{\partial \mu_{\rm F}^2} \begin{pmatrix} f_q(x,\mu_{\rm F}) \\ f_g(x,\mu_{\rm F}) \end{pmatrix} = \frac{\alpha_s(\mu_{\rm F})}{2\pi} \sum_{q'} \int_x^1 \frac{dz}{z} \begin{pmatrix} P_{q'q}(\frac{x}{z},\alpha_s(\mu_{\rm F})) & P_{gq}(\frac{x}{z},\alpha_s(\mu_{\rm F})) \\ P_{q'g}(\frac{x}{z},\alpha_s(\mu_{\rm F})) & P_{gg}(\frac{x}{z},\alpha_s(\mu_{\rm F})) \end{pmatrix} \begin{pmatrix} f_{q'}(z,\mu_{\rm F}) \\ f_g(z,\mu_{\rm F}) \end{pmatrix} \tag{2.21}$$

where the parton labels q and g to distinguish among quark and gluon PDFs have been introduced. In this notation, q and q' can take any of the $2n_f$ flavour and antiflavour values. Equation 2.21 shows how a PDF of a given flavour is closely connected to other parton PDFs through the AP splitting functions P_{ji} (where here again i and j refer to any parton). The dependence of P_{ji} on $\alpha_s(\mu_{\rm F})$ has been highlighted to stress that AP kernels admit a perturbative expansion, starting from $\mathcal{O}(\alpha_s^0)$ (with the exception of $P_{q'q}$, whose α_s^0 contribution is zero unless $q'=q$). The splitting functions P_{ji} entering DGLAP equations differ from the $\hat{P}_{ji}$ ones of Eq. (2.13): the former are also referred to as *regularized* AP splitting functions, that at LO read

$$\begin{aligned} P_{qq}^{(0)}(z) &= C_F\left[\frac{1+z^2}{[1-z]_+} + \frac{3}{2}\delta(1-z)\right] \qquad P_{gq}^{(0)}(z) = \hat{P}_{gq}^{(0)}(z) \qquad P_{qg}^{(0)}(z) = \hat{P}_{qg}^{(0)}(z) \\ P_{gg}^{(0)}(z) &= 2C_A\left[\frac{1-z}{z} + \frac{z}{[1-z]_+} + z(1-z)\right] + \delta(1-z)\frac{11C_A - 4n_f T_R}{6}\,, \end{aligned} \tag{2.22}$$

where soft singularities associated to the limit $z \to 1$ have been regularized by means of the δ and plus distributions, the latter defined so that

$$\int_0^1 dz \frac{f(z)}{[1-z]_+} = \int_0^1 dz \frac{f(z)-f(1)}{1-z} \quad \text{and} \quad \frac{1}{[1-z]_+} = \frac{1}{1-z} \text{ for } x \in [0,1[\tag{2.23}$$

for any smooth function f. By making use of regularized splitting functions, *unitarity* is enforced in DGLAP equations: the cancellation of soft singularities among the virtual and the unresolved real contributions is automatically satisfied.

When also initial-state collinear singularities have been taken into account, we can rewrite the NLO-accurate expression for the expectation value of O of Eq. (2.16) as

[4] Equation 2.21 leads to a very involved evolution for the PDFs through a system of coupled integro-differential equations, which is very hard to solve. The situation can be significantly simplified by expressing the DGLAP equation in the so-called *evolution basis* [74], instead of the *flavour basis*. In this new basis the splitting matrix can be entirely diagonalized except for a 2x2 block, involving a flavour singlet combination of the quark PDFs and the gluon PDF; the remaining basis elements do not mix with the gluon density and with each others.

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\rm n}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^N \left[\mathcal{B}(\mathbf{\Phi}_{\rm n}) + \alpha_s\, \mathcal{V}_b(\mathbf{\Phi}_{\rm n})\right] O_n(\mathbf{\Phi}_{\rm n}) \\ &+ \int d\mathbf{\Phi}_{\rm n+1}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^{N+1}\, \mathcal{R}(\mathbf{\Phi}_{\rm n+1}) O_{n+1}(\mathbf{\Phi}_{\rm n+1}) \\ &+ \int d\mathbf{\Phi}_{\rm n,\oplus}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^{N+1}\, \mathcal{G}_{\oplus,b}(\mathbf{\Phi}_{\rm n,\oplus}) O_n(\bar{\mathbf{\Phi}}_{\rm n}) \\ &+ \int d\mathbf{\Phi}_{\rm n,\ominus}\, f_{h_1}(x_\oplus) f_{h_2}(x_\ominus)\, \alpha_s^{N+1}\, \mathcal{G}_{\ominus,b}(\mathbf{\Phi}_{\rm n,\oplus}) O_n(\bar{\mathbf{\Phi}}_{\rm n}) \end{aligned} \tag{2.24}$$

where:

$$\begin{aligned} d\mathbf{\Phi}_{\rm n,\oplus} &= dx_\oplus\, dx_\ominus\, dz\, d\Phi_{\rm n}(zk_\oplus + k_\ominus; k_1, \ldots, k_n) \\ d\mathbf{\Phi}_{\rm n,\ominus} &= dx_\oplus\, dx_\ominus\, dz\, d\Phi_{\rm n}(k_\oplus + zk_\ominus; k_1, \ldots, k_n) \end{aligned} \tag{2.25}$$

and $\bar{\mathbf{\Phi}}_{\rm n} = \{\bar{x}_\oplus, x_\ominus, k_1, \ldots, k_n\}$ with $\bar{x}_\oplus = z x_\oplus$ is the *underlying Born* configuration associated to $\mathbf{\Phi}_{\rm n,\oplus}$ (with a similar definition for the one corresponding to $\mathbf{\Phi}_{\rm n,\ominus}$), on which the observable in the last two lines of Eq. (2.24) has been evaluated, due to its infrared safety. In Eq. (2.24) the dependence on $\mu_{\rm F}$ entering the bare collinear counterterms is understood. Just as with $\mu_{\rm R}$, physical observables can not depend on the unphysical scale $\mu_{\rm F}$, which essentially separates emissions included in the PDFs (and therefore unresolved) from the ones participating to the hard scattering. But once the perturbative series is truncated, a residual dependence on $\mu_{\rm F}$ survives and by varying this scale around a central value $\mu_{\rm F,0}$ one can get an estimate of the size of missing higher order terms. Therefore, what is usually done in practise is to set $\mu_{\rm R,0} = \mu_{\rm F,0} = Q$, with Q the characteristic scale of the process, and to simultaneously vary the two scales up and down their central values by a factor of two. If we define $\mu_{\rm R} = K_{\rm R}\mu_{\rm R,0}$ and $\mu_{\rm F} = K_{\rm F}\mu_{\rm F,0}$, the two factors $K_{\rm R}$ and $K_{\rm F}$ take values in the seven-point set:

$$(K_{\rm R}, K_{\rm F}) = \{(2,2), (2,1), (1,2), (1,1), (1,1/2)\,, (1/2,1)\,, (1/2,1/2)\}\ . \tag{2.26}$$

Then, according to this 7-*point scale variation* approach, one obtains an estimate of theoretical uncertainties by taking the maximum and minimum values of the cross section over the seven combinations of scales.

With the requirement for O to be infrared safe and for its Born configuration $O_n(\mathbf{\Phi}_{\rm n})$ to be finite, equation 2.24 is overall finite, even though its individual terms are not. We have described in some details how the cancellation of IR singularities occurs at NLO in QCD. Despite the more intricate structure, the some idea also holds for NNLO calculations. In that case, IR divergences coming from loop integrals in the one-loop squared $|\mathcal{M}_n^{(1)}|^2$ and two-loop $2Re[\mathcal{M}_n^{(0)\star} \cdot \mathcal{M}_n^{(2)}]$ contributions will cancel against the one-loop real term $2Re[\mathcal{M}_{n+1}^{(0)\star} \cdot \mathcal{M}_{n+1}^{(1)}]$, integrated over $\mathbf{\Phi}_{\rm n+1}$, and the double real term $|\mathcal{M}_{n+2}^{(0)}|^2$, integrated over $\mathbf{\Phi}_{\rm n+2}$. By keeping all contributions into account, the final result is finite, but comprises of individually divergent terms.

2.1.3 Subtraction methods

Equation 2.24 can not be practically used in any computer program, since it comprises of individually divergent integrals that only give a finite result when added up together. For this equation to be useable, the IR cancellation of singularities has to be made manifest and separately finite contributions need to be provided to a computer program which performs the multidimensional phase-space integrals over $\mathbf{\Phi}_{\mathrm{n}}$ and $\mathbf{\Phi}_{\mathrm{n+1}}$ via numerical approaches based on Monte Carlo methods.

A first class of methods to achieve that relies on *phase space slicing*. An observable τ sensitive to the kinematics of the radiation is chosen, defined such that the Born kinematics is recovered for $\tau \to 0$. Then, τ is used as a slicing parameter to partition the real phase space in a resolved and an unresolved radiation region:

$$\int d\mathbf{\Phi}_{\mathrm{n+1}}\, \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}}) = \int_0^{\tau_{\mathrm{cut}}} d\tau' \int d\mathbf{\Phi}_{\mathrm{n+1}}\, \frac{d\mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})}{d\tau'} + \int_{\tau_{\mathrm{cut}}} d\tau' \int d\mathbf{\Phi}_{\mathrm{n+1}}\, \frac{d\mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})}{d\tau'}\,. \tag{2.27}$$

In the previous formula, a symbolic notation is used, where PDFs and α_s factors are understood and the τ *spectrum* is defined as:

$$\frac{d\mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})}{d\tau'} = \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})\delta(\tau' - \tau'(\mathbf{\Phi}_{\mathrm{n+1}}))\,. \tag{2.28}$$

The integral over the resolved region $\tau' > \tau_{\mathrm{cut}}$ is finite, since the singularities of the spectrum for $\tau' \to 0$ have been cut away. On the other hand, the intregrand in the unresolved region $0 < \tau' < \tau_{\mathrm{cut}}$ is approximated by making use of its soft-collinear limit ($\mathcal{R}_{\mathrm{sing}}$). If an expression for its integral over the radiation kinematics is known analytically (typically from resummation approaches), it can be combined with the virtual term to give a finite result. We might summarize all that as:

$$\begin{aligned}\int d\mathbf{\Phi}_{\mathrm{n+1}}\, \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}}) &\approx \int_0^{\tau_{\mathrm{cut}}} d\tau' \int d\mathbf{\Phi}_{\mathrm{n+1}}\, \frac{d\mathcal{R}_{\mathrm{sing}}(\mathbf{\Phi}_{\mathrm{n+1}})}{d\tau'} + \int_{\tau_{\mathrm{cut}}} d\tau' \int d\mathbf{\Phi}_{\mathrm{n+1}}\, \frac{d\mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})}{d\tau'} \\ &+ \mathcal{O}(\tau_{\mathrm{cut}}) = \int d\mathbf{\Phi}_{\mathrm{n}} \sigma^{\mathrm{sing}}_{\mathcal{R}}(\mathbf{\Phi}_{\mathrm{n}}, \tau_{\mathrm{cut}}) + \int d\mathbf{\Phi}_{\mathrm{n+1}}\, \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}})\theta(\tau'(\mathbf{\Phi}_{\mathrm{n+1}}) - \tau_{\mathrm{cut}}) + \mathcal{O}(\tau_{\mathrm{cut}})\,.\end{aligned} \tag{2.29}$$

The $\sigma^{\mathrm{sing}}_{\mathcal{R}}(\mathbf{\Phi}_{\mathrm{n}}, \tau_{\mathrm{cut}})$ term, also known as *cumulant*, effectively acts a global term for the subtraction of IR divergences, since its form does not match the singular structure of the real contribution point-by-point in the phase space. This clearly leads to a less efficient cancellation of divergences from a numerical point of view. Moreover, due to the soft-collinear approximation carried out in the unresolved region, Eq. (2.29) is exact up to power corrections $\mathcal{O}(\tau_{\mathrm{cut}})$, that would be captured by the integral of the non-singular spectrum $d\mathcal{R}_{\mathrm{reg}}/d\tau'$ over the unresolved region. The size of these corrections has to be monitored in order to keep the accuracy of the result under control: smaller values of τ_{cut} reduce the size of neglected power corrections, but lead to numerical instabilities due to the non-local nature of the subtraction approach. Therefore, despite the conceptual simplicity, the usage of slicing methods to handle IR divergences can be quite delicate

and local subtraction strategies should be preferred, where available. As we will see, that is the case for NLO calculations, but not for NNLO ones, where the complex structure of the singular regions makes the implementation of simple and efficient local subtraction approaches very challenging. Therefore, in the context of NNLO calculations slicing methods to remove additional singularities not already accounted for at NLO level are still extremely useful. The most common ones are for instance the q_{T} subtraction [75, 76] or the N-jettiness slicing [77–79], where the slicing parameter is set to the transverse momentum of a given colourless system and to the N-jettiness variable [80], respectively. The q_{T} subtraction scheme has been fully automated in the MATRIX framework [81, 82], a parton-level Monte Carlo program where many reactions involving Higgs, vector bosons and top quarks in the final state can be computed at NNLO in QCD (and even NLO EW). This program has been extensively used for many of the results presented in Chapter 4. On the other hand, N-jettiness subtraction has been widely employed in the MCFM program to achieve NNLO accuracy for colour-singlet production [83] (with a more recent extension to q_{T} subtraction, as well, as discussed in Ref. [84]) or in the GENEVA framework [85].

As anticipated, numerically more reliable approaches to make the cancellation of IR divergences explicit exist and are simply named *subtraction* methods. In this case a local subtraction term $\mathcal{C}$ which reproduces the singular behavior of the real amplitude is constructed to remove point-by-point in the phase space its singular poles. Despite this obvious constraint and the requirement for $\mathcal{C}$ to be analytically integrable over the radiation kinematics, the exact form of $\mathcal{C}$ is arbitrary and different choices define different *subtraction schemes*. These methods are nowadays the standard ones for NLO calculations, and in the remainder of this section we examine them in a bit more details. Indeed, besides their importance, that will allow us to set up some useful notation for the following considerations, especially for Section 3.1. To this end, we will mainly follow the structure of the discussion of Ref. [86].

Given the real emission phase space $\mathbf{\Phi}_{\mathrm{n+1}}$, different singular regions α_r can be associated to it, depending on whether a given final-state parton becomes soft (S) or collinear to a final (FSC) or initial (ISC) parton. With {S}, {FSC} and {ISC} we denote the sets of soft, final-state collinear and initial-state collinear singular configurations, respectively. $\mathbf{\Phi}_{\mathrm{n+1}}$ can be mapped to one of these singular configurations by means of a mapping $M^{(\alpha_r)}$

$$\tilde{\mathbf{\Phi}}^{(\alpha_r)}_{\mathrm{n+1}} = M^{(\alpha_r)}(\mathbf{\Phi}_{\mathrm{n+1}})\,, \quad \tilde{\mathbf{\Phi}}^{(\alpha_r)}_{\mathrm{n+1}} = \{\tilde{x}^{(\alpha_r)}_{\oplus}, \tilde{x}^{(\alpha_r)}_{\ominus}, \tilde{k}^{(\alpha_r)}_{1}, \ldots, \tilde{k}^{(\alpha_r)}_{n+1}\} \tag{2.30}$$

having one unresolved final-state parton: a final-state parton with zero four-momentum if $\alpha_r \in \{\mathrm{S}\}$, two final-state partons with parallel three-momenta if $\alpha_r \in \{\mathrm{FSC}\}$, or a final-state parton with a three-momentum parallel to an initial one if $\alpha_r \in \{\mathrm{ISC}\}$. The only requirements on $M^{(\alpha_r)}$ are to be smooth and equal to the identity in correspondence of the singular region: except for that, the exact definition of the singular mapping in Eq. (2.30) is to a large extent arbitrary. A subtraction term $\mathcal{C}^{(\alpha_r)}$ is then associated to each of the singular regions, defined on the remapped real phase space $\tilde{\mathbf{\Phi}}^{(\alpha_r)}_{\mathrm{n+1}}$. According to the subtraction approach, this term is added and subtracted back in Eq. (2.24), in

such a way that, as far as the real contribution is concerned, we have:

$$\int d\mathbf{\Phi}_{\rm n+1}\,\mathcal{L}(x_\oplus,x_\ominus)\,\mathcal{R}(\mathbf{\Phi}_{\rm n+1})O_{n+1}(\mathbf{\Phi}_{\rm n+1}) = \sum_{\alpha_r}\int d\mathbf{\Phi}_{\rm n+1}\big[\mathcal{L}(x_\oplus,x_\ominus)O_n(\mathbf{\Phi}_{\rm n+1})\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\big]_{\alpha_r}$$
$$+\int d\mathbf{\Phi}_{\rm n+1}\left\{\mathcal{L}(x_\oplus,x_\ominus)\,\mathcal{R}(\mathbf{\Phi}_{\rm n+1})O_{n+1}(\mathbf{\Phi}_{\rm n+1}) - \sum_{\alpha_r}\big[\mathcal{L}(x_\oplus,x_\ominus)O_n(\mathbf{\Phi}_{\rm n+1})\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\big]_{\alpha_r}\right\} \tag{2.31}$$

where we have introduced the luminosity function $\mathcal{L}(x_\oplus, x_\ominus) = f_{h_1}(x_\oplus)f_{h_2}(x_\ominus)$ and made use of the context convention $[\dots]_{\alpha_r}$ of Ref. [86] to indicate that all arguments of the functions appearing in the bracket context are evaluated on the singular phase-space region α_r. By construction, the second line of Eq. (2.31) is IR finite, with at most integrable singularities, provided that the definition of O_n suppresses any singular Born-like configuration, as mentioned at the end of Section 2.1.2. The remaining sum over $\mathcal{C}^{(\alpha_r)}$ should be used to cancel the IR singularities of the virtual term. To make the cancellation explicit, the subtraction terms should be analytically integrable over the radiation kinematics. First of all, we can introduce a second mapping $M'^{(\alpha_r)}$

$$M'^{(\alpha_r)}(\mathbf{\Phi}_{\rm n+1}) = \{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)},\,\Phi_{\rm rad}^{(\alpha_r)}\} \quad \text{with } \bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)} = [\{\bar{x}_\oplus,\,\bar{x}_\ominus,\,\bar{k}_1,\,\dots,\,\bar{k}_n\}]_{\alpha_r} \tag{2.32}$$

which expresses the $\mathbf{\Phi}_{\rm n+1}$ phase space in terms of an *underlying Born* $\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}$ part and a three-dimensional phase-space element $\Phi_{\rm rad}^{(\alpha_r)}$, associated to the kinematics of the radiated parton getting unresolved for α_r. The $\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}$ phase space contains n resolved final-state partons, and it is obtained from $\mathbf{\Phi}_{\rm n+1}$ by removing the zero-momentum parton for $\alpha_r \in \{\rm S\}$, by replacing the momenta of the two collinear partons with their momentum sum for $\alpha_r \in \{\rm FSC\}$, and by dropping the parton collinear to an initial-state one, where the momentum fraction of the latter before radiation is replaced by its momentum fraction after radiation, for $\alpha_r \in \{\rm ISC\}$. Momentum conservation is enforced among underlying Born variables:

$$\bar{x}_\oplus^{(\alpha_r)}K_\oplus + \bar{x}_\ominus^{(\alpha_r)}K_\ominus = \sum_{i=1}^{n}\bar{k}_i^{(\alpha_r)}\,. \tag{2.33}$$

The mapping $M'^{(\alpha_r)}$ allows us to factorize the phase-space volume $d\mathbf{\Phi}_{\rm n+1}$ in the first term on the right-hand side of Eq. (2.31) as:

$$d\mathbf{\Phi}_{\rm n+1} = J(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}d\tilde{\Phi}_{\rm rad}^{(\alpha_r)} \equiv d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}d\Phi_{\rm rad}^{(\alpha_r)} \tag{2.34}$$

up to a jacobian factor $J(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})$ conventionally included in $d\Phi_{\rm rad}^{(\alpha_r)} \equiv J(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})d\tilde{\Phi}_{\rm rad}^{(\alpha_r)}$. Then, since for $\alpha_r \in \{\rm S,\,FSC\}$ we have $\bar{x}_\oplus^{(\alpha_r)} = \tilde{x}_\oplus^{(\alpha_r)}$, the luminosity function can be easily factored out from the integral over $\Phi_{\rm rad}$ so as to obtain:

$$\left[\int d\mathbf{\Phi}_{\rm n+1}\mathcal{L}(x_\oplus,x_\ominus)O_n(\mathbf{\Phi}_{\rm n+1})\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\right]_{\alpha_r\in\{\rm S,\,FSC\}} =$$
$$= \int d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}\mathcal{L}(\bar{x}_\oplus^{(\alpha_r)},\bar{x}_\ominus^{(\alpha_r)})O_n(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})\big[\bar{\mathcal{C}}(\bar{\mathbf{\Phi}}_{\rm n})\big]_{\alpha_r\in\{\rm S,\,FSC\}} \tag{2.35}$$

with

$$\left[\bar{\mathcal{C}}(\bar{\mathbf{\Phi}}_{\rm n})\right]_{\alpha_r\in\{\rm S,\,FSC\}} \equiv \left[\int d\Phi_{\rm rad}\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\right]_{\alpha_r\in\{\rm S,\,FSC\}} \tag{2.36}$$

where in Eq. (2.35) we have set $O_n(\tilde{\mathbf{\Phi}}_{\rm n+1}^{(\alpha_r)}) = O_n(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})$ thanks to the infrared safety of O. For $\alpha_r \in \{\rm ISC\}$, one has to account for the fact that $\bar{x}_\oplus^{(\alpha_r)} < \tilde{x}_\oplus^{(\alpha_r)}$ and $\bar{x}_\ominus^{(\alpha_r)} = \tilde{x}_\ominus^{(\alpha_r)}$ for initial-state radiation collinear to the $\oplus$ direction ($\alpha_r \in \{\rm ISC_\oplus\}$) and similarly $\bar{x}_\oplus^{(\alpha_r)} = \tilde{x}_\oplus^{(\alpha_r)}$ and $\bar{x}_\ominus^{(\alpha_r)} < \tilde{x}_\ominus^{(\alpha_r)}$ for initial-state radiation collinear to the $\ominus$ direction ($\alpha_r \in \{\rm ISC_\ominus\}$). Therefore one can write:

$$\begin{aligned}
&\left[\int d\mathbf{\Phi}_{\rm n+1}\mathcal{L}(x_\oplus,x_\ominus)O_n(\mathbf{\Phi}_{\rm n+1})\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\right]_{\alpha_r\in\{\rm ISC_\oplus\}} = \\
&= \int d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)} O_n(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})\left[\int d\Phi_{\rm rad}\frac{dz}{z}\mathcal{L}(x_\oplus,x_\ominus)\mathcal{C}(\bar{\mathbf{\Phi}}_{\rm n})\,z\,\delta(z-\bar{x}_\oplus/\tilde{x}_\oplus)\right]_{\alpha_r\in\{\rm ISC_\oplus\}} = \\
&= \int d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}\frac{dz}{z} O_n(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})\mathcal{L}(\bar{x}_\oplus^{(\alpha_r)}/z,\bar{x}_\ominus^{(\alpha_r)})\left[\bar{\mathcal{C}}(\bar{\mathbf{\Phi}}_{\rm n},z)\right]_{\alpha_r\in\{\rm ISC_\oplus\}},
\end{aligned} \tag{2.37}$$

with

$$\left[\bar{\mathcal{C}}(\bar{\mathbf{\Phi}}_{\rm n},z)\right]_{\alpha_r\in\{\rm ISC_\oplus\}} = \left[\int d\Phi_{\rm rad}\mathcal{C}(\bar{\mathbf{\Phi}}_{\rm n})\,z\,\delta(z-\bar{x}_\oplus/\tilde{x}_\oplus)\right]_{\alpha_r\in\{\rm ISC_\oplus\}} \tag{2.38}$$

with a similar treatment for the case $\alpha_r \in \{\rm ISC_\ominus\}$. Therefore, we get that $\tilde{x}_\oplus^{(\alpha_r)} = \bar{x}_\oplus^{(\alpha_r)}/z$ for $\alpha_r \in \{\rm ISC_\oplus\}$ and $\tilde{x}_\ominus^{(\alpha_r)} = \bar{x}_\ominus^{(\alpha_r)}/z$ for $\alpha_r \in \{\rm ISC_\ominus\}$. By looking at Eq. (2.25), one can easily get convinced that the following identifications hold:

$$d\mathbf{\Phi}_{\rm n,\oplus} = d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}\frac{dz}{z} \text{ for } \alpha_r \in \{\rm ISC_\oplus\} \quad d\mathbf{\Phi}_{\rm n,\ominus} = d\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}\frac{dz}{z} \text{ for } \alpha_r \in \{\rm ISC_\ominus\} \tag{2.39}$$

which allow us to rewrite Eq. (2.24) as

$$\begin{aligned}
\langle O\rangle &= \int d\mathbf{\Phi}_{\rm n}\,\mathcal{L}(x_\oplus,x_\ominus)\,\alpha_s^N\left[\mathcal{B}(\mathbf{\Phi}_{\rm n}) + \alpha_s\,\mathcal{V}_b(\mathbf{\Phi}_{\rm n})\right]O_n(\mathbf{\Phi}_{\rm n}) \qquad (2.40)\\
&+ \int d\mathbf{\Phi}_{\rm n+1}\,\alpha_s^{N+1}\left\{\mathcal{L}(x_\oplus,x_\ominus)\,\mathcal{R}(\mathbf{\Phi}_{\rm n+1})O_{n+1}(\mathbf{\Phi}_{\rm n+1}) - \sum_{\alpha_r}\left[\mathcal{L}(x_\oplus,x_\ominus)O_n(\mathbf{\Phi}_{\rm n+1})\mathcal{C}(\mathbf{\Phi}_{\rm n+1})\right]_{\alpha_r}\right\}\\
&+ \sum_{\alpha_r\in\{\rm S,\,FSC\}}\left[\int d\bar{\mathbf{\Phi}}_{\rm n}\,\alpha_s^{N+1}\,\mathcal{L}(\bar{x}_\oplus,\bar{x}_\ominus)O_n(\bar{\mathbf{\Phi}}_{\rm n})\bar{\mathcal{C}}(\bar{\mathbf{\Phi}}_{\rm n})\right]_{\alpha_r}\\
&+ \int d\mathbf{\Phi}_{\rm n,\oplus}\,\mathcal{L}\left(\frac{\bar{x}_\oplus}{z},\bar{x}_\ominus\right)\alpha_s^{N+1}\left\{\mathcal{G}_{\oplus,b}(\mathbf{\Phi}_{\rm n,\oplus}) + \sum_{\alpha_r\in\{\rm ISC_\oplus\}}\bar{\mathcal{C}}^{(\alpha_r)}(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)},\,z)\right\}O_n(\bar{\mathbf{\Phi}}_{\rm n})\\
&+ \int d\mathbf{\Phi}_{\rm n,\ominus}\,\mathcal{L}\left(\bar{x}_\oplus,\frac{\bar{x}_\ominus}{z}\right)\alpha_s^{N+1}\left\{\mathcal{G}_{\ominus,b}(\mathbf{\Phi}_{\rm n,\ominus}) + \sum_{\alpha_r\in\{\rm ISC_\ominus\}}\bar{\mathcal{C}}^{(\alpha_r)}(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)},\,z)\right\}O_n(\bar{\mathbf{\Phi}}_{\rm n})\,.
\end{aligned}$$

The bare collinear counterterms $\mathcal{G}_{b,\oplus}$ combine with the integrated ISC counterterms to explicitly cancel ISC singularities:

$$\mathcal{G}_{b,\oplus}(\mathbf{\Phi}_{\mathrm{n},\oplus}) + \sum_{\alpha_r \in \{\mathrm{ISC}_\oplus\}} \bar{\mathcal{C}}^{(\alpha_r)}(\bar{\mathbf{\Phi}}_{\mathrm{n}}^{(\alpha_r)}, z) = \mathcal{G}_\oplus(\mathbf{\Phi}_{\mathrm{n},\oplus}) + \delta(1-z)\mathcal{G}_\oplus^{\mathrm{soft}}(\bar{\mathbf{\Phi}}_{\mathrm{n}}) \tag{2.41}$$

where $\mathcal{G}_\oplus(\mathbf{\Phi}_{\mathrm{n},\oplus})$ are the two finite *collinear remnants*, while the remaining divergent piece $\mathcal{G}_\oplus^{\mathrm{soft}}(\bar{\mathbf{\Phi}}_{\mathrm{n}})$ is an endpoint singularity (i.e defined for $z = 1$) of soft origin. The latter combines with the bare virtual contribution, together with the remaining integrated counterterms, to completely account for IR loop divergences:

$$\mathcal{V}(\mathbf{\Phi}_{\mathrm{n}}) = \mathcal{V}_b(\mathbf{\Phi}_{\mathrm{n}}) + \left[\sum_{\alpha_r \in \{\mathrm{FSC,S}\}} \bar{\mathcal{C}}^{(\alpha_r)}(\bar{\mathbf{\Phi}}_{\mathrm{n}}^{(\alpha_r)}) + \mathcal{G}_\oplus^{\mathrm{soft}}(\bar{\mathbf{\Phi}}_{\mathrm{n}}) + \mathcal{G}_\ominus^{\mathrm{soft}}(\bar{\mathbf{\Phi}}_{\mathrm{n}})\right]_{\bar{\mathbf{\Phi}}_{\mathrm{n}}=\mathbf{\Phi}_{\mathrm{n}}} \tag{2.42}$$

where the integrated counterterms have to be evaluated on an underlying Born kinematics matching the Born one point-by-point in the phase space. We can further simplify the notation by including all luminosity factors and α_s powers in the definition of our NLO contributions:

$$\begin{aligned}
&B(\mathbf{\Phi}_{\mathrm{n}}) \equiv \alpha_s^N \mathcal{L}(x_\oplus, x_\ominus)\, \mathcal{B}(\mathbf{\Phi}_{\mathrm{n}}) \qquad V(\mathbf{\Phi}_{\mathrm{n}}) \equiv \alpha_s^{N+1} \mathcal{L}(x_\oplus, x_\ominus)\, \mathcal{V}(\mathbf{\Phi}_{\mathrm{n}}) \\
&R(\mathbf{\Phi}_{\mathrm{n+1}}) \equiv \alpha_s^{N+1} \mathcal{L}(x_\oplus, x_\ominus)\, \mathcal{R}(\mathbf{\Phi}_{\mathrm{n+1}}) \quad C^{(\alpha_r)}(\tilde{\mathbf{\Phi}}_{\mathrm{n+1}}^{(\alpha_r)}) \equiv \alpha_s^{N+1} \mathcal{L}(x_\oplus^{(\alpha_r)}, x_\ominus^{(\alpha_r)})\, \mathcal{C}^{(\alpha_r)}(\tilde{\mathbf{\Phi}}_{\mathrm{n+1}}^{(\alpha_r)}) \\
&G_\oplus(\mathbf{\Phi}_{\mathrm{n},\oplus}) \equiv \alpha_s^{N+1} \mathcal{L}\left(\frac{\bar{x}_\oplus}{z}, \bar{x}_\ominus\right) \mathcal{G}_\oplus(\mathbf{\Phi}_{\mathrm{n},\oplus}) \quad G_\ominus(\mathbf{\Phi}_{\mathrm{n},\ominus}) \equiv \alpha_s^{N+1} \mathcal{L}\left(\bar{x}_\oplus, \frac{\bar{x}_\ominus}{z}\right) \mathcal{G}_\ominus(\mathbf{\Phi}_{\mathrm{n},\ominus})
\end{aligned} \tag{2.43}$$

so that we can finally write an expression for the expectation value of O which is NLO accurate and which can be fed into a proper Monte Carlo integrator:

$$\begin{aligned}
\langle O \rangle = &\int d\mathbf{\Phi}_{\mathrm{n}} \left[B(\mathbf{\Phi}_{\mathrm{n}}) + V(\mathbf{\Phi}_{\mathrm{n}})\right] O_n(\mathbf{\Phi}_{\mathrm{n}}) + \\
&+ \int d\mathbf{\Phi}_{\mathrm{n+1}} \left\{ R(\mathbf{\Phi}_{\mathrm{n+1}}) O_{n+1}(\mathbf{\Phi}_{\mathrm{n+1}}) - \sum_{\alpha_r} \left[O_n(\mathbf{\Phi}_{\mathrm{n+1}}) C(\mathbf{\Phi}_{\mathrm{n+1}})\right]_{\alpha_r} \right\} \\
&+ \int d\mathbf{\Phi}_{\mathrm{n},\oplus} G_\oplus(\mathbf{\Phi}_{\mathrm{n},\oplus}) O_n(\bar{\mathbf{\Phi}}_{\mathrm{n}}) + \int d\mathbf{\Phi}_{\mathrm{n},\ominus} G_\ominus(\mathbf{\Phi}_{\mathrm{n},\ominus}) O_n(\bar{\mathbf{\Phi}}_{\mathrm{n}}) \,.
\end{aligned} \tag{2.44}$$

The formalism that we have recalled in quite some details is common to any subtraction method. As we said, what distinguishes one method from the other is the choice of $C^{(\alpha_r)}$. In the Frixione-Kunst-Signer (FKS) method [87, 88], one rewrites the real amplitude as a sum of contributions associated to only one singular region where one parton gets soft and/or collinear. Then, for each term singularities are handled separately and subtracted by making use of the formalism of plus distributions. This method is the one which is mainly used within the POWHEG framework, and it will be discussed in more details in Section 3.1. In the Catani-Seymour (CS) scheme [61] the counterterms, which are named *dipole functions* in this approach, are built explicitly in such a way to match the real amplitude behaviour in its soft and/or collinear limit. Each dipole function is constructed using the momenta of three partons, named the emitted parton, which is the

one getting unresolved, the emitter and the spectator. In this formalism, multiple dipoles differing in their spectator partons can contribute to the same singular region. The FKS and CS schemes have been the most extensively used for NLO calculations. Nevertheless, many other scheme choices can be made and alternative subtraction schemes have been developed, in particular in view of their extension to the NNLO case, that, as mentioned at the beginning of this section, is a highly non-trivial issue. Just as an example, the antenna subtraction [89–91] is a scheme which is closely related to the CS one. Here subtraction terms are obtained from proper *antenna functions*, which contain all singular configurations for a given parton getting unresolved. Also in this case, an antenna function is defined on three parton momenta, corresponding to the one of the emitted parton and of two hard radiators, which are the partons colour connected to the emitted one. The two hard radiators play the role of the emitter and the spectator of the CS scheme, which are treated symmetrically in the antenna formalism. Antenna subtraction is the method of choice of NNLOJET [92, 93], a parton-level Monte Carlo generator where many inclusive and jet-production cross sections can be computed at NNLO QCD accuracy.

As evidence of the fact that the development of a local NNLO subtraction scheme is still a very active research area, we have the many proposals that have been designed to tackle this problem, each of which has its strengths and weaknesses. Additionally to the NNLO antenna subtraction, we might mention the STRIPPER framework [94], the projection-to-Born method [95], the techniques of Nested Soft-Collinear subtractions [96, 97], the CoLoRFulNNLO method [98–100], or new ideas which are still under development, such as the geometric [101] and the local analytic [102] subtractions, together with many more.

2.2 Parton showers

So far we have described the FO approach to compute the expectation value of an observable O. We have seen that a FO calculation is based on the possibility to truncate the perturbative expansion of the amplitude in Eq. (2.5). Therefore, it is reliable provided that the higher order contributions which are not included in the calculation are negligible. If perturbation theory is applicable, higher order corrections are expected to become smaller and smaller as soon as the order of the α_s contributions we are accounting for increases, as evidence that the perturbative series is starting to converge.

Unfortunately, the stability of the perturbative series can be threatened on many fronts. First of all, in quantum field theory the perturbative expansion of an observable $O = \sum_n O_n \alpha_s^n$ is known to be affected by *renormalon divergences* [103], i.e by a factorial growth of the coefficients $O_n \sim n!$, which should manifest itself at sufficiently high orders. Due to the relatively small number of terms included in current FO calculations, this issue is not visible nowadays, but it effectively sets a limit to the accuracy that can be reached by a simple perturbative approach.

Moreover, FO calculations can provide reasonable predictions for observables that are sufficiently inclusive over QCD radiation: as we have seen in the previous section, in such cases the cancellation of IR divergences between the real and virtual contributions

is complete. Anyway, in more realistic event descriptions, we might want to look at more exclusive observables in particular phase-space regions or we might need to compare theory predictions with experimental measurements where kinematic cuts are applied. Such cuts, that restrict the available phase space, are experimentally required for instance to reduce the background to the measurement of a given process or to account for the geometry of the detectors. All that introduces multiple physical scales Q_i in the problem, additionally to the hard scale Q characteristic of the process. On the theory side, that results in an incomplete cancellation of the singularities between the integral of the real emission term, which is affected by the definition of the observable, and the virtual term, whose loop integral is unconstrained. This manifests itself with the appearence of potentially large logarithmic terms $L \equiv \log(Q_i/Q)$ of the ratio of the physical scales of the problem. Since IR divergences can be of either soft or collinear origin, we can have at most two powers of L multiplying one power of α_s, but such combination will appear at all orders in perturbation theory, with contributions behaving like $\sim (\alpha_s L^2)^n$. Then, if $\alpha_s L^2 \sim 1$, such terms will become equally important at any perturbative order and the reliability of the FO approach will break down. Indeed, to obtain physically sound predictions, these terms have to be accounted for to all orders through some *resummation* techniques.

The accuracy of a resummation procedure for a given observable O can be defined in terms of the cumulative cross section $\Sigma(v)$, giving the probability for $\tilde{O}$ to have a value smaller than v:

$$\Sigma(v) \equiv \int_0^v dv' \frac{d\sigma}{dv'} \,, \tag{2.45}$$

where $\tilde{O}$ is the adimensional version of O defined in such a way that the associated large logarithmic terms can be written as $L \equiv \log(1/v)$. Then, we can express the logarithmically enhanced part of $\Sigma(v)$ for $v \ll 1$ as a double series expansion:

$$\Sigma(v) = \sum_{n=0}^{\infty} \sum_{m=0}^{2n} h_{nm} \alpha_s^n L^m + \mathcal{O}(\alpha_s v) \,. \tag{2.46}$$

The resummed result is said to be leading logarithmic (LL) accurate if all terms with $m = 2n$ are accounted for, next-to-leading logarithmic (NLL) accurate if also terms with $m = 2n - 1$ are resummed, and so on. A different counting of the logarithmic accuracy can also be obtained once the observable O admits an exponential form of the resummation formula:

$$\Sigma(v) = \exp\left[\sum_{n=1}^{\infty} \sum_{m=0}^{n+1} G_{nm} \alpha_s^n L^m\right] + \mathcal{O}(\alpha_s v) \,. \tag{2.47}$$

Despite the fact that Eq. (2.47) represents a more powerful formulation of resummation than Eq. (2.46), it is also less general, since not all known observables present logarithmically enhanced terms whose coefficients are consistent with exponentiation (see for instance Ref. [104]). Indeed, LL contributions in the language of Eq. (2.46) are all

resummed by the first term G_{12} of the exponentiated double series in Eq. (2.47). So, according to Eq. (2.47), N^{k}LL accuracy is reached once terms with $m = n + 1 - k$ are included at all orders.

Analytic calculations based on factorization and on considerations on the behaviour of the matrix element in the IR regime can achieve resummation keeping exactly track of the logarithmic accuracy of the final result. Analytically resummed results have been obtained for instance using the common *b-space* resummation, the more recent *momentum-space* resummation [105] (whose applicability is restricted to transverse observables), or approaches based on soft-collinear effective field theories (SCET) (see for example Ref. [106] for a nice introduction). On the other hand, with these approaches each observable requires a non-trivial dedicated calculation. Semi-analytic methods have also been developed in the past, capable of resumming a relatively broad class of observables meeting the requirements of continuously globalness and recursive collinear and infrared safety. That has been done in the CAESAR framework [104], which was then upgraded to the ARES one [107] for e^+e^- collisions. Despite that, fully numerical methods based on parton shower (PS) simulations, which can be directly embedded into a GPMC program, are nowadays still the most flexible tools to achieve resummation. Indeed, they provide a completely exclusive event description, which is fully differential in the kinematics of the final-state particles. Then, any observable can be evaluated on the set of final-state momenta which are provided at the end of the event simulation. The full generality of the approach clearly comes with the price of less control on the logarithmic accuracy of the final result. Parton showers are known to resum LL contributions, but whether a NLL or a higher accuracy can be claimed for some observables it is often not clear and pretty hard to understand. Some studies to get a better analytic understanding of PS accuracy have been carried out for instance in Ref. [108], with still some ongoing work in trying to develop new NLL PS algorithms, especially within the PanScales collaboration [109–111]. Due to the importance that PS simulations have in this work and in the results presented in Chapter 4, we want to briefly recall some general concepts of PS simulations.

The key idea of PS algorithms [112] is to start from DGLAP equations in Eq. (2.21) for the simulation of splitting processes that lead the system to evolve from hard down to hadronization scales. But first of all, we want to express DGLAP equations in terms of the unregularized splitting functions of Eq. (2.13):

$$\begin{aligned} t\frac{\partial}{\partial t} f_i(x,t) &= \frac{\alpha_s(t)}{2\pi} \sum_j \int_x^1 \frac{dz}{z} P_{ji}(z,\alpha_s(t)) f_j\left(\frac{x}{z},t\right) \\ &= \frac{\alpha_s(t)}{2\pi} \sum_j \left\{ \int_x^{z_{\max}} \frac{dz}{z} \hat{P}_{ji}(z,\alpha_s(t)) f_j\left(\frac{x}{z},t\right) - f_j(x,t) \int_{z_{\min}}^{z_{\max}} dz \hat{P}_{ji}(z,\alpha_s(t)) \right\}, \end{aligned} \tag{2.48}$$

where the last term in the second line of the previous equation originates from dropping the plus prescription in the definition of P_{ji}. We notice that μ_F^2 has been set to a generic energy scale t, that for the time being can be thought of as the (absolute) value of the virtuality of the splitting particle. Moreover, in order to make the integral over $\hat{P}_{ji}$ well-defined, we need to introduce some IR cutoffs $z_{\min}$ and $z_{\max}$ to regularize $\hat{P}_{ji}$

singularities by hand [5]. Then, if we define the so called *Sudakov form factor* $\Delta_i(t, t_c)$:

$$\Delta_i(t, t_c) \equiv \exp\left[-\sum_j \int_{t_c}^{t} \frac{dt'}{t'} \int_{z_{\min}}^{z_{\max}} dz \frac{\alpha_s(t')}{2\pi} \hat{P}_{ji}(z, \alpha_s(t'))\right], \tag{2.49}$$

where t_c is a constant such that $t_c < t$, we can rewrite Eq. (2.48) as:

$$t\frac{\partial}{\partial t}\left(\frac{f_i(x,t)}{\Delta_i(t,t_c)}\right) = \frac{1}{\Delta_i(t,t_c)} \frac{\alpha_s(t)}{2\pi} \sum_j \int_x^{z_{\max}} \frac{dz}{z} \hat{P}_{ji}(z, \alpha_s(t)) f_j(x/z, t)\,, \tag{2.50}$$

or in an equivalent integral formulation of the Fredholm type:

$$f_i(x,t) = f_i(x,t_c)\Delta_i(t,t_c) + \sum_j \int_{t_c}^{t} \frac{dt'}{t'} \frac{\Delta_i(t,t_c)}{\Delta_i(t',t_c)} \int_{z_{\min}}^{z_{\max}} \frac{dz}{z} \frac{\alpha_s(t')}{2\pi} \hat{P}_{ji}(z, \alpha_s(t')) f_j\left(\frac{x}{z}, t'\right). \tag{2.51}$$

Written in this form, DGLAP equations have an appealing probabilistic interpretation which is at the heart of PS algorithms. In the first term of Eq. (2.51), where f_i is evaluated at t_c, the PDF flavour i is unchanged, which means there is no evolution. On the other hand, in the second term of Eq. (2.51) a splitting process turning the flavour i into j takes place. That suggests to interpret $\Delta_i(t, t_c)$ as the probability for no splitting to occur from scale t_c to t for the parton line i, and the ratio $\Delta_i(t, t_c)/\Delta_i(t', t_c)$ as the probability of having no splitting from t' to t. It becomes clear that the scale t can be interpreted as an evolution (or ordering) variable. In the context of DGLAP evolution, a parton coming out of a hadron with a probability described by a proper PDF evolves from a low scale t_c, with a virtuality $\bar{t}_c = -t_c$, towards the hard scale $t_0 > t_c$, so participating to the hard process with a more negative virtuality $\bar{t}_0 = -t_0$. Evolution from lower to higher scales is characteristic of the so called *space-like showers*. When considering parton evolution from higher to lower scales, as it occurs for radiation from final states, we refer to the PS as *time-like*. In what follows, we will mainly consider the time-like evolution, where it is simpler to introduce many of PS underlying ideas and we will discuss the space-like case further in due time. In the context of time-like evolution, $\Delta_i(t, t_c)$ can be read as the probability for no splitting to take place from the scale t down to t_c and $\Delta_i(t, t_c)/\Delta_i(t', t_c)$ as the probability for no splitting to happen from t to t' (notice the direction of evolution compared to the space-like case).

It is evident that the two constants t_0 and t_c define the boundaries of the PS evolution: t_0 refers to its starting scale and is typically set to the hard scale of the process, while t_c is its final (or cutoff) scale. The value of t_c can either be the hadronization scale, if the PS generates radiation for partons directly coming out of the hard production process, or the width of an unstable particle, whose decay products are dressed with PS radiation (in this case the period of PS radiation cannot exceed the lifetime of the resonance particle [2]).

[5] As also discussed in the fifth chapter of Ref. [113], the exact relation between regularized and unregularized splitting functions is mathematically recovered in the limits $z_{\max} \to 1$ and $z_{\min} \to 0$. Therefore, the second line of Eq. (2.48) is a valid alternative formulation of DGLAP evolution equations provided that the IR cutoffs $z_{\max}$ and $z_{\min}$ do not depart too much from these limits.

The evolution in DGLAP equations is based on AP functions, which emerge from the collinear limit of the matrix element in Eq. (2.12). For the generation of multiple splittings, one can simply iterate the formula by writing the amplitude as a product of a squared matrix element for the generation of a given number of resolved partons and multiple AP splitting factors accounting for all collinear emissions. This approximation effectively neglects interference effects at the level of the squared matrix element, which would arise from the possibility for a radiated parton to be collinear to multiple final-state partons. These contributions are highly suppressed in the kinematic regime of *strongly ordered emissions* $t_0 > t_1 > \cdots > t_n > t_c$: each branching in the PS cascade can be treated as an independent $1 \to 2$ splitting [6] just providing the starting kinematic conditions for the subsequent splitting process (in this sense PS algorithms are *Markov chains*). The strong ordering condition is a fundamental requirement for PS to work. Moreover, strongly ordered configurations are the most singular ones responsible for the highest power of large logarithms, therefore accounting for terms like $\alpha_s^n \log^n(t_0/t_c)$ for n splitting processes.

In the formalism described by Eq. (2.51), named *parton branching method*, the kinematics of a splitting process is controlled by the term:

$$\frac{d\phi}{2\pi}\frac{dt'}{t'}dz\,\Delta_i(t,t')\frac{\alpha_s(t')}{2\pi}\hat{P}_{ji}(z,\alpha_s(t')) \quad \text{with} \quad \Delta_i(t,t') \equiv \frac{\Delta_i(t,t_c)}{\Delta_i(t',t_c)}\,. \tag{2.52}$$

Note that, when considering time-like showers for final-state radiation, there is a missing $1/z$ term in the previous formula compared to Eq. (2.51): this extra factor is required to obtain the correct splitting probability when PDFs are involved, as for initial-state radiation (as evident from Eq. (2.53) later in this section). Since the AP splitting functions which are commonly used in PS algorithms are spin-averaged, the azimuthal angle ϕ around the direction of the splitting parton is uniformly generated [7] within the range $[0,\, 2\pi[$. Then, the hardness t' at which the splitting occurs is generated according to the Sudakov factor $\Delta_i(t,t')$. By noticing that $\Delta_i(t,t') \in]0,1]$, where the extremes are approached respectively for $t' \to t_c$ and $t' = t$, we can generate a uniformly distributed random number $r \in [0,1]$ and solve the equation $r = \Delta_i(t,t')$ for t'. A new resolvable branching takes place at t' if $t' > t_c$; otherwise the parton line i is not allowed to split any longer [8]. Finally, the parton energy fraction z is generated according to $\hat{P}_{ji}(z,\alpha_s(t'))$. This simple parton shower algorithm was first proposed in Ref. [120], and allows to populate the final-state phase space with an arbitrary number of QCD partons.

The factorized form of the splitting process $i \to j$ described by $\hat{P}_{ji}(z,\alpha_s(t))$ just works in the strictly collinear limit, where the virtuality m_i^2 of the splitting particle i is zero.

[6] The possibility to improve parton showers by including $1 \to 3$ splittings, where three partons get simultaneously collinear, has been explored in some works, like Ref. [114].

[7] In more realistic branchings, a correlation between the polarization of the splitting parton and the plane of branching exists. An algorithm to fully account for these spin correlation effects was proposed in Ref. [115–118] and is known with the name of *Collins-Knowles algorithm*, which was later adopted by the Herwig PS program [119].

[8] Similarly, in a space-like shower evolving from t_c to t_0 (*forward evolution*), the splitting scale t' is obtained by solving the equation $r = \Delta_i(t,t')$, where now $\Delta_i(t,t') \equiv \Delta_i(t',t_c)/\Delta_i(t,t_c)$, and where the splitting is accepted provided that $t' < t_0$.

In the PS language, this means that values of t different from zero are not allowed and that only final-state configurations with partons having zero angular separation can be described. This fact for instance would prevent GPMC generators based on PS algorithms to account for the momentum distribution within collimated bunches of final-state hadrons (named *jets*), which is observed experimentally. To remedy this problem, on-shell splitting particles are allowed to get a non-zero virtuality provided that momentum conservation is enforced via a *momentum reshuffling* procedure. By doing that, the recoil of the splitting process is either absorbed by a spectator parton (*local recoil*) or more generally distributed among the other particles of the process (*global recoil*). The exact way this is done defines different *recoil schemes*. The impact of different choices is beyond the LL accuracy that PS aims to reach by numerically resumming logarithmically enhanced terms in the Sudakov form factor of Eq. (2.65). Nevertheless, recoil schemes are known to play an important role when trying to move to higher logarithmic accuracy.

Different PS algorithms are defined not just by the choice of a recoil scheme, but also by the exact definition of the ordering variable t, which might differ from the virtuality of the splitting particle. As soon as $t = h(z)\, m_i^2$, with h a quite general function of z, the LL accuracy of the shower is preserved by different t choices, as a consequence of the fact that the singular behaviour encoded in the factor dt/t of Eq. (2.52) will not be affected [121]. Once t has been defined, the extremes of the z integration $z_{\max}$ and $z_{\min}$ are fixed by kinematic constraints, if t is chosen to be the virtuality or the transverse momentum of the splitting pair, while z remains unconstrained if t is set for instance to the angle formed by the splitting products. In the latter case, a value for $z_{\min}$ has to be manually set to prevent the PS evolution to hit the Landau pole of α_s. Moreover, the definition of z itself is to some extent arbitrary, as soon as it preserves its meaning of momentum fraction in the collinear limit.

Finally, the exact form of the splitting kernels in different PS algorithms can also deviate from the AP functions. This is even better understood if we consider that by now we have only described the problem of resumming collinear radiation and built the PS algorithm around the collinear limit of the squared matrix elements. Indeed, when trying to account for soft radiation, the situation is more delicate. As evident from Eq. (2.14) and Eq. (2.15), a soft emission can not be easily thought of as a $1 \to 2$ splitting process, since intereference effects play a crucial role there. Therefore, different strategies are in use to account for soft emissions, at least in an approximate way. A first class of approaches is based on angular ordering (whose derivation is nicely presented in Ref. [1]). A gluon emitted at an angle θ_a which is larger than the angular separation θ_b between the two partons involved in the previous splitting step can not resolve the distance among the two partons of the pair: upon azimuthal average, the amplitude for such a configuration is zero. It turns out that contributions to the emission process just arise from gluons within the cone defined by the previous pair, i.e for $\theta_a < \theta_b$. Phase-space configurations which violate angular ordering are therefore strongly suppressed. That suggests a way to account for large-angle soft radiation, which makes use of the angle of the splitting process as an evolution variable t for the shower. This is the default choice of all HERWIG parton shower programs [119, 122]. In the PYTHIA program [123, 124], the virtuality or

the transverse momentum of the splitting pair is used as t, but emissions whose angles do not respect the angular ordering prescription are vetoed. In a second class of methods, named *dipole-based* showers [125], the main idea is to build the PS algorithm around the soft approximation of the squared matrix element and so to conceive the splitting intrinsically as a $2 \to 3$ process involving colour-connected partons. Then, in an inverted logic compared to AP-based showers, hard collinear emissions are accounted for via correction factors. Dipole shower algorithms are implemented and used by default in many PS programs like VINCIA [126], ARIADNE [127], and SHERPA [128, 129].

Even thought in our discussion so far we have referred to the generation of radiation from final-state partons, an algorithm for initial-state radiation has also been developed long ago in Ref. [130]. In this approach, instead of starting the PS cascade from the scale t_c at which a parton from a given PDF is resolved, and of evolving the system till it reaches the hard scale t_0, the evolution proceeds in the opposite direction. It starts from the initial-state partons entering the hard process and then evolves them backward, by computing the probability for a parton to have emerged from a splitting process. This strategy, dubbed *backward evolution*, is preferably applied to space-like showers, since it allows for a more efficient event generation, where the PS is built around the hard process. Indeed, any algorithm starting the evolution from partons emerging directly from PDFs would be extremely inefficient, since the cross section for processes which are relevant for LHC searches are order of magnitudes smaller compared to the total hadronic cross section and very few events with a parton having exactly the energy to contribute to the hard process will be generated in this way (see for instance discussion in Ref. [131]). One of the main differences when moving from a time-like to a backward-evolved space-like shower for initial-state radiation is the appropriate modification of the Sudakov form factor, which now reads:

$$\Delta_i^{\rm ISR}(t, t_c; x) \equiv \exp\left[-\sum_j \int_{t_c}^{t} \frac{dt'}{t'} \int_{z_{\rm min}}^{z_{\rm max}} \frac{dz}{z} \frac{\alpha_s(t')}{2\pi} \hat{P}_{ji}(z, \alpha_s(t)) \frac{f_i(x/z, t')}{f_j(x, t')}\right], \tag{2.53}$$

where the PDF ratio guarantees that the backward evolution is done consistently with DGLAP equations [9]. Furthermore, the inclusion of PDFs in $\Delta_i^{\rm ISR}(t, t_c; x)$ automatically enforces momentum conservation and effectively leads to the contraint $z_{\rm min} \geq x$, with x the momentum fraction of the parton which might have been produced from a possible splitting at earlier times.

[9] In Ref. [113] it was discussed that backward evolution as implemented in many PS programs can potentially be inconsistent with DGLAP evolution. With some manipulations, Eq. (2.51) can be rewritten as

$$\Delta_i(t, t') = \frac{f_i(x, t')}{f_i(x, t)} \Delta_i^{\rm ISR}(t, t'; x)\,, \tag{2.54}$$

which provides a strong consistency condition for PS evolution. Indeed, the right-hand side of Eq. (2.54) has to be x-independent. In Ref. [113], it is shown that many PS programs violate this x-independence requirement. That is a consequence of enforcing four-momentum conservation in the shower evolution, which leads to values of the IR cutoffs $z_{\rm max}$ and $z_{\rm min}$ significantly away from one and zero, respectively (see discussion in the footnote 5 of this section). Some attempts in trying to incorporate this requirement in PS programs have been pursed for instance in Ref. [132].

As we mentioned many times, PS methods are expected to resum large logarithmic terms at LL accuracy. Despite that, NLL accuracy might still be reached for some observables: this is possible because some higher-order and NLL contributions are indeed included in the PS evolution. When computing next-to-leading order corrections to the AP functions, there are two kinds of universal contributions which arise from the soft limit $z \to 1$ of the NLO splitting kernel $\hat{P}_{ji}$ (see for instance Ref. [133]):

$$\alpha_s(m_i) \cdot \lim_{z\to 1} \hat{P}_{ji}(z, \alpha_s(m_i)) = \alpha_s(m_i) \frac{2C_i}{1-z} \left[1 - \alpha_s(m_i) \left(\beta_0 \log(1-z) + \frac{K_g}{2\pi} \right) \right] + \mathcal{O}(\alpha_s^3) \,, \tag{2.55}$$

with

$$K_g = C_A \left(\frac{67}{18} - \frac{\pi^2}{6} \right) - T_R n_f \frac{10}{9} \,, \tag{2.56}$$

and where α_s has been evaluated at the virtuality m_i of the splitting particle. Clearly, $C_i = C_F$ for $q \to qg$ splittings and $C_i = C_A$ for $g \to gg$ splittings. The first $\mathcal{O}(\alpha_s^2)$ term of Eq. (2.55) arises from the incomplete cancellation between the virtual and real corrections [10] to the AP function. Since the transverse momentum of the splitting process k_T can be related to the virtuality m_i via the relation $k_\mathrm{T}^2 \approx z(1-z)m_i^2$, this kind of contributions can be accounted for to all orders by evaluating α_s at k_T instead of m_i. It is worth noting that, for $z \to 1$, $k_\mathrm{T}^2 \to (1-z)m_i^2$, as required by the argument of the logarithm in Eq. (2.55). But by making use of k_T, the singular behaviour for $z \to 0$ typical of the $g \to gg$ splitting is also captured. These corrections which are resummed by a proper choice of the renormalization scale for α_s are crucial for a meaningful description of events: despite entering at higher orders, their numerical impact can be large, being LL contributions. The second α_s^2 correction in Eq. (2.55) is proportional to K_g, named the *two-loop cusp anomalous dimension*, and originates from the finite reminder of the double real correction to the AP function. This class of genuinely NLL terms is usually included by making use of the Catani-Marchesini-Webber (CMW) or Bremsstrahlung scheme for α_s, i.e $\alpha_s^{\mathrm{CMW}}(k_\mathrm{T}) = \alpha_s(k_\mathrm{T}) + \alpha_s^2(k_\mathrm{T}) K_g/(2\pi)$ [134].

Even though Eq. (2.50) and Eq. (2.51) already provide many of the ingredients entering PS algorithms, we saw that PS evolution differs in many respects from the one described by DGLAP equations. The latter are derived in the strictly collinear limit and just keep track of one of the particles involved in the splitting process. One can actually formalized the definition of PS, whose underlying idea we have shortly illustrated above, in terms of an operator $\mathcal{S}$ acting on the space of particle configurations (see for instance Ref. [135, 136]). A state of n particles with momenta $k_1, \ldots, k_n$ and quantum numbers (such as flavour, colour and spin) $\rho_1, \ldots, \rho_n$ can be denoted by $|k_1, \rho_1; \ldots; k_n, \rho_n)$, with a normalization:

$$(k_1, \rho_1; \ldots; k_n, \rho_n | k'_1, \rho'_1; \ldots; k'_{n'}, \rho'_{n'}) = \delta_{n,n'} \prod_{i=1}^{n} \delta^3(k_i - k'_i) \delta_{\rho_i, \rho'_i} \,. \tag{2.57}$$

[10] More precisely, this type of correction originates only from the combination of the gluon self-energy and the double real contributions arising from the gluon splitting process [104].

Then, the shower $\mathcal{S}$ can be seen as a sum over all possible configurations, weighted by a probability factor S:

$$\mathcal{S} = \sum_{n=1}^{+\infty} \sum_{\rho_1,\ldots,\rho_n} \int d^3k_1 \ldots d^3k_n S(k_1,\rho_1;\ldots;k_n,\rho_n)(k_1,\rho_1;\ldots;k_n,\rho_n| \,. \tag{2.58}$$

In this formalism, an observable O can be defined in terms of its values O_n over the different particle configurations:

$$O = \sum_{n=1}^{+\infty} \sum_{\rho_1,\ldots,\rho_n} \int d^3k_1 \ldots d^3k_n O_n(k_1,\rho_1;\ldots;k_n,\rho_n)(k_1,\rho_1;\ldots;k_n,\rho_n| \,, \tag{2.59}$$

in such a way that its expectation value will be simply given by $\langle O \rangle = \mathcal{S} \cdot O$. The Markov chain nature of PS algorithms is enforced by requiring $\mathcal{S}$ to satisfy a recursive equation (we just refer here to time-like showers, but all of the discussion can also be adapted to the space-like case). If we denote by $\mathcal{S}_i(t)$ the PS evolving the parton line i from scale t, then the recursion relation reads:

$$\begin{aligned} \mathcal{S}_i(t) &= \Delta_i(t,t_c)\mathcal{S}_i(t_c) + \\ \sum_{(jl)} \int_{t_c}^{t} \frac{dt'}{t'} \int_0^{2\pi} \frac{d\phi}{2\pi} \int_{z_{\min}}^{z_{\max}} dz \frac{\alpha_s(t')}{2\pi} \hat{P}_{i,jl}(z,\alpha_s(t')) &\Delta_i(t,t')\mathcal{S}_j(\tilde{t}(t',z))\mathcal{S}_l(\tilde{t}(t',1-z)) \,. \end{aligned} \tag{2.60}$$

The first term just accounts for the probability of no further splittings down to the cutoff scale t_c, leaving the initial state $\mathcal{S}_i$ untouched. The second term describes the probability of having no splitting till a scale t' with the Sudakov factor $\Delta_i(t,t')$ and then a splitting $i \to jl$ (a sum over all possible splittings (jl) is carried out) which initiates two new sub-showers $\mathcal{S}_j$ and $\mathcal{S}_l$. The starting scales $\tilde{t}$ for the two new showers $\mathcal{S}_j$ and $\mathcal{S}_l$ are functions of the momentum fractions z and $1-z$, respectively, but the exact dependence of $\tilde{t}$ on t' is clearly affected by the definition of the ordering variable. We notice that the exclusive AP splitting functions $\hat{P}_{i,jl}$ have been used in the equation: since they keep track of both particles of the splitting pair, their definitions slighlty differ from the inclusive AP ones (see Ref. [136]). From Eq. (2.60) one can also derive an evolution equation for the PS operator:

$$\begin{aligned} t\frac{\partial \mathcal{S}_i(t)}{\partial t} =& \sum_{(jl)} \int_0^{2\pi} \frac{d\phi}{2\pi} \int_{z_{\min}}^{z_{\max}} dz \frac{\alpha_s(t)}{2\pi} \hat{P}_{i,jl}(z,\alpha_s(t)) \mathcal{S}_j(\tilde{t}(t,z))\mathcal{S}_l(\tilde{t}(t,1-z)) \\ &- \left[\sum_{(jl)} \int_{z_{\min}}^{z_{\max}} dz \frac{\alpha_s(t)}{2\pi} \hat{P}_{i,jl}(z,\alpha_s(t)) \right] \mathcal{S}_i(t) \,, \end{aligned} \tag{2.61}$$

which shows how increasing the evolution scale t makes it more likely for the shower to split in two sub-showers than to remain the same. The latter behaviour is instead the one described by the second term of the equation, originating from the derivative acting on the Sudakov exponent.

Moreover, from the recursive relation of Eq. (2.60) one can derive a strong formulation of the conservation of probability (also named *shower unitarity*). Since the sum of probability factors over all shower configurations has to be normalized to one, i.e:

$$\sum_{n=1}^{+\infty} \sum_{\rho_1,\ldots,\rho_n} \int d^3k_1 \ldots d^3k_n S(k_1,\rho_1;\ldots;k_n,\rho_n) = 1\,, \tag{2.62}$$

we can apply this sum to both sides of Eq. (2.60), so as to obtain:

$$1 = \Delta_i(t,t_c) + \sum_{(jl)} \int_{t_c}^{t} \frac{dt'}{t'} \int_0^{2\pi} \frac{d\phi}{2\pi} \int_{z_{\min}}^{z_{\max}} dz \frac{\alpha_s(t')}{2\pi} \hat{P}_{i,jl}(z,\alpha_s(t'))\Delta_i(t,t')\,, \tag{2.63}$$

which simply states how the probabilities of having no resolved splitting and one resolved splitting have to sum up to one. As a consequence, PS algorithms do not change the total inclusive cross section of a process, which is properly recovered when integrating over the details of the final-state kinematics. Moreover, the unitarity principle in Eq. (2.63) encodes the KLN theorem, which manifests itself in the requirement for the IR singularities to cancel to all orders between the unresolved and resolved contributions. It is worth mentioning that, in order for unitarity to work, the integrand of the second term of Eq. (2.63) has to be an exact differential, which means that the exponent of the Sudakov factor has to match the splitting kernel multiplying it.

2.3 Combining fixed-order and parton shower predictions

In Sec. 2.1 we have described the main ingredients entering fixed-order calculations: they are based on the exact form of the matrix element at a given perturbative order, in such a way that all quantum interference effects are taken into account. Their predictive power is anyway restricted to sufficiently inclusive observables, since in exclusive event descriptions the convergence of the perturbative series is threaten by large logarithmic contributions. Moreover, the amount of final-state particles n that can be described by this exact approach is limited by the rise of the calculation complexity as soon as n increases. On the other hand, no such limitation exists for PS algorithms, whose basics have been recalled in Sec. 2.2. PS can generate in principle an arbitrary final-state multiplicity and can resum to all orders enhanced logarithmic contributions causing FO predictions to fail. The flexibility of PS tools comes with the price of a lower and hard-to-control accuracy: their applicability is limited to the soft and/or collinear regime, due to the approximation they are based on. Moreover, despite the ongoing effort in trying to incorporate quantum interference [109, 137–139], together with spin [140–142]

and colour [11] [144–147] correlation, these effects are only included in an approximate way.

From these considerations, it emerges quite clearly how the FO and PS approaches are complementary: the former should be better used to describe hard radiation, and so the high-energy core of the process, while the latter is more suitable for the description of additional soft and/or collinear radiation. Therefore, any realistic event simulation of LHC processes, where disparate energy scales are probed, will benefit from the combination of the two methods. Unfortunately, many subtleties are involved in a possible combination procedure, which has to guarantee a smooth transition between the FO- and PS-dominated regimes and to avoid double-counting issues, by preventing radiation contributions from being included multiple times in the kinematic regions where the FO and the PS descriptions overlap. Finding the best way of combining the two approaches is a very active and broad research field by itself in the context of high-precision phenomenology. Two main orthogonal directions for this combination can be outlined, namely *matching*, described in Subsection 2.3.1, and *merging*, that we present in Subsection 2.3.2. In what follows we are just recalling the main ideas and phylosophy of these two combination strategies, without claiming to give a comprehensive description of the vast variety of methods belonging to them.

2.3.1 Matching at NLO

The first attempt to improve the PS description of hard-kinematics regions has been the *matrix element correction* (MEC) approach [148], where the hardest PS emission is corrected by the full LO matrix element. This is essentially done by generating PS radiation in the usual way, with a probability controlled by the PS splitting kernels, and then by accepting the branching with a probability given by the ratio of the LO matrix element over the PS kernel itself (a procedure known as *veto algorithm*, nicely described for instance in Ref. [149]). Despite its simplicity, extending this method to any parton multiplicity so as to correct an arbitrary number of emissions for the tree-level matrix elements is very involved. Some attempts in this direction have been pursed in Ref. [150] and made available in the Vincia antenna shower model [126].

Matching strategies start from a different perspective. Their main goal is trying to improve PS methods in such a way that the inclusive cross section is recovered at NLO accuracy. In order to reach the desired NLO normalization, the LO matrix element has to be supplemented by the virtual and real contributions. On the other hand, the LL accuracy of the PS has to be preserved. Therefore, one has to make sure to avoid double-counting issues, since the real matrix element correctly reproduces the distribution of radiation in both the hard and soft/collinear regime, the latter already accounted

[11] Among other approximations, standard parton showers keep track of the colour structure of the event by working in the *large-N_c limit*, where N_c refers to the number of colours. This limit is often used also in many FO calculations, where taking $N_c \to \infty$ drastically simplifies the treatment of the QCD colour algebra, allowing to neglect terms like $\mathcal{O}(1/N_c^n)$ with $n \geq 1$ (*leading colour approximation*). In MC programs, making use of the large-N_c limit enables to evolve the colour flow across the parton cascade, from parton level all the way down to the non-perturbative regime. Indeed, the colour information is an indispensible input for all hadronization models (see for instance introductory discussion in Ref. [143]).

for by the PS. It is also clear that, as a spinoff of the methods, matching approaches also correct the hardest radiation for the tree-level matrix elements, similarly to MEC approaches.

Even though different variants of matching strategies exist, we can still describe the main idea in a quite general way [85, 151]. We can define the quantity

$$\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n}) = B(\mathbf{\Phi}_{\rm n}) + V(\mathbf{\Phi}_{\rm n}) + \left[\int d\Phi_{\rm rad}\left\{R_{\rm S}(\mathbf{\Phi}_{\rm n+1}) - \sum_{\alpha_r} C^{(\alpha_r)}(\mathbf{\Phi}^{(\alpha_r)}_{\rm n+1})\right\} + \int \frac{dz}{z}\left\{G_{\oplus}(\mathbf{\Phi}_{\rm n,\oplus}) + G_{\ominus}(\mathbf{\Phi}_{\rm n,\ominus})\right\}\right]_{\bar{\mathbf{\Phi}}_{\rm n}=\mathbf{\Phi}_{\rm n}}, \tag{2.64}$$

where $R_{\rm S}$ contains all IR singular contributions of the real amplitude R, in such a way that $R - R_{\rm S}$ is IR finite. Therefore, Eq. (2.64) corresponds to the inclusive cross section (up to the integration over the real kinematics of the $R - R_{\rm S}$ contribution) at fixed Born kinematics $\mathbf{\Phi}_{\rm n}$. We can also define a modified Sudakov form factor

$$\Delta_{\rm match}(\mathbf{\Phi}_{\rm n}, t) \equiv \exp\left\{-\int \frac{[d\Phi_{\rm rad} R_{\rm S}(\mathbf{\Phi}_{\rm n+1})\,\theta(t(\mathbf{\Phi}_{\rm n+1}) - t)]_{\bar{\mathbf{\Phi}}_{\rm n}=\mathbf{\Phi}_{\rm n}}}{B(\mathbf{\Phi}_{\rm n})}\right\}, \tag{2.65}$$

which exponentiates the singular real contributions. The variable t is a hardness variable, which might or might not match the PS evolution variable, and the θ-condition in the integrand meets the usual PS ordering condition.

With this notation, we can write a quite general expression for matching NLO and PS calculations in terms of the fully differential cross section $d\sigma$ as:

$$\begin{aligned} d\sigma =& \bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Delta_{\rm match}(\mathbf{\Phi}_{\rm n}, t_c) d\mathbf{\Phi}_{\rm n} + \left\{\int d\Phi_{\rm rad}(R(\mathbf{\Phi}_{\rm n+1}) - R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))\,\theta(t_c - t_\Phi)\right\} d\mathbf{\Phi}_{\rm n} \\ &+ \bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Delta_{\rm match}(\mathbf{\Phi}_{\rm n}, t_\Phi)\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}\,\theta(t_\Phi - t_c) d\mathbf{\Phi}_{\rm n+1} \\ &+ (R(\mathbf{\Phi}_{\rm n+1}) - R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))\,\theta(t_\Phi - t_c) d\mathbf{\Phi}_{\rm n+1}\,, \end{aligned} \tag{2.66}$$

where we have set $t_\Phi \equiv t(\mathbf{\Phi}_{\rm n+1})$. The first term in Eq. (2.66) is responsible for the generation of events with Born-like kinematics, where the Sudakov factor $\Delta_{\rm match}(\mathbf{\Phi}_{\rm n}, t_c)$ suppresses any further radiation down to the cutoff scale t_c. Born-like events are distributed according to $\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})$ and are usually generated by a Monte Carlo program with the standard hit-and-miss technique. The third term of Eq. (2.66) generates the first emission at scales $t_\Phi > t_c$. The combination of the first and third terms, once the $\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})$ function is factored out, exactly reproduces the PS structure for the generation of radiation, with the ratio $R_{\rm S}(\mathbf{\Phi}_{\rm n+1})/B(\mathbf{\Phi}_{\rm n})$ replacing the PS splitting kernel. Since this ratio exactly matches the exponent of $\Delta_{\rm match}(\mathbf{\Phi}_{\rm n}, t_\Phi)$, the radiation kinematics can be generated in a PS fashion, i.e relying on the veto algorithm. The second term of Eq. (2.66) (the FO matching term) has Born-like kinematics and it is needed to compensate for the replacement in $\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})$ of R with its singular approximation $R_{\rm S}$ in the region $t_\Phi < t_c$, where the IR cancellation of singularities should occur. Finally, the last term of Eq. (2.66) generates the first radiation away from the soft/collinear limit. In this latter case, the

radiation kinematics is distributed according to $R - R_{\rm S}$ and can be generated again by hit-and-miss methods (this term effectively amounts to a matrix element correction). Once the first radiation kinematics has been generated according to the matching formula, extra radiation with hardness t lower than t_Φ can be handled by the PS.

It is important to show that Eq. (2.66) correctly reproduces NLO accurate results for sufficiently inclusive observables. We can compute the expectation value $\langle O\rangle$ as:

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\rm n}\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Big\{\Delta_{\rm match}(\mathbf{\Phi}_{\rm n},t_c)O_n(\mathbf{\Phi}_{\rm n})\\ &+\int_{t_\Phi>t_c} d\Phi_{\rm rad}\Delta_{\rm match}(\mathbf{\Phi}_{\rm n},t_\Phi)\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}O_{n+1}(\mathbf{\Phi}_{\rm n+1})\Big\}\\ &+\int d\mathbf{\Phi}_{\rm n}\int d\Phi_{\rm rad}(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))\,\theta(t_c-t_\Phi)O_{n+1}(\mathbf{\Phi}_{\rm n+1})\\ &+\int d\mathbf{\Phi}_{\rm n+1}(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))\,\theta(t_\Phi-t_c)O_{n+1}(\mathbf{\Phi}_{\rm n+1})\,. \end{aligned}\tag{2.67}$$

The last two lines in the previous expression naively combine. Then, we can rewrite the previous equation by adding and subtracting a term proportional to $O_n(\mathbf{\Phi}_{\rm n})$ as:

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\rm n}\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Big\{\Delta_{\rm match}(\mathbf{\Phi}_{\rm n},t_c)+\int_{t_\Phi>t_c} d\Phi_{\rm rad}\Delta_{\rm match}(\mathbf{\Phi}_{\rm n},t_\Phi)\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}\Big\}O_n(\mathbf{\Phi}_{\rm n})\\ &+\int d\mathbf{\Phi}_{\rm n}\int_{t_\Phi>t_c} d\Phi_{\rm rad}\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Delta_{\rm match}(\mathbf{\Phi}_{\rm n},t_\Phi)\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}(O_{n+1}(\mathbf{\Phi}_{\rm n+1})-O_n(\mathbf{\Phi}_{\rm n}))\\ &+\int d\mathbf{\Phi}_{\rm n+1}(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))O_{n+1}(\mathbf{\Phi}_{\rm n+1})\,. \end{aligned}\tag{2.68}$$

Due to shower unitarity (see Eq. (2.63)), the term in curly brackets in the first line simply evaluates to one. Moreover, the difference $O_{n+1}(\mathbf{\Phi}_{\rm n+1})-O_n(\mathbf{\Phi}_{\rm n})$ suppresses singular contributions from $R_{\rm S}(\mathbf{\Phi}_{\rm n+1})$ in the second line of Eq. (2.68): thanks to IR safety, $O_{n+1}(\mathbf{\Phi}_{\rm n+1})-O_n(\mathbf{\Phi}_{\rm n})$ vanishes for soft/collinear configurations and, therefore, no logarithmically enhanced term is exponentiated by the Sudakov factor in these phase-space regions. That means that, at NLO accuracy, we can make use of the expansion $\Delta_{\rm match}=1+\mathcal{O}(\alpha_s)$. For the same reason, the lower integration bound in the $\Phi_{\rm rad}$ integration can be set to zero. With that in mind, we can write:

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\rm n}\Big\{\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})O_n(\mathbf{\Phi}_{\rm n})+\int d\Phi_{\rm rad}\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}(O_{n+1}(\mathbf{\Phi}_{\rm n+1})-O_n(\mathbf{\Phi}_{\rm n}))\Big\}\\ &+\int d\mathbf{\Phi}_{\rm n+1}(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))O_{n+1}(\mathbf{\Phi}_{\rm n+1})+\mathcal{O}(\alpha_s^{N+2})=\\ =&\int d\mathbf{\Phi}_{\rm n}\Big\{\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})O_n(\mathbf{\Phi}_{\rm n})+\int d\Phi_{\rm rad}R_{\rm S}(\mathbf{\Phi}_{\rm n+1})(O_{n+1}(\mathbf{\Phi}_{\rm n+1})-O_n(\mathbf{\Phi}_{\rm n}))\Big\}\\ &+\int d\mathbf{\Phi}_{\rm n+1}(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))O_{n+1}(\mathbf{\Phi}_{\rm n+1})+\mathcal{O}(\alpha_s^{N+2})\,, \end{aligned}\tag{2.69}$$

where we have neglected $\mathcal{O}(\alpha_s)$ contributions from the ratio $\bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})/B(\mathbf{\Phi}_{\rm n})=1+\mathcal{O}(\alpha_s)$ which multiplies the $R_{\rm S}$ term. At this stage, it is easy to see that from Eq. (2.69) we simply recover the NLO accurate expression for O of Eq. (2.44).

The matching formula in Eq. (2.66) is the master equation for *accumulative* (or *additive*) *matching* schemes, where the tree-level real corrections $R - R_S$ are added as a separate event sample to Born-like events. The first matching procedure of this kind to have been formulated is the famous MC@NLO method of Ref. [152]. This method has the advantage of leaving the PS untouched: not just the generation of extra radiation attached to the $(R - R_S)$-distributed events (dubbed hard events) is entirely left to the PS program, but also of all radiation in the soft/collinear regime attached to $\bar{B}_S(\mathbf{\Phi}_n)$-distributed events (named standard events). Therefore, in Eq. (2.66) the Sudakov factor Δ_{match} coincides with the PS one: no need to modify PS emissions is required. The latter is an appealing property for a matching scheme, especially due to recent progresses in trying to improve PS accuracy beyond LL (see for instance Ref. [108–111]). On the other hand, to avoid any double-counting, PS corrections up to order $\mathcal{O}(\alpha_s)$ have to be subtracted back. That prevents from spoiling the NLO accuracy of the FO calculation, but eventually makes the matching procedure PS-dependent: the R_S terms in MC@NLO (referred to as Monte Carlo subtraction terms) are extracted from the splitting kernels used by the specific PS algorithm (typically given by AP splitting functions times the Born amplitude). Two main complications arise from this choice of R_S. The first one is due to the fact that the PS approximation of the real amplitude is not guaranteed to reproduce the full IR singular structure of R, which effectively spoils the IR cancellation of singularities occuring at NLO level in Eq. (2.64). In this case, the FO matching term of Eq. (2.66) will generally be non-vanishing. To mitigate this issue, R_S are more commonly chosen in such a way to interpolate between the behaviour of the PS splitting kernels, reproduced for $t_\Phi > t_c$, and the subtraction counterterms C^{α_r}, for $t_\Phi < t_c$ [12]. A second problem related to the choice of R_S for $t_\Phi > t_c$ is that the PS approximation of R can overestimate the real amplitude in some phase-space regions, with no guarantee for $R - R_S$ to be positive. Even though the overall result will still be positive, events with negative weights can be generated, which significantly threatens the efficiency of the event generation procedure [13].

A second standard matching procedure is the POWHEG one, introduced in Ref. [86, 135]. The POWHEG method essentially differs from MC@NLO by the choice of R_S, which is now given by the exact real amplitude times a *damping factor* $h(\mathbf{\Phi}_{n+1})$. The factor $h(\mathbf{\Phi}_{n+1})$ is used to separate the hard and singular regions of R, in such a way that close

[12] The mismatch between the PS splitting kernels and the subtraction terms can introduce some residual dependence on the cutoff scale t_c, which might in principle deteriorate the FO accuracy of the result, as discussed in Ref. [85].

[13] As also shown in Ref. [153], if we have a sample of N events with constant weights $w_i = \pm c$, where the fraction of negative weights is ϵ, the effective number of events which controls the scaling of the accuracy in a MC integration is reduced:

$$\frac{\sqrt{\sum_i w_i^2}}{\sum_i w_i} = \frac{\sqrt{c^2\,\mathrm{N}^2}}{c\,(1-2\epsilon)\,\mathrm{N}} = \frac{1}{\sqrt{(1-2\epsilon)^2\,\mathrm{N}}}\,, \tag{2.70}$$

in such a way that the error of a MC integration will scale with the inverse of the squared root of $\mathrm{N}_{\text{eff}} = (1-2\epsilon)^2\,\mathrm{N}$, instead of N. Therefore, larger event samples are needed to obtained a statistical accuracy comparable to samples free of negative weights. In the context of MC@NLO, the possibility to reduce the fraction ϵ has been explored in the so called MC@NLO-Δ variant of the algorithm [154].

to a singular configuration of the real phase space $h(\mathbf{\Phi}_{n+1}) \to 1$ and R_S exactly matches the IR singular structure of R by definition. That said, standard and hard events in POWHEG are generated as usual with a hit-and-miss method, but MC subtraction terms are not needed, since the kinematics of the hardest radiation attached to standard events is also generated by the MC program. This is done in a PS-like approach, by making use of a modified Sudakov form factor Δ_{match} which exponentiates the singular real contributions R_S [14]. In this approach the hardest PS emission is therefore corrected and generated with a t_Φ value according to the full real amplitude, with emissions having $t < t_\Phi$ still accounted for by the PS. The POWHEG method has different advantages. First of all, it can be easily interfaced to any PS algorithm, without any ad-hoc MC subtraction term to be extracted from the specific implementation of the PS. Moreover, the difference $R - R_S$ is by definition positive and event samples are free of negative weights (with the clear exception of those events associated to phase-space regions where perturbation theory becomes unreliable). On the other hand, some new complications have to be taken into account. The first one has to do with the fact that the ratio R_S/B can become unphysically large close to resonance peaks: this issue has been faced in a modified POWHEG method, specifically designed to deal with processes involving resonances [155]. Moreover, the POWHEG approach is meant to correct the hardest PS emission with the exact matrix element. If the PS is ordered in a variable t which can be associated to the hardness of the radiation (typically its transverse momentum), then the POWHEG formula simply amounts to correct the first PS emission; if a different ordering variable, such as the angle of the splitting pair, is used, the hardest emission will generally follow a certain number of soft large-angle emissions. This set of soft emissions is required to correctly reproduce the final-state multiplicity and the colour coherent pattern of QCD radiation. A workaround to this problem was proposed in Ref. [135] by letting the PS evolve from the hard scale of the process down to the scale of the hardest emission, instead of the usual PS cutoff (a feature that gives the name to this special PS evolution, dubbed *truncated shower*). Then, the kinematics of the radiation with the largest transverse momentum $p_{T,h}$ can be corrected with the POWHEG formalism. Finally, extra radiation down to t_c can be generated by means of *vetoed showers*: only radiation having a transverse momentum $p_T < p_{T,h}$ is accepted, while vetoed branchings with $p_T > p_{T,h}$ are just used to reset the shower evolution scale (in order to preserve unitarity).

When it comes to NLO matching, the MC@NLO and POWHEG methods, together with all their variants, are still the mainstream ones, due to their high flexibility and generality, and nowadays they can boast a huge amount of applications in many physical contexts. Nevertheless, alternative matching approaches have also been proposed more recently. One of those is the KrkNLO method [156], which relies on a *multiplicative*

[14] In a previous POWHEG version the entire R contribution was exponentiated without any distinction between standard and hard events. That was corrected afterwards to better describe phase-space regions away from the PS soft-collinear approximation.

matching, given schematically by:

$$\begin{aligned} d\sigma =& B(\mathbf{\Phi}_{\mathrm{n}}) \cdot (1+\delta_{\mathrm{NLO}}) \cdot \Big\{ \Delta_{\mathrm{match}}(\mathbf{\Phi}_{\mathrm{n}}, t_c) d\mathbf{\Phi}_{\mathrm{n}} \\ &+ \left[\Delta_{\mathrm{match}}(\mathbf{\Phi}_{\mathrm{n}}, t_\Phi) \frac{R_{\mathrm{S}}(\mathbf{\Phi}_{\mathrm{n+1}})}{B(\mathbf{\Phi}_{\mathrm{n}})} \right] \frac{R(\mathbf{\Phi}_{\mathrm{n+1}})}{R_{\mathrm{S}}(\mathbf{\Phi}_{\mathrm{n+1}})} \, \theta(t_\Phi - t_c) d\mathbf{\Phi}_{\mathrm{n+1}} \Big\}, \end{aligned} \tag{2.71}$$

where, as with MC@NLO, the generation of radiation is completely left to the PS, so that Δ_{match} and R_{S} coincide with the PS Sudakov and splitting kernel, respectively. The shower is attached to events distributed according to the LO differential cross section $B(\mathbf{\Phi}_{\mathrm{n}})$. NLO accuracy over inclusive observables is then achieved with an overall weight $(1+\delta_{\mathrm{NLO}})$ and the factor R/R_{S}, which corrects the kinematics of the hardest PS emission in a multiplicative way. The advantages of this method are its conceptual simplicity (NLO accuracy is achieved just with proper correction weights), the absence of negative weights (as in POWHEG) and the fact of leaving the PS evolution untouched (as in MC@NLO). Anyway, the NLO corrections δ_{NLO} should be provided in a proper factorization scheme for the PDFs, named *Monte-Carlo scheme*, introduced in Ref. [157] (which simplifies the form of the NLO corrections by including initial-state collinear singularities in the PDF definition). Moreover, the PS is required to properly cover the full radiation phase space, which is not the case if for instance R_{S} vanishes in some phase-space regions. This requirement on the PS definitely limits the applicability of the method. Finally, since the multiplicative matching generates weighted events by definition, the efficiency of the method can be limited by the efficiency of a process-specific unweighting procedure [15]. All that has made concrete applications of this matching procedure quite involved, except for simple processes like Drell-Yan or Higgs production.

Therefore, even though NLO matching is a well-understood problem, ongoing efforts in trying to improve the accuracy of current PS algorithms require to find better ideas and to explore new possibilities for a matching procedure combing all advantages of existing approaches. As a proof of how this research field is still extremely active, a very recent proposal for a new matching scheme was presented in Ref. [151] with the name of MAcNLOPS. Its goal is to combine the MC@NLO additive scheme for the generation of hard events and the KrkNLO multiplicative scheme in phase-space regions where the PS overestimates the real matrix element.

[15] In nature events clearly come without a weight, but the occurrence of a given kinematic configuration is proportional to the value of the differential cross section on that phase-space point. Therefore, if event generators provide weighted events, in order to become usable in realistic analyses, those events first have to be unweighted. That means that a set of N events with weights w_i is converted into a set of M < N events with unit weights. If $w_{\mathrm{max}} = \mathrm{max}_i\{w_i | i = 1, \ldots, \mathrm{N}\}$ and r is a random number $r \in [0,1]$, then each event i in the sample is processed and accepted only if $w_i/w_{\mathrm{max}} > r$. In so doing, the effective size of the sample is reduced: that is why the statistical accuracy of weighted events is generally bound to be lower. Different ways to reduce the fraction of rejected events are available, which should be tuned on a process-by-process basis (see for instance Ref. [158]).

2.3.2 Multijet merging

Another direction to improve PS simulations is merging multiple event samples characterized by a different jet multiplicity. If the PS is corrected by tree-level matrix elements in each event sample, then the merged event set can describe hard and well-separated emissions at LO accuracy, while accounting for extra soft and collinear radiation with parton showers. Among other reasons, the interest in improving multijet simultations is motivated by the fact that multijet events represent a relevant source of background in many BSM searches.

The number of merging approaches is extremely vast, but, as discussed in Ref. [159], they all share some common features. First of all, an event with a given multiplicity n (with a maximum value $n_{\max} \sim 6$ due to computational feasibility reasons) is generated according to the exact tree-level matrix element. The multiplicity n is selected on an event-by-event basis with a probability P_n given by:

$$P_n = \frac{\sigma_n}{\sum_{m=1}^{n_{\max}} \sigma_m}\,, \tag{2.72}$$

where σ_n is the cross section for the production of n final-state jets as obtained from a FO calculation. All merging strategies essentially agree till this point. What can be significantly different is instead the way the transition between the hard regime described by the full matrix element and the PS regime is accessed. Despite that, all methods are unavoidably required to define a *merging scale* $Q_{\rm MS}$ to separate the two kinematic regimes.

From a historical point of view, one of the first merging approaches to have been formulated is the MLM one [160]. The strength of the method is its simplicity and the fact that no constraint on the way the PS is run is required. Despite the many variants of the method, the key idea of the original MLM algorithm is the following. The LO partonic events are showered starting from the hard scale of the process. Then, partons generated by PS evolution are clustered into jets with a proper clustering algorithm, which defines an interparton separation measurement y_{ij} between partons i and j. In this context, this jet-resolution (or clustering) scale plays the role of a merging variable, once a value $Q_{\rm MS}$ for the merging scale is chosen: two partons are considered as resolved if $y_{ij} > Q_{\rm MS}$, otherwise they are combined into the same jet. At this point, the energy-ordered list of these jets is required to match the set of LO partons in number and in terms of a parton-jet distance criterion (i.e in order to be identifiable, the separation between the candidate jet and the parton can not exceed a certain separation scale). If these conditions are not met, the showered event is discarded. The simplicity of this procedure comes with the price of little analytic control. Indeed, a veto procedure based on a parton-jet matching criterion does not have a simple translation into a well-defined QCD object, even though its effect is to suppress extra radiation just like PS Sudakov form factors. The mismatch between the MLM veto and the PS Sudakov factor does not allow for a smooth transition between the hard and the PS regime, which effectively results in a residual dependence on the unphysical merging scale $Q_{\rm MS}$.

A different set of approaches for LO merging is based on the CKKW method, first proposed in Ref. [161]. A smooth transition between the matrix elements and PS evo-

lution is pursed by a proper reweighting of the partonic events. This is done by first reconstructing a *parton shower history* for the partons generated according to the FO matrix element. In the original CKKW approach, such a history is obtained by iteratively combining with a $k_\perp$-clustering (or Durham) algorithm [162] the tree-level parton lines. The distance measurement y_{ij} for partons i and j is computed as:

$$y_{ij} \equiv 2 \min(E_i^2, E_j^2)(1 - \cos\theta_{ij})/Q^2 \,, \tag{2.73}$$

where E_i and E_j are the parton energies and θ_{ij} their angular separation [16], while Q is the hard scale of the process. At each step, only partons which are the closest in terms of y_{ij} are clustered together (a history reconstruction strategy named *winner-takes-all*): that corresponds to selecting the most probable configuration with which a PS evolution would have generated the LO partons. In so doing, the recursive clustering defines a set of ordered scales $q_k \equiv \sqrt{y_{ij}}Q$ for $k \in \{1, \ldots, n\}$, where q_k corresponds to the relative transverse momentum of the clustered pair i and j. The clustering scales $\{q_k\}$ are first used to rescale matrix elements, where the α_s coupling is evaluated at the hard scale Q, by the ratio $\prod_{k=1}^{n} \alpha_s(q_k)/\alpha_s^n(Q)$ (for hadron-initiated processes, a similar factor involving ratios of PDFs is also required). Such a reweighting procedure is also applied in some implementations of the MLM merging [164], to allow matrix elements to better mimic the PS behaviour. What is peculiar of the CKKW strategy is the introduction of a reweighting term involving the product of analytic Sudakov factors $\Delta(q_k, q_{k+1})$, one for each parton produced at scale q_k and splitting at scale q_{k+1} within the reconstructed parton shower history. The reweighted event is then passed to the parton shower program. A crucial ingredient of the CKKW method is setting the shower starting scale t_0. If the shower ordering variable t matches the clustering variable q, which is used to define the merging scale $Q_{\rm MS} = \sqrt{y_{\rm cut}}Q$ in terms of a jet-resolution value $y_{\rm cut}$, t_0 can be safely set to $Q_{\rm MS}$. Even if an exact correspondence is never fulfilled in current shower algorithms, no big issues are expected for transverse-momentum ordered showers. If the shower is instead ordered in a different variable, such as the angle of the splitting pair, setting t_0 to $Q_{\rm MS}$ would result in a radiation gap, due to missing PS emissions with transverse momentum lower than the merging scale, but shower evolution values larger than the merging scale. This radiation gap effectively introduces an unphysical dependence on the merging scale. This issue was already addressed in the original publication: in Ref. [161] the PS for each parton line was invoked from the scale where the parton was first produced (instead of from $Q_{\rm MS}$), with a veto over radiation with transverse momentum larger than $Q_{\rm MS}$. A further improvement to also account for missing radiation from the intermediate states of the reconstructed shower history was proposed in Ref. [165], which adapted the idea of truncated showers from Ref. [135] to the CKKW framework.

A quite different variant of the CKKW approach, named CKKW-L, was first introduced in Ref. [166]. The idea is to ameliorate the problem of the mismatching between the analytic Sudakov reweighting and the PS evolution with a clustering algorithm for

[16] This is the original choice of y_{ij} used in Ref. [161], where hadronic final states for e^+e^- collisions where considered. For hadron initiated events, the proper definition of the Durham jet measurement should be used [163].

the reconstruction of the PS history which exactly inverts the PS branchings. The reconstructed branching kinematics and the scales q_k clearly become dependent on the specific PS that is employed, but that allows for an exact equivalence between the clustering and the evolution scales. Moreover, when the shower history is built, a *probabilistic strategy* is used, where all possible shower histories are considered and only at the very end the one with highest probability (given by the product of the associated splitting kernels) is chosen. Finally, when it comes to reweighting the matrix elements, the product of Sudakov factors is not included analytically, but it is effectively accounted for by allowing the PS to generate radiation from the reconstructed states between scales q_k and q_{k+1}: the PS is invoked from scale q_k and any radiation above q_{k+1} is vetoed. Despite these adjustments, a residual dependence on $Q_{\rm MS}$ still survives, as it can be seen for instance by looking at the exclusive partonic cross section σ_n. For simplicity, we can consider σ_3 as an illustrative example. We denote by $\bar{\Delta}(q_i, q_j)$ the product of Sudakov form factors giving the non-emission probability from scale q_i to q_j for the different parton lines of the history and by $\Gamma_{2\to3}(t)$ the PS splitting kernel integrated over the momentum fraction. We can schematically write the 3-jet cross section differential in the kinematics of the lowest multiplicity ($n = 2$) phase space as obtained from CKKW-L merging as:

$$\frac{d\sigma_3}{d\Phi_2} = \int dt\, \bar{\Delta}(Q,t) \cdot \left\{ \frac{\alpha_s(t)}{\alpha_s(Q)}\ |\mathcal{M}_3^{(0)}|^2(Q)\ \theta(t - Q_{\rm MS}) \right. \\ \left. + |\mathcal{M}_2^{(0)}|^2(Q)\ \Gamma_{2\to3}(t)\ \theta(Q_{\rm MS} - t) \right\} \cdot \bar{\Delta}(t, t_c)\,. \tag{2.74}$$

The first term of the previous equation describes the radiation kinematics using the correct tree-level matrix element $|\mathcal{M}_3^{(0)}|^2$ for values of the radiation hardness $t > Q_{\rm MS}$ (possible PDF reweighting factors for hadron-initiated processes are understood). For $t < Q_{\rm MS}$, the kinematics of the third jet is described by the PS attached to the lower multiplicity matrix element $|\mathcal{M}_2^{(0)}|^2$. Therefore, the size of the residual $Q_{\rm MS}$ dependence is related to how well the PS can describe the third emission also away from the strict soft/collinear regime: in the limit case where $\Gamma_{2\to3}(t) \sim |\mathcal{M}_3^{(0)}|^2/|\mathcal{M}_2^{(0)}|^2$ the merging scale dependence completely cancels out.

In Ref. [167] it was also pointed out that this residual $Q_{\rm MS}$ dependence in the CKKW-L algorithm potentially violates unitarity: when all exclusive jet-multiplicity samples are combined, the inclusive cross section $\sigma^{\rm incl}$ is not exactly recovered. We consider again a simple case where two jet samples with multiplicities $n = 2$ (the one with the lowest multiplicity) and $n = 3$ are merged using the CKKW-L approach. Then, $\sigma^{\rm incl}$ can be computed by integrating over the hardness t of the emitted parton in the 3-jet sample down to the merging scale (for radiation having $t < Q_{\rm MS}$, unitarity is guaranteed by PS):

$$\sigma^{\rm incl} = \int d\Phi_2 \left\{ \frac{d\sigma_2}{d\Phi_2} + \int_{Q_{\rm MS}}^{Q} dq_1\, \frac{d\sigma_3}{d\Phi_2 dq_1} \right\} = \\ = \int d\Phi_2 \left\{ |\mathcal{M}_2^{(0)}|^2 \cdot \bar{\Delta}(Q, Q_{\rm MS}) + \int_{Q_{\rm MS}}^{Q} dq_1 \frac{\alpha_s(q_1)}{\alpha_s(Q)} |\mathcal{M}_3^{(0)}|^2(Q) \cdot \bar{\Delta}(Q, q_1) \right\}, \tag{2.75}$$

where q_1 is the reconstructed scale for the first emission within the shower history [17]. As with Eq. (2.74), when considering the kinematics of the first emission, we just highlighted the dependence on the hardness. Then, expanding Eq. (2.75) to order α_s gives:

$$\sigma^{\text{incl}} = \int d\Phi_2 \Bigg\{ |\mathcal{M}_2^{(0)}|^2 \left(1 - \int_{Q_{\text{MS}}}^{Q} dq_1 \Gamma_{2\to 3}(q_1) + \mathcal{O}(\alpha_s^2)\right) + \int_{Q_{\text{MS}}}^{Q} dq_1 \frac{\alpha_s(q_1)}{\alpha_s(Q)} |\mathcal{M}_3^{(0)}|^2(Q)(1+\mathcal{O}(\alpha_s))\Bigg\} \neq \int d\Phi_2 |\mathcal{M}_2^{(0)}|^2 + \mathcal{O}(\alpha_s^2)\,, \tag{2.76}$$

where we see again that the amount of unitarity violation is related to the accuracy of the PS in the description of the radiation kinematics. Since PS algorithms are mostly LL accurate, the non-cancellation in Eq. (2.76) generally introduces a dependence on the merging scale of the form $\sim \alpha_s L_{\text{cut}}$, with $L_{\text{cut}} = \log(Q/Q_{\text{MS}})$. If $Q_{\text{MS}} \ll Q$, then these contributions become sizeable, spoiling the FO accuracy of the inclusive cross section. Even if the first PS emission is corrected using MEC approaches, the problem will unavoidably appear again when merging higher multiplicity samples. This issue was the main motivation for the introduction of the Unitary Matrix Element plus Parton Shower (UMEPS) merging in Ref. [167], where unitarity is restored by subtracting back some ad-hoc events with the radiation kinematics integrated out. The price to be paid for that is the introduction of negative-weighted events, with a considerable loss of efficiency (see footnote 13 in this chapter).

This brief and necessarily incomplete analysis of different merging approaches has allowed us to describe the main subtleties that are involved when trying to supplement the PS description with matrix elements. In particular the problem of reducing the dependence on the unphysical merging scale is common to all approaches and we will come back to this again in Section 3.2. Other LO merging methods that we did not cover in our summary are for instance pseudo-showers [168] or the MESS approach [169], the latter formulated within the VINCIA sector shower framework.

Despite all complications involved in merging jet samples corrected by tree-level matrix elements, a LO accuracy in the description of jet observables is often not enough for a precise phenomenology. The most natural line of development is trying to merge different jet multiplicities at NLO accuracy. This can be done by improving existing NLO matching methods with merging strategies in such a way that NLO accuracy is reached also for the production of associated jets. Examples of NLO merging algorithms are for instance FXFX [170], which combines MC@NLO with MLM merging, MEPS@NLO [171], where MC@NLO is supplemented with suitably modified truncated showers, or UNLOPS [172], which is based on an extension of the unitarity paradigm of UMEPS. Another relevant framework to achieve multijet merging is the GENEVA one [173], where FO calculations for different jet multiplicities are combined together with the use of analytic resummation. All these methods are fairly general, since multiple jet samples can be combined at NLO, and many of their advantages/disadvantages are inherited from the matching and merging

[17] The absence of the Sudakov factor $\bar{\Delta}(q_1, Q_{\text{MS}})$ in the reweighting term for $|\mathcal{M}_3^{(0)}|^2$ in Eq. (2.75) is required for the highest multiplicity sample in order to allow the shower to generate more than $n = 3$ jets above the merging scale, as noticed in Ref. [163, 168].

schemes they are based on. Another method that we will discuss more extensively in Section 3.2 is MiNLO′ [174, 175], which is a Ckkw-inspired NLO merging approach embedded in the Powheg matching framework. As we will see, it has the advantage of preserving unitarity, since no unphysical merging scale is introduced. As a drawback, the method is less general than previous approaches and only two jet samples (the 0- and 1-jet) can be merged, even though some attempts to extend it to three jet samples were pursued in Ref. [176].

The NLO merging algorithms are in their turn the starting point for the formulation of NNLO+PS generators, where NNLO accuracy is achieved for observables inclusive over QCD radiation, on top of NLO accuracy for observables involving one jet. As we will discuss more extensively in Section 3.3, reaching NNLO+PS accuracy is still an open problem, since no general method is available and just few approaches with limited applicability have been formulated so far.

3 Matching NNLO QCD calculations with parton showers

Armed with a general overview of the different methods to achieve precise phenomenology, we are now ready to specialize the discussion to the theoretical framework that was used in this work to produce the results presented in Chapter 4. In Section 3.1 we present the POWHEG method, that was already introduced in Subsection 2.3.1, and some technical details of this NLO matching procedure. The MINLO′ method, that we mentioned in Subsection 2.3.2 in the context of NLO merging, is then described in Section 3.2 in some of its basic aspects. Finally, MINNLO$_{\rm PS}$ is introduced in Section 3.3, where we illustrate the key idea to reach NNLO accuracy for QCD calculations matched with parton showers.

3.1 The POWHEG framework

In Eq. (2.66) we reported a quite general expression to achieve NLO matching, which allowed us to describe the two mainstream matching methods, i.e MC@NLO and POWHEG, with a uniform notation, as done in Ref. [151]. We mentioned that, in POWHEG, the singular approximation of the real amplitude $R_{\rm S}$ is chosen exactly equal to R, up to a damping factor to separate its finite and singular contributions. Since the FO matching term of Eq. (2.66) vanishes with this choice [85], the POWHEG NLO matching formula can be simply rewritten as:

$$\begin{aligned} d\sigma =& \bar{B}_{\rm S}(\mathbf{\Phi}_{\rm n})\Big\{\Delta_{\rm pow}(\mathbf{\Phi}_{\rm n},t_c)d\mathbf{\Phi}_{\rm n}+\Delta_{\rm pow}(\mathbf{\Phi}_{\rm n},k_{\rm T}(\mathbf{\Phi}_{\rm n+1}))\frac{R_{\rm S}(\mathbf{\Phi}_{\rm n+1})}{B(\mathbf{\Phi}_{\rm n})}\theta(k_{\rm T}(\mathbf{\Phi}_{\rm n+1})-t_c)d\mathbf{\Phi}_{\rm n+1}\Big\} \\ &+(R(\mathbf{\Phi}_{\rm n+1})-R_{\rm S}(\mathbf{\Phi}_{\rm n+1}))\theta(k_{\rm T}(\mathbf{\Phi}_{\rm n+1})-t_c)d\mathbf{\Phi}_{\rm n+1}\,. \end{aligned} \tag{3.1}$$

As we explicitly verified in Section 2.3.1, the previous formula guarantees a NLO accurate description of infrared safe observables with no singularities associated to the Born phase space $\mathbf{\Phi}_{\rm n}$. The first line of Eq. (3.1) is responsible for the generation of radiation in the singular regions starting from a Born-like kinematics and the second line describes the hard radiation. The term in the curly bracket of Eq. (3.1) exactly mimics the PS formula, where the Sudakov form factor is chosen to be

$$\Delta_{\rm pow}(\mathbf{\Phi}_{\rm n},p_{\rm T})\equiv\exp\Big\{-\int\frac{[d\Phi_{\rm rad}R_{\rm S}(\mathbf{\Phi}_{\rm n+1})\theta(k_{\rm T}(\mathbf{\Phi}_{\rm n+1})-p_{\rm T})]_{\bar{\mathbf{\Phi}}_{\rm n}=\mathbf{\Phi}_{\rm n}}}{B(\mathbf{\Phi}_{\rm n})}\Big\}\,, \tag{3.2}$$

with a Sudakov exponent matching the choice of the splitting kernel used in POWHEG, i.e the singular part of the real amplitude $R_{\rm S}$. The hardness $k_{\rm T}(\mathbf{\Phi}_{\rm n+1})$ of the QCD radiation which is matched to the shower corresponds, in the soft/collinear limit, to the

transverse momentum of the emitted particle with respect to the emitter. If this scale is used to set the factorization and renormalization scales in the α_s coupling and the PDF factors in the integrand of Eq. (3.2), then it was shown in Ref. [86] that $\Delta_{\text{pow}}(\mathbf{\Phi}_{\text{n}}, k_{\text{T}})$ has LL accuracy (with some small adjustments even NLL accuracy can be achieved for some processes): this is one important ingredient which ensures that the LL accuracy of the PS is not spoiled. The POWHEG formula is used to generate the hardest radiation, distributed according to the tree-level matrix element, under the assumption that this radiation event has been the first one to occur in the PS cascade. If the POWHEG ordering variable equals the one of the PS, the matching procedure continues straightforwardly: the hardness of the radiation generated by POWHEG serves as a starting scale for the PS evolution. This is the case for showers ordered in the transverse momentum of the radiation, such as the PYTHIA8 shower [124]. The PS strong ordering assumption would instead be violated by POWHEG if, for instance, the shower is angular ordered. If that is the case, the correct radiation pattern is reproduced by making use of truncated showers, accounting for large-angle soft radiation occurring before the hardest POWHEG emission at k_{T}, as we described in Section 2.3.1. Then, the ordering condition for the shower starting at k_{T} scale is enforced by a veto.

Together with the ordering variable, the choice of the cutoff t_c in the POWHEG and the PS Sudakov has to agree, as well. The standard choice of POWHEG is to set the minimum allowed transverse momentum for the radiation close to the hadronization scale (specifically $t_c^2 = p_{\text{T,min}}^2 = 0.8\,\text{GeV}^2$).

The generation of radiation in POWHEG is closely related to the NLO subtraction procedure, that we described in Section 2.1.3. Even though a formulation with the CS scheme can be quite easily achieved, the FKS subtraction turns out to be the most natural choice in POWHEG. Indeed, the matrix element R_{S} is written as a sum of contributions having just one singularity for a parton i getting soft or collinear to only one final-state parton j ($\alpha_{r,ij}$ singular configuration in the set of singular regions $\{\alpha_{\tilde{r}}\}$) or to an initial-state parton ($\alpha_{r,i}$ singular configuration within the sets $\{\alpha_{\ominus}\}$ and $\{\alpha_{\oplus}\}$, depending on the initial parton direction). In what follows, if needed, we will refer to a generic singular configuration within the two sets $\alpha_{\tilde{r}}$ and $\alpha_{\oplus/\ominus}$ as α_r. The aforementioned separation of contributions to R_{S} can be easily obtained by isolating the real-emission singular regions by means of some FKS functions $\mathcal{F}$ of the real phase space $\Phi_{\text{n+1}}$, satisfying a partition-of-unity condition:

$$\sum_{\alpha_{r,i}\in\{\alpha_{\oplus}\}} \mathcal{F}_i^{\oplus} + \sum_{\alpha_{r,i}\in\{\alpha_{\ominus}\}} \mathcal{F}_i^{\ominus} + \sum_{\alpha_{r,ij}\in\{\alpha_{\tilde{r}}\}} \mathcal{F}_{ij} = 1 \tag{3.3}$$

with the sums ranging over all possible singular configurations, and fulfilling the limits

$$\begin{aligned} &\lim_{k_m^0\to 0}\left(\mathcal{F}_i^{\oplus} + \mathcal{F}_i^{\ominus} + \sum_j \mathcal{F}_{ij}\right) = \delta_{im} \qquad \lim_{\vec{k}_m\parallel\vec{k}_l} (\mathcal{F}_{ij} + \mathcal{F}_{ji}) = \delta_{im}\delta_{jl} + \delta_{il}\delta_{jm} \\ &\lim_{\vec{k}_m\parallel\vec{k}_{\oplus/\ominus}} \mathcal{F}_i^{\oplus/\ominus} = \delta_{im} \qquad \lim_{\vec{k}_m\parallel\vec{k}_{\oplus/\ominus}} \mathcal{F}_{ij} = 0 \qquad \lim_{\vec{k}_m\parallel\vec{k}_{\oplus/\ominus}} \mathcal{F}_i^{\ominus/\oplus} = 0\,, \end{aligned} \tag{3.4}$$

which guarantee that each FKS function isolates one singular region.

By means of the FKS functions, the real amplitude can be separated into a sum of contributions $R_i^{\alpha_\oplus}$, $R_i^{\alpha_\ominus}$ and $R_{ij}^{\alpha_{\tilde{r}}}$ which are divergent in only one singular region:

$$\begin{aligned} R_{\mathrm{S}} &= \sum_{\alpha_{r,i}\in\{\alpha_\oplus\}} R_i^{\alpha_\oplus} + \sum_{\alpha_{r,i}\in\{\alpha_\ominus\}} R_i^{\alpha_\ominus} + \sum_{\alpha_{r,ij}\in\{\alpha_{\tilde{r}}\}} R_{ij}^{\alpha_{\tilde{r}}} \\ &\equiv \sum_{\alpha_{r,i}\in\{\alpha_\oplus\}} \mathcal{F}_i^{\oplus} R_{\mathrm{S}} + \sum_{\alpha_{r,i}\in\{\alpha_\ominus\}} \mathcal{F}_i^{\ominus} R_{\mathrm{S}} + \sum_{\alpha_{r,ij}\in\{\alpha_{\tilde{r}}\}} \mathcal{F}_{ij} R_{\mathrm{S}}\,. \end{aligned} \tag{3.5}$$

As soon as the FKS functions allow to separate the different singular regions, their exact definition is arbitrary. One can introduce some objects

$$\begin{cases} d_i^{\oplus/\ominus} = 0 & \Longleftrightarrow \quad k_i^0 = 0 \text{ or } \vec{k}_i \parallel \vec{k}_{\oplus/\ominus} \\ d_{ij} = 0 & \Longleftrightarrow \quad k_i^0 = 0 \text{ or } k_j^0 = 0 \text{ or } \vec{k}_i \parallel \vec{k}_j \end{cases} \tag{3.6}$$

where the energies k_i^0 and vector components $\vec{k}_i$ are computed in the center of mass (CM) frame of the incoming partons. Then, by defining:

$$\mathcal{D} = \sum_k \frac{1}{d_k^{\oplus}} + \sum_k \frac{1}{d_k^{\ominus}} + \sum_{kl} \frac{1}{d_{kl}} \tag{3.7}$$

we can write:

$$\mathcal{F}_i^{\oplus/\ominus} = \frac{1}{\mathcal{D} d_k^{\oplus/\ominus}} \qquad \mathcal{F}_{ij} = \frac{1}{\mathcal{D} d_{ij}} g\Big(\frac{k_i^0}{k_i^0 + k_j^0}\Big) \tag{3.8}$$

where the extra function g, satisfying the conditions

$$\lim_{k\to 0} g(k) = 1 \quad \lim_{k\to 1} g(k) = 0 \quad g(k) + g(1-k) = 1\,, \tag{3.9}$$

is required to avoid double countings of singular regions if both $\mathcal{F}_{ij}$ and $\mathcal{F}_{ji}$ exist (which is the case if i and j are gluons arising from a gluon splitting).

To get even closer to the actual POWHEG implementation of Eq. (3.1), it is important to remember that in a GPMC event generator one has to keep track of the flavour structure of the NLO event: different flavours effectively describe different events and give rise to different shower evolutions. We define as ℓ_{B} the list of particle flavours of the Born event and ℓ_{R} that of the real event. Once a real flavour structure is given, the different singular regions $\{\alpha_r|\ell_{\mathrm{R}}\} \equiv \{\alpha_{\oplus/\ominus}|\ell_{\mathrm{R}}\} \cup \{\alpha_{\tilde{r}}|\ell_{\mathrm{R}}\}$ associated to it can be identified. We also define as $\{\alpha_r|\ell_{\mathrm{B}}\} \equiv \{\alpha_{\oplus/\ominus}|\ell_{\mathrm{B}}\} \cup \{\alpha_{\tilde{r}}|\ell_{\mathrm{B}}\}$ the set of singular regions whose underlying Born has ℓ_{B} as a flavour structure. If the flavour ℓ_{R} can not develop any singular configuration, we define its contribution to the real amplitude as a *regular* one; otherwise ℓ_{R} simply contributes to the singular part of R. Therefore, the set of real flavour structures can be divided into a regular $\mathrm{Rg}(\ell_{\mathrm{R}})$ and singular $\mathrm{Sg}(\ell_{\mathrm{R}})$ set as $\{\ell_{\mathrm{R}}\} = \mathrm{Rg}(\ell_{\mathrm{R}}) \cup \mathrm{Sg}(\ell_{\mathrm{R}})$. Then, within each singular region, we can further identify a strictly *singular* contribution R_{S} (or also R^{sing} in what follows), which is the only one to be exponentiated in Δ_{pow}, and a *remnant* one R^{remn}. The separation between singular and remnant contributions is

achieved by means of a damping factor $h^{(\alpha_r)} = h(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}, \Phi_{\rm rad}^{(\alpha_r)})$ (that we have already introduced in Section 2.3.1), which depends on the details of the mapping of Eq. (2.32) in each singular region. Therefore, we can decompose the real amplitude as:

$$
\begin{aligned}
R = \sum_{\{\ell_{\rm R}\}} R_{\ell_{\rm R}} &= \sum_{\ell_{\rm R}\in{\rm Sg}(\ell_{\rm R})} R_{\ell_{\rm R}} + \sum_{\ell_{\rm R}\in{\rm Rg}(\ell_{\rm R})} R^{\rm reg}_{\ell_{\rm R}} = \\
&= \sum_{\ell_{\rm R}\in{\rm Sg}(\ell_{\rm R})} \Bigg\{ \sum_{\alpha_{\tilde r}\in\{\alpha_{\tilde r}|\ell_{\rm R}\}} R^{\alpha_{\tilde r}}_{\ell_{\rm R}} + \sum_{\alpha_{\oplus/\ominus}\in\{\alpha_{\oplus/\ominus}|\ell_{\rm R}\}} R^{\alpha_{\oplus/\ominus}}_{\ell_{\rm R}} \Bigg\} + \sum_{\ell_{\rm R}\in{\rm Rg}(\ell_{\rm R})} R^{\rm reg}_{\ell_{\rm R}} = \\
&= \sum_{\ell_{\rm R}\in{\rm Sg}(\ell_{\rm R})} \Bigg\{ \sum_{\alpha_{\tilde r}\in\{\alpha_{\tilde r}|\ell_{\rm R}\}} h^{\alpha_{\tilde r}} R^{\alpha_{\tilde r}}_{\ell_{\rm R}} + \sum_{\alpha_{\oplus/\ominus}\in\{\alpha_{\oplus/\ominus}|\ell_{\rm R}\}} h^{\alpha_{\oplus/\ominus}} R^{\alpha_{\oplus/\ominus}}_{\ell_{\rm R}} + \\
&+ \sum_{\alpha_{\tilde r}\in\{\alpha_{\tilde r}|\ell_{\rm R}\}} (1-h^{\alpha_{\tilde r}}) R^{\alpha_{\tilde r}}_{\ell_{\rm R}} + \sum_{\alpha_{\oplus/\ominus}\in\{\alpha_{\oplus/\ominus}|\ell_{\rm R}\}} (1-h^{\alpha_{\oplus/\ominus}}) R^{\alpha_{\oplus/\ominus}}_{\ell_{\rm R}} \Bigg\} + \sum_{\ell_{\rm R}\in{\rm Rg}(\ell_{\rm R})} R^{\rm reg}_{\ell_{\rm R}} = \\
&\equiv \sum_{\ell_{\rm R}\in{\rm Sg}(\ell_{\rm R})} \Bigg\{ \sum_{\alpha_{\tilde r}\in\{\alpha_{\tilde r}|\ell_{\rm R}\}} R^{{\rm sing},\,\alpha_{\tilde r}}_{\ell_{\rm R}} + \sum_{\alpha_{\oplus/\ominus}\in\{\alpha_{\oplus/\ominus}|\ell_{\rm R}\}} R^{{\rm sing},\,\alpha_{\oplus/\ominus}}_{\ell_{\rm R}} \Bigg\} + \\
&+ \sum_{\ell_{\rm R}\in{\rm Sg}(\ell_{\rm R})} \Bigg\{ \sum_{\alpha_{\tilde r}\in\{\alpha_{\tilde r}|\ell_{\rm R}\}} R^{{\rm remn},\,\alpha_{\tilde r}}_{\ell_{\rm R}} + \sum_{\alpha_{\oplus/\ominus}\in\{\alpha_{\oplus/\ominus}|\ell_{\rm R}\}} R^{{\rm remn},\,\alpha_{\oplus/\ominus}}_{\ell_{\rm R}} \Bigg\} + \sum_{\ell_{\rm R}\in{\rm Rg}(\ell_{\rm R})} R^{\rm reg}_{\ell_{\rm R}}
\end{aligned}
\tag{3.10}
$$

where we have defined

$$
R^{{\rm sing},\,\alpha_r}_{\ell_{\rm R}} = h^{\alpha_r} R^{\alpha_r}_{\ell_{\rm R}} \qquad R^{{\rm remn},\,\alpha_r}_{\ell_{\rm R}} = (1-h^{\alpha_r}) R^{\alpha_r}_{\ell_{\rm R}}\,. \tag{3.11}
$$

Despite being formally correct, the organization of the sum of the singular and remnant contributions in Eq. (3.10) is slighlty different from its actual implementation. In POWHEG the generation of the real kinematics starts from the generation of Born-like variables, which are interpreted as the underlying Born variables of a given singular real flavour structure in the set ${\rm Sg}(\ell_{\rm R})$. Since an underlying Born configuration can be associated to multiple singular real flavour structures, different radiation patterns are possible. In this reverted logic, we can rewrite Eq. (2.64) for $\bar B_{\rm S}$ (that we rename here as $\bar B$ for simplicity) for a fixed Born flavour structure $\ell_{\rm B}$ as:

$$
\begin{aligned}
\bar B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n}) = B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n}) + V_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n}) + \sum_{\alpha_r\in\{\alpha_r|\ell_{\rm B}\}} \Bigg[\int d\Phi_{\rm rad} \{ R^{\rm sing}_{\ell_{\rm R}}(\mathbf{\Phi}_{\rm n+1}) - C(\mathbf{\Phi}_{\rm n+1}) \} \Bigg]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}=\mathbf{\Phi}_{\rm n}} \\
+ \sum_{\alpha_\oplus\in\{\alpha_\oplus|\ell_{\rm B}\}} \int \frac{dz}{z} G^{\alpha_\oplus}_{\oplus}(\mathbf{\Phi}_{{\rm n},\oplus}) + \sum_{\alpha_\ominus\in\{\alpha_\ominus|\ell_{\rm B}\}} \int \frac{dz}{z} G^{\alpha_\ominus}_{\ominus}(\mathbf{\Phi}_{{\rm n},\ominus})\,.
\end{aligned}
\tag{3.12}
$$

If we define a flavour-aware Sudakov form factor $\Delta^{\ell_{\rm B}}_{\rm pow}$ as:

$$
\Delta^{\ell_{\rm B}}_{\rm pow}(\mathbf{\Phi}_{\rm n}, p_{\rm T}) \equiv \exp\Bigg\{ - \sum_{\alpha_r\in\{\alpha_r|\ell_{\rm B}\}} \int \frac{\big[d\Phi_{\rm rad} R^{\rm sing}_{\ell_{\rm R}}(\mathbf{\Phi}_{\rm n+1}) \theta(k_{\rm T}(\mathbf{\Phi}_{\rm n+1}) - p_{\rm T}) \big]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}=\mathbf{\Phi}_{\rm n}}}{B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n})} \Bigg\}\,,
\tag{3.13}
$$

we can expand the flavour sum of Eq. (3.1) and rewrite it as:

$$
\begin{aligned}
d\sigma = \sum_{\ell_{\mathrm{B}}} \bar{B}_{\ell_{\mathrm{B}}}(\mathbf{\Phi}_{\mathrm{n}})\, d\mathbf{\Phi}_{\mathrm{n}} \Bigg\{ & \Delta^{\ell_{\mathrm{B}}}_{\mathrm{pow}}(\mathbf{\Phi}_{\mathrm{n}}, p_{\mathrm{T,min}}) \\
& + \sum_{\alpha_r \in \{\alpha_r | \ell_{\mathrm{B}}\}} \frac{\left[d\Phi_{\mathrm{rad}} \Delta^{\ell_{\mathrm{B}}}_{\mathrm{pow}}(\mathbf{\Phi}_{\mathrm{n}}, k_{\mathrm{T}}(\mathbf{\Phi}_{\mathrm{n+1}})) R^{\mathrm{sing}}_{\ell_{\mathrm{R}}}(\mathbf{\Phi}_{\mathrm{n+1}}) \theta(k_{\mathrm{T}}(\mathbf{\Phi}_{\mathrm{n+1}}) - p_{\mathrm{T,min}}) \right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\mathrm{n}}^{(\alpha_r)} = \mathbf{\Phi}_{\mathrm{n}}}}{B_{\ell_{\mathrm{B}}}(\mathbf{\Phi}_{\mathrm{n}})} \Bigg\} \\
& + \sum_{\ell_{\mathrm{B}}} \sum_{\alpha_r \in \{\alpha_r | \ell_{\mathrm{B}}\}} \left[R^{\mathrm{remn}}_{\ell_{\mathrm{R}}}(\mathbf{\Phi}_{\mathrm{n+1}}) \theta(k_{\mathrm{T}}(\mathbf{\Phi}_{\mathrm{n+1}}) - p_{\mathrm{T,min}}) d\mathbf{\Phi}_{\mathrm{n}} d\Phi_{\mathrm{rad}} \right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\mathrm{n}}^{(\alpha_r)} = \mathbf{\Phi}_{\mathrm{n}}} \\
& + \sum_{\ell_{\mathrm{R}} \in \mathrm{Rg}(\ell_{\mathrm{R}})} R^{\mathrm{reg}}_{\ell_{\mathrm{R}}}(\mathbf{\Phi}_{\mathrm{n+1}}) \theta(k_{\mathrm{T}}(\mathbf{\Phi}_{\mathrm{n+1}}) - p_{\mathrm{T,min}}) d\mathbf{\Phi}_{\mathrm{n+1}} \,. \qquad (3.14)
\end{aligned}
$$

To summarize, according to the previous equation, the regular contributions are obtained by adding up all real flavour structures $\ell_{\mathrm{R}} \in \mathrm{Rg}(\ell_{\mathrm{R}})$: the corresponding real kinematics is generated with a distribution $R^{\mathrm{reg}}_{\ell_{\mathrm{R}}}$ using a hit-and-miss technique and then directly passed to the shower program. Each flavour channel is selected with a probability proportional to its contribution to the overall sum. For flavours in the list $\mathrm{Sg}(\ell_{\mathrm{R}})$, when generating a remnant contribution, the real kinematics is obtained in pretty much the same way as with the regular part. Otherwise, POWHEG starts by generating an underlying Born kinematics with a probability $\bar{B}_{\ell_{\mathrm{B}}}$, where again a flavour channel is chosen according to its contribution to $\bar{B}$. Then, the kinematics of the extra radiation with respect to the Born phase space its accounted for by the POWHEG Sudakov. Notice that ℓ_{R} in Eq. (3.12), Eq. (3.13) and Eq. (3.14), used for $R^{\mathrm{sing}}_{\ell_{\mathrm{R}}}$ and $R^{\mathrm{remn}}_{\ell_{\mathrm{R}}}$ when summing over $\alpha_r \in \{\alpha_r | \ell_{\mathrm{B}}\}$, refers to the real flavour structure having α_r as a singular region.

In order not to spoil the strong ordering condition, the $k_{\mathrm{T}}^{(\alpha_r)}$ generated by POWHEG in each singular region α_r is required not to be harder than any other radiation which might be part of the underlying Born kinematics (which is obtained from the mapping shown in Eq. (2.32)). With the separation of regions of Eq. (3.5), each contribution to the sum is strongly suppressed away from its specific singular configuration. Therefore, $R^{\mathrm{sing},\,\alpha_r}$ is used to generate radiation with a transverse momentum $k_{\mathrm{T}}^{(\alpha_r)}$ which is the smallest within the corresponding real kinematics, with unordered configurations suppressed. That can be enforced by a specific choice of d_{ij} and $d_i^{\oplus/\ominus}$ in Eq. (3.6), which can be defined in such a way that the separation of singular contributions is done according to the transverse momentum of the radiation. Therefore, we can write:

$$
\begin{cases}
d_i^{\oplus/\ominus} = (k_i^0)^{2b}\, 2^b\, (1 \mp \cos\theta_i)^b \\
d_{ij} = \left(\frac{k_i^0 k_j^0}{k_i^0 + k_j^0} \right)^{2b} 2^b\, (1 - \cos\theta_{ij})^b
\end{cases} \qquad (3.15)
$$

with θ_i and θ_{ij} the angles that the radiated parton forms with the beam axis and the emitter, respectively. Again, all quantities are computed in the CM frame of the incoming partons. The parameter b can be used in POWHEG to control the strength of the separation of regions.

All of the features that we have described in this section have been fully automatized in a general framework, named POWHEG-BOX [177], which is a FORTRAN-based code to handle NLO calcultations and match them to a GPMC program with the POWHEG method.

3.1.1 Treatment of resonances

When resonance particles are involved in the process, the procedure that we have outlined above can be jeopardized on different levels. First of all, the convergence of the subtraction procedure can be spoilt for these processes. As we discussed in Section 2.1.3, the definition of the counterterms for the subtration of IR singularities is strictly related to the mapping of Eq. (2.32), where the real phase-space kinematics is factored into an underlying Born and a radiation kinematics. Such remappings are usually not designed to preserve the virtuality of the resonance lines. Since the real term is defined on the real phase space $\mathbf{\Phi}_{\mathrm{n+1}}$ and the counterterms on the factorized one, the virtualities of the resonances on which the two contributions are computed might differ. This mismatch can potentially affect the efficiency of the subtraction procedure in the cancellation of singularities.

When considering the matching of the NLO calculation with the parton shower, further subtleties arise. The POWHEG algorithm relies on the fact that R^{sing} correctly approaches its soft and/or collinear approximation in the relevant singular regions, getting close to the Born amplitude B up to a splitting kernel factor. If the two terms are not computed on the same resonance virtuality, R^{sing} might significantly departs from its soft/collinear approximation and the ratio R^{sing}/B in the Sudakov factor Δ_{pow} can become large enough to induce distorsions in the radiation pattern.

Moreover, since at FO interference contributions among partons radiated from different resonances are present (in a finite-width calculation), no information on the resonance assignment of partons for a given flavour configuration can be passed to the shower program, which consequently does not preserve the invariant mass of the resonance, leading to further shape distorsions.

Solutions to account for these problems were first discussed in Ref. [178], where resonances were treated in the zero-width approximation. The extension to the finite-width case was first proposed in Ref. [155], with the introduction of substancial modifications to the POWHEG algorithm which have been then implemented in a dedicated framework, dubbed POWHEG-BOX-RES. In the latter approach, given a Born flavour structure ℓ_{B} (named *bare* flavour structure), which just contains the information on the initial- and final-state flavours, all possible tree-level graphs (or *resonance histories*) are identified [1]: the nodes of these graphs coincide with the lines of the resonant particles. Therefore, given a bare flavour structure, different *full* flavour structures $\hat{\ell}_{\mathrm{B}}$ are possible. They

[1] As described in details in Appendix A.1 of Ref. [179], the automatic identification of resonance histories is based on the *richer-resonance-structure* criterion. If a resonance history involves a fermion line emitting two EW bosons (namely a W, a Z or a Higgs) that could directly couple to each other (for instance through a triple coupling), that resonance structure is discarded. Indeed, the same resonance structure can be described by a richer topology, where the two electroweak bosons arise from the decay of a single electroweak boson.

embody the details of the entire resonance history: the flavours of the intermediate resonances are made explicit, together with the information on the resonance of origin (sometimes dubbed *mother*) for all final states (sometimes *siblings* of a given mother resonance). The full flavour information is used in the POWHEG-BOX-RES framework to improve the POWHEG-BOX code on many respects.

First of all, that is used to perform a resonance-aware phase space sampling, which is employed, in the context of MC integration, to evaluate an integral more frequently in the regions of phase space that give the larger contribution to the overall result (where the integrand is predominantly localized, e.g close to a resonace peak). Moreover, when radiation from resonance decay products is considered, the efficiency of the infrared subtraction is improved by means of a resonance-aware subtraction algorithm, where the mapping from a real to its underlying Born configuration preserves the virtuality of intermediate resonances. Finally, in a parton-shower context, the distortion of resonances through recoil effects is avoided by supplying the PS with details on the resonance cascade chain.

The key idea behind the algorithm used in the POWHEG-BOX-RES framework is to decompose the cross section into contributions associated to a well-defined resonance structure, which are enhanced on that particular cascade chain. We have already exploited the fact that the Born cross section B can be written as the sum over all bare flavour structures $B_{\ell_{\rm B}}$. After introducing the set $T(\ell_{\rm B})$ of all $\hat{\ell}_{\rm B}$ having $\ell_{\rm B}$ as a bare flavour structure (which essentially means they share the same initial- and final-state flavour configurations), we can go one step further and break down $B_{\ell_{\rm B}}$ into a weighted sum over $\hat{\ell}_{\rm B}$ using weight functions $\mathcal{P}_{\hat{\ell}_{\rm B}}$:

$$B_{\ell_{\rm B}} = \sum_{\hat{\ell}_{\rm B}\in T(\ell_{\rm B})} \mathcal{P}_{\hat{\ell}_{\rm B}} B_{\ell_{\rm B}}, \quad \text{with} \sum_{\hat{\ell}_{\rm B}\in T(\ell_{\rm B})} \mathcal{P}_{\hat{\ell}_{\rm B}} = 1\,. \tag{3.16}$$

The weight functions $\mathcal{P}_{\hat{\ell}_{\rm B}}$ are chosen such that Eq. (3.16) expresses $B_{\ell_{\rm B}}$ as a sum over resonance-peaked terms $\mathcal{P}_{\hat{\ell}_{\rm B}} B_{\hat{\ell}_{\rm B}}$, which develop the expected resonance enhancement of $\hat{\ell}_{\rm B}$. Clearly, there is a certain freedom in their explicit expression. In the POWHEG-BOX-RES code, the quantities $P_{\hat{\ell}_{\rm B}}$ are introduced:

$$P_{\hat{\ell}_{\rm B}} = \prod_{i\in{\rm Nd}(\hat{\ell}_{\rm B})} \frac{m_i^4}{(s_i - m_i^2)^2 + \Gamma_i^2 m_i^2}\,, \tag{3.17}$$

where the product is performed over all nodes i in the set ${\rm Nd}(\hat{\ell}_{\rm B})$ associated to the graph described by $\hat{\ell}_{\rm B}$, and where s_i, m_i and Γ_i are the invariant mass of the decay product system, the resonance mass and width, respectively. Once $P_{\hat{\ell}_{\rm B}}$ are defined, the following choice of $\mathcal{P}_{\hat{\ell}_{\rm B}}$ is made:

$$\mathcal{P}_{\hat{\ell}_{\rm B}} = \frac{P_{\hat{\ell}_{\rm B}}}{\sum_{\hat{\ell}'_{\rm B}\in T(\ell_{\rm B}(\hat{\ell}_{\rm B}))} P_{\hat{\ell}'_{\rm B}}}\,, \tag{3.18}$$

where the sum in the denominator runs over all configurations in the set of graphs $T(\ell_{\rm B}(\hat{\ell}_{\rm B}))$ having a bare flavour structure $\ell_{\rm B}(\hat{\ell}_{\rm B})$ consistent with the initial and final

flavours of $\hat{\ell}_{\mathrm{B}}$. The same discussion done for the Born term also applies to the virtual correction $V_{\ell_{\mathrm{B}}}$ of Eq. (3.12).

For the real contribution, as detailed in Ref. [155], the decomposition into resonance enhanced regions has to be combined with the separation of singularities. After introducing a similar notation for the bare ℓ_{R} and full $\hat{\ell}_{\mathrm{R}}$ real flavour structures, this problem is faced in POWHEG-BOX-RES by considering only singular regions α_r within the set $\mathrm{Sr}(\hat{\ell}_{\mathrm{R}})$ of those which are *compatible* with a given resonance history $\hat{\ell}_{\mathrm{R}}$: the two particles becoming collinear should originate from the decay products of the same resonance or they should have both been generated at production level [2]. If we denote by $\ell_{\mathrm{R}}(\alpha_r)$ and $\hat{\ell}_{\mathrm{R}}(\alpha_r)$ the full and bare flavours associated to the singular region α_r, we can define the FKS decomposition of the real amplitude as:

$$R^{\alpha_r} = \frac{P_{\hat{\ell}_{\mathrm{R}}(\alpha_r)} d^{-1}_{\hat{\ell}_{\mathrm{R}}}(\alpha_r)}{\sum_{\hat{\ell}'_{\mathrm{R}} \in T(\ell_{\mathrm{R}}(\hat{\ell}_{\mathrm{R}}))} P_{\hat{\ell}'_{\mathrm{R}}} \sum_{\alpha'_r \in \mathrm{Sr}(\hat{\ell}'_{\mathrm{R}})} d^{-1}_{\hat{\ell}'_{\mathrm{R}}}(\alpha'_r)} R_{\ell_{\mathrm{R}}(\alpha_r)}\,, \tag{3.19}$$

where the $P_{\hat{\ell}_{\mathrm{R}}(\alpha_r)}$ weights are constructed like the corresponding Born weights. With this choice, in the soft and collinear limits $P_{\hat{\ell}_{\mathrm{R}}(\alpha_r)}$ reproduces the corresponding weight factor $P_{\hat{\ell}_{\mathrm{B}}}$, with $\hat{\ell}_{\mathrm{B}}$ the associated underlying Born structure of $\hat{\ell}_{\mathrm{R}}(\alpha_r)$. The terms $d_{\hat{\ell}_{\mathrm{R}}}(\alpha_r)$ have the same form presented in Eq. (3.15) (by choosing the appropriate definition for initial and final state singularities), but they also enclose a dependence on the full flavour structure $\hat{\ell}_{\mathrm{R}}$. Indeed, consistently with a resonance-aware remapping that preserves the invariant mass of the resonance, in POWHEG-BOX-RES not all of the $d_{\hat{\ell}_{\mathrm{R}}}(\alpha_r)$ are computed in the CM frame: for final-state singularities at the level of the decay products of a resonance, they are evaluated in the rest frame of the corresponding resonance.

3.1.2 Event generation

Since we will refer to it on different occasions, in this section we want to briefly sketch the general POWHEG procedure. The output of a run for a process implemented within the POWHEG-BOX framework is a text file containing the set of events distributed according to the master formula in Eq. (3.14). This file is written in a standardized form suitable for being passed to the Les Houches Interface for User Processes (LHIUP) [180], which allows for a simple interface to any GPMC program. Indeed, once the Les Houches Event (LHE) file is available, extra radiation beyond the hardest one provided by POWHEG can be generated by using that file as an input for the desired parton shower program.

In order to generate the LHE output, the POWHEG workflow can be described as consisting of four different stages, which have much in common with standard Monte Carlo programs. Anyway, as described in detail in Ref. [177], the expression in Eq. (3.12) is still not suitable for generating Born-like events distributed according to $\bar{B}_{\ell_{\mathrm{B}}}$. Indeed,

[2] The impact of considering only singular structures compatible with the resonance history is discussed in Ref. [155]. Indeed, such a restriction does not properly account for all soft non-collinear singularities and consequently, upon integration, some logarithmic terms $\log\Gamma$ of the resonance decay width arise. Anyway, the effect of these soft terms is claimed to be negligible in Ref. [155] within the POWHEG framework.

any hit-and-miss technique would require the evaluation of the $\bar{B}_{\ell_B}$ function several times and, for each evaluation, an integration over the radiation variables is needed, making the algorithm extremely inefficient. Therefore, for each singular region the radiation kinematics in $\Phi_{\rm rad}^{(\alpha_r)}$ is parametrized in terms of three variables $X_{\rm rad} = \{X_{\rm rad}^{(1)}, X_{\rm rad}^{(2)}, X_{\rm rad}^{(3)}\}$ ranging in the unit cube $\mathcal{C}$, so that one can introduce

$$\begin{aligned}\tilde{B}_{\ell_B}(\mathbf{\Phi}_n, X_{\rm rad}) = {} & B_{\ell_B}(\mathbf{\Phi}_n) + V_{\ell_B}(\mathbf{\Phi}_n) \\ & + \sum_{\alpha_r \in \{\alpha_r | \ell_B\}} \left[\left| \frac{\partial \Phi_{\rm rad}}{\partial X_{\rm rad}} \right| \{ R^{\rm sing}_{\ell_R}(\mathbf{\Phi}_{n+1}) - C(\mathbf{\Phi}_{n+1}) \} \right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_n^{(\alpha_r)} = \mathbf{\Phi}_n} \\ & + \sum_{\alpha_\oplus \in \{\alpha_\oplus | \ell_B\}} \frac{1}{z} \left| \frac{\partial z}{\partial X_{\rm rad}^{(1)}} \right| G_\oplus^{\alpha_\oplus}(\mathbf{\Phi}_{n,\oplus}) + \sum_{\alpha_\ominus \in \{\alpha_\ominus | \ell_B\}} \frac{1}{z} \left| \frac{\partial z}{\partial X_{\rm rad}^{(1)}} \right| G_\ominus^{\alpha_\ominus}(\mathbf{\Phi}_{n,\ominus}) \,. \end{aligned} \tag{3.20}$$

Moreover, in order to make use of Monte Carlo integration techniques (which are the standard ones when several integration dimensions are involved), the Born phase space $\mathbf{\Phi}_n$ is also parametrized in terms of a set $X_{\rm born}$ of $3n-2$ variables $X_{\rm born}^{(i)}$ in a unit hypercube, in such a way that

$$\bar{B}_{\ell_B}(\mathbf{\Phi}_n) = \left| \frac{\partial \mathbf{\Phi}_n}{\partial X_{\rm born}} \right| \cdot \bar{B}_{\ell_B}(X_{\rm born}) = \left| \frac{\partial \mathbf{\Phi}_n}{\partial X_{\rm born}} \right| \cdot \int_{\mathcal{C}} d^3 X_{\rm rad} \, \tilde{B}_{\ell_B}(X_{\rm born}, X_{\rm rad}) \,. \tag{3.21}$$

By trading $\bar{B}_{\ell_B}(\mathbf{\Phi}_n)$ for $\tilde{B}_{\ell_B}(X_{\rm born}, X_{\rm rad})$, points $(X_{\rm born}, X_{\rm rad})$ are generated: only at this stage the $X_{\rm rad}$ variables are discarded, which amounts to integrating over them, but in a computationally more efficient way. Then, the corresponding phase-space point in $\mathbf{\Phi}_n$ is computed from $X_{\rm born}$. A similar discussion holds for the remnant and regular contributions of Eq. (3.14), where the real phase space $\mathbf{\Phi}_{n+1}$ is also expressed using a set $X_{\rm real}$ of $(3n-2)+3$ variables, but where all of them are kept once they are generated.

By rewriting $\bar{B}_{\ell_B}$ as fully defined on a unit hypercube of dimension $(3n-2)+3$, the integration and generation of events can be handled via the `MINT` program [181], built in the POWHEG-BOX framework. `MINT` is a modification of the `SPRING-BASES` package [182], which reduces the storage requirement and has a special handle for non-positive valued functions. When a program for a process implemented in POWHEG-BOX is started, an importance sampling grid for the computation of the inclusive cross section is created and stored with `MINT` using a standard `VEGAS` algorithm [183]: that is done in the so-called POWHEG *stage 1*. Three separate grids are generated for the Born, remnant and regular contributions. An integration grid can be simply defined as a partition into smaller hypercubes (or cells) of the integration volume: in a MC approach, random points for the evaluation of the integrand are generated, which are distributed so that on average the number of points falling inside any hypercube is the same. By starting from an equally spaced partition (named *rectangular grid*), importance sampling grids are obtained by iteratively adjusting the cell dimensions in a way that minimizes the uncertainty of the integral evaluation: that results in having a finer partition in regions of phase space where the integral receives its larger contributions.

Monte Carlo techniques are the methods of choice when it comes to performing integrations over several dimensions (as we have already mentioned), but also when the

integration phase space is particularly involved. This is often the case for processes whose Born phase space has a non-trivial structure, as we will explicitly encounter for instance in Section 4.1.1. Indeed, in order for Eq. (2.24) to be infrared finite after that the subtraction procedure has been applied, $O_n(\mathbf{\Phi}_\mathrm{n})$ has to be finite as well. If some singularities associated to the Born phase space are still present, some minimal phase-space restrictions (named *generation cuts*) are required to obtain a finite result (see for instance Ref. [184]). Generation cuts should be carefully used and their values should be chosen to be much smaller than possible cuts (often *analyisis* or *fiducial cuts*) defining the fiducial volume, i.e the region of phase space which is eventually considered when inspecting the event sample. This precaution is needed, because choosing cuts of comparable size at generation and fiducial level can cause phase-space points which might have passed the fiducial cuts thanks to shower recoil effects to be cut away at generation level. A different (or sometimes complementary) approach to handle Born singularities is introducing a *suppression factor* F, which multiplies the differential cross section during integration (and which should be multiplied away at event generation level to recover the physical result). The analytic expression of F should be chosen in such a way to suppress the value of the integrand close to singular regions, so as to redirect the numerical sampling where the result of the integral is finite. Moreover, a proper form of F can have the further benefit of avoiding sampling large statistics in phase-space regions which are eventually removed by the fiducial cuts at analysis level. A drawback of this approach is that suppressed events will be very rare (since $F \to 0$ for singular configurations), but they will contribute with very large $1/F$ weights, appearing in distributions as undesired spikes which should be properly removed.

Another issue that was pointed out in Ref. [184] is that the function $\tilde{B}_{\ell_\mathrm{B}}(\mathbf{\Phi}_\mathrm{n}, X_\mathrm{rad})$ in Eq. (3.20) can become negative for some values of X_rad, even if the value of $\bar{B}_{\ell_\mathrm{B}}(\mathbf{\Phi}_\mathrm{n})$ is positive over the whole phase-space $\mathbf{\Phi}_\mathrm{n}$. That generates spurious negative-weighted events, which can be mostly removed by the `MINT` *folding* procedure, introduced in Ref. [181]. In essence, folding the integrand $\tilde{B}_{\ell_\mathrm{B}}(\mathbf{\Phi}_\mathrm{n}, X_\mathrm{rad})$ $n_{X^{(1)}_\mathrm{rad}}$ times along the $X^{(1)}_\mathrm{rad}$ direction, for instance, amounts to replacing it with the function

$$\tilde{B}^{\mathrm{folded}}_{\ell_\mathrm{B}}(\mathbf{\Phi}_\mathrm{n}, X^{(1)}_\mathrm{rad}, X^{(2)}_\mathrm{rad}, X^{(3)}_\mathrm{rad}) = \frac{1}{n_{X^{(1)}_\mathrm{rad}}} \sum_{i=0}^{n_{X^{(1)}_\mathrm{rad}}-1} \tilde{B}_{\ell_\mathrm{B}}\left(\mathbf{\Phi}_\mathrm{n}, \frac{X^{(1)}_\mathrm{rad}+i}{n_{X^{(1)}_\mathrm{rad}}}, X^{(2)}_\mathrm{rad}, X^{(3)}_\mathrm{rad}\right), \quad (3.22)$$

which integrates to the same result of the unfolded function, but where negative and positive values combine together. Therefore, if $n_{X^{(1)}_\mathrm{rad}}$ (or folding number) is large enough, the desired positive result is restored and $\tilde{B}_{\ell_\mathrm{B}}$ becomes a smoother function of $X^{(1)}_\mathrm{rad}$. In Ref. [185] another advantage of this procedure was highlighted. If the folding is restricted to the variables X_rad, with folding numbers $(n_{X^{(1)}_\mathrm{rad}}, n_{X^{(2)}_\mathrm{rad}}, n_{X^{(3)}_\mathrm{rad}})$, applying the folding equation 3.22 independently along the three radiation axes corresponds to evaluate the real contribution $n_{X^{(1)}_\mathrm{rad}} \cdot n_{X^{(2)}_\mathrm{rad}} \cdot n_{X^{(3)}_\mathrm{rad}}$ times more than the virtual part. The integral over the virtual contribution converges faster than the one over the real part, due to the simpler phase-space structure, but its computation is generally more demanding. Therefore, it is clear that more efficient codes can be obtained by properly tuning the

number of points used for the MC evaluation of the integral and the values of the folding numbers.

Once the importance sampling grids are available, the value of the inclusive cross section is effectively computed at POWHEG *stage 2*. During this integration stage, the `MINT` program also computes the upper bound of the integrand in each cell of the grids. Only this set of values is stored, differently from the `SPRING-BASES` algorithm, where the value of the integral in each cell is computed and saved. The stored upper bounds are used by `MINT` to compute an upper bounding envelope $\tilde{B}_{\text{env}}$ for $\tilde{B}$ (summed over the flavour components) in the form of a multidimensional step function. In order to reduce the memory requirement, such function is computed as a product of one-dimensional step functions $f^{(i)}$ along each integration direction i, as

$$\tilde{B}(X_{\text{born}}, X_{\text{rad}}) \leq \tilde{B}_{\text{env}}(X_{\text{born}}, X_{\text{rad}}) = \prod_{i=1}^{3n-2} f^{(i)}(X^{(i)}_{\text{born}}) \prod_{i=1}^{3} f^{(i)}(X^{(i)}_{\text{real}}) \,. \tag{3.23}$$

Similar upper bounding envelopes are computed for the remnant and regular contributions, as well. These functions are used in POWHEG *stage 4* in a hit-and-miss algorithm to generate Born-like events (where the X_{rad} variables are ignored) or events where the radiation is not associated to a singular region, distributed according to R^{remn} and R^{reg}. As a first step, the point $X = (X_{\text{born}}, X_{\text{rad}})$ is easily generated according to the envelope, together with a uniform random number r in the interval $0 \leq r \leq \tilde{B}_{\text{env}}(X)$. Only at this level the program computes, for instance, $\tilde{B}(X)$ and the event X is accepted only if $r \leq \tilde{B}(X)$, otherwise the configuration is discarded. By doing so, the kinematics of the events is correctly generated according to the desired distribution $\tilde{B}$. The efficiency of the hit-and-miss generation is strictly related to the number of configurations that have to be discarded before an event is accepted: the larger this number, the smaller the efficiency. Improvements in this direction can be achieved by choosing a better form for the envelope compared to the one of Eq. (3.23), as first proposed in Ref. [185]. Finally, a specific flavour structure is selected with a probability proportional to its contribution to the cross section. In the POWHEG-BOX-RES framework, the `MINT` program also handles a discrete summation index to label the different resonance histories, which are also selected accordingly at event generation level and provided as an extra information to the shower program.

If we end up with a Born kinematics $\mathbf{\Phi}_{\text{n}}$ and a Born flavour structure ℓ_{B}, the radiation kinematics has to be generated according to the term in curly brackets of Eq. (3.14). That is done in a way that is closely related to the generation of radiation in common PS programs. Specifically, the probability distribution of radiation in a given singular region $\alpha_r \in \{\alpha_r | \ell_{\text{B}}\}$ is given by

$$\left[\Delta^{\ell_{\text{B}}}_{\text{pow}}(\mathbf{\Phi}_{\text{n}}, k_{\text{T}}(\mathbf{\Phi}_{\text{n+1}})) \frac{R^{\text{sing}}_{\ell_{\text{R}}}(\mathbf{\Phi}_{\text{n+1}})}{B_{\ell_{\text{B}}}(\mathbf{\Phi}_{\text{n}})} \theta(k_{\text{T}}(\mathbf{\Phi}_{\text{n+1}}) - p_{\text{T,min}})\right]_{\alpha_r}^{\bar{\mathbf{\Phi}}^{(\alpha_r)}_{\text{n}} = \mathbf{\Phi}_{\text{n}}} d\Phi^{(\alpha_r)}_{\text{rad}} \,. \tag{3.24}$$

By noticing that the Sudakov form factor of Eq. (3.13) can be written as:

$$\Delta^{\ell_{\text{B}}}_{\text{pow}}(\mathbf{\Phi}_{\text{n}}, k_{\text{T}}) = \prod_{\alpha_r \in \{\alpha_r | \ell_{\text{B}}\}} \Delta^{\ell_{\text{B}}}_{\text{pow},\, \alpha_r}(\mathbf{\Phi}_{\text{n}}, k_{\text{T}}) \,, \tag{3.25}$$

with

$$\Delta^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}(\mathbf{\Phi}_{\rm n},k_{\rm T})\equiv\exp\left\{-\int\frac{\left[d\Phi_{\rm rad}R^{\rm sing}_{\ell_{\rm R}}(\mathbf{\Phi}_{\rm n+1})\theta(k_{\rm T}{}'(\mathbf{\Phi}_{\rm n+1})-k_{\rm T})\right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}=\mathbf{\Phi}_{\rm n}}}{B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n})}\right\}, \quad (3.26)$$

we can generate radiation distributed according to Eq. (3.24) using the *highest-$k_{\rm T}$ bid procedure*. In Ref. [86] it is shown that choosing the singular region α_r having the largest $k_{\rm T}^{(\alpha_r)}(\mathbf{\Phi}_{\rm n+1})$ correctly reproduces Eq. (3.24) [3]. Therefore, the problem boils down to generating radiation for each α_r according to

$$\left[\Delta^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}(\mathbf{\Phi}_{\rm n},k_{\rm T}(\mathbf{\Phi}_{\rm n+1}))\frac{R^{\rm sing}_{\ell_{\rm R}}(\mathbf{\Phi}_{\rm n+1})}{B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n})}\theta(k_{\rm T}(\mathbf{\Phi}_{\rm n+1})-p_{\rm T,min})\right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}=\mathbf{\Phi}_{\rm n}}d\Phi^{(\alpha_r)}_{\rm rad}\,, \quad (3.27)$$

which is achieved in POWHEG with a proper variant of the veto algorithm, illustrated in details in Ref. [177]. In essence, a uniform random number $r\in[0,1]$ is generated, and the equation

$$\frac{\tilde{\Delta}^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}(\mathbf{\Phi}_{\rm n},k_{\rm T}^{(\alpha_r)}(\mathbf{\Phi}_{\rm n+1}))}{\tilde{\Delta}^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}(\mathbf{\Phi}_{\rm n},p_{\rm T,max})}=r \quad (3.28)$$

is solved for $k_{\rm T}^{(\alpha_r)}$. Here, $\tilde{\Delta}^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}$ is obtained from $\Delta^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}$ by replacing the ratio $R^{\rm sing}_{\ell_{\rm R}}/B_{\ell_{\rm B}}$ under the integration sign by a proper upper bounding function $\mathcal{U}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})$ satisfying

$$\mathcal{K}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})\equiv\left[J(\bar{\mathbf{\Phi}}_{\rm n})\frac{R^{\rm sing}_{\ell_{\rm R}}(\mathbf{\Phi}_{\rm n+1})}{B_{\ell_{\rm B}}(\mathbf{\Phi}_{\rm n})}\right]_{\alpha_r}^{\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)}=\mathbf{\Phi}_{\rm n}}\leq\mathcal{U}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})\,, \quad (3.29)$$

which has to be chosen such that Eq. (3.28) can be easily solved. In Eq. (3.29), $J(\bar{\mathbf{\Phi}}_{\rm n}^{(\alpha_r)})$ is the jacobian factor that we encountered in Eq. (2.34). Moreover, in Eq. (3.28) $p_{\rm T,max}$ refers to the maximally-allowed transverse momentum, which is set to a value such that $\tilde{\Delta}^{\ell_{\rm B}}_{{\rm pow},\,\alpha_r}(\mathbf{\Phi}_{\rm n},p_{\rm T,max})=1$ at the beginning of the veto procedure. A proper expression for $\mathcal{U}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})$ can be obtained from the soft/collinear limit of $R^{\rm sing}_{\ell_{\rm R}}$ up to a normalization $\mathcal{N}(\Phi^{(\alpha_r)}_{\rm rad})$, whose value is computed during POWHEG *stage 3* and stored.

Once a $\bar{k}_{\rm T}^{(\alpha_r)}$ is found, the veto algorithm continues by generating $\Phi^{(\alpha_r)}_{\rm rad}$ variables according to $\mathcal{U}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})\theta(k_{\rm T}^{(\alpha_r)}(\mathbf{\Phi}_{\rm n+1})-\bar{k}_{\rm T}^{(\alpha_r)})$. Then, a second uniform random number r' is used in such a way that if $r'>\mathcal{K}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})/\mathcal{U}(\mathbf{\Phi}_{\rm n},\Phi^{(\alpha_r)}_{\rm rad})$ the radiation is discarded, $p_{\rm T,max}$ is set to $\bar{k}_{\rm T}^{(\alpha_r)}$ and the veto procedure starts again from Eq. (3.28). The algorithm stops once the radiation is accepted or when the constraint $\bar{k}_{\rm T}^{(\alpha_r)}>p_{\rm T,min}$ is violated: in the latter case, an event with no radiation is generated.

[3] More precisely, generation of radiation is obtained in POWHEG by further grouping the singular regions α_r into sets sharing the same emitter. The highest-$k_{\rm T}$ bid procedure is therefore applied in order to select a specific emitter and then a singular configuration α_r within that emitter region is chosen with a probability proportional to $R^{{\rm sing},\,\alpha_r}_{\ell_{\rm R}}$. See Ref. [177] for the detailed steps of the algorithm implemented in the code.

3.2 The MiNLO′ formalism

We have already discussed in Section 2.3.2 how different tree-level accurate jet multiplicities can be combined together to improve the PS description of hard radiation. On the other hand, we have also mentioned that precise phenomenology demands to go beyond a simple tree-level merging, requiring the accuracy for different jet samples to be pushed to NLO or higher perturbative orders. In this section we address the minimal problem of NLO merging, which serves as a basis for what is presented in Section 3.3. Moreover, we further specialize the discussion to event generators for colour-singlet production: we start by considering a single colourless particle `B`, where `B` could be a Higgs, a Z or W gauge boson, and in due course we generalize the discussion to any colour-singlet state `F`. That will allow us to introduce the theoretical framework used for the results presented in Chapter 4.

We can rephrase our NLO merging problem in the following way. We start from different event generators `BnJ`, which are NLO+PS accurate in the description of a colourless particle in association with n jets. Since a `BnJ` generator is based on a NLO calculation for $pp \to B + n\,\text{jets}$, it is NLO accurate in the description of observables involving n hard jets (and inclusive over additional radiation), LO accurate for quantities requiring m hard jets, with $m = n + 1$ (since the $n + 1$ jet process is just described at tree-level by the real contribution entering the full NLO result), and LL (or at most NLL) accurate when $m > n + 1$, since all extra radiation is handled by the PS program to which the generator is interfaced. A `BnJ` generator can not describe regions of the phase space where one or more jets of the n starting ones become unresolved: unphysical predictions are returned for more inclusive observables requiring only $B + n'$ jets, with $n' < n$. That is clear from the fact that `BnJ` generators are based on NLO calculations whose tree-level Born processes involve n jets, so that no virtual corrections for lower multiplicity processes are included to cancel divergences that appear when one or more of the n starting partons become soft or collinear. Such corrections are instead part of a `B(n-1)J` generator or, more generally, of `BmJ` generators with $m < n$. Conversely, to improve the description of radiation beyond n jets, `BmJ` generators with $m = n + 1$ or higher are required. From this discussion, it clearly emerges that having one event generator that can reach NLO accuracy simultaneously for different jet multiplicities can be extremely useful.

Despite the different theoretical frameworks, the vast majority of merging approaches starts from the generation of independent event samples, that, in all practical cases, are obtained from the generators `B`, `BJ` and `BJJ`. For the first two generators `B` and `BJ` not corresponding to the highest-considered multiplicity, only events where a NLO accurate description is achieved (namely involving 0 or 1 jet respectively) are kept, while the others are discarded. That is done by making use of a jet resolution variable or merging scale Q_{MS}, which discriminates the jet multiplicity of an event. Finally, the different exclusive event samples are joined together, so as to recover an inclusive description. If $d\sigma_n^{\text{MC}}/d\Phi_n$ is the fully differential Monte Carlo sample describing exactly n jets and $d\sigma_{\geq n}^{\text{MC}}/d\Phi_n$ the inclusive sample for n or more jets, an inclusive sample for B production

can be obtained as:

$$\frac{d\sigma^{\text{MC}}_{\geq 0}}{d\Phi_0} = \frac{d\sigma^{\text{MC}}_{0}}{d\Phi_0} + \frac{d\sigma^{\text{MC}}_{0+1}}{d\Phi_{0+1}} + \frac{d\sigma^{\text{MC}}_{\geq 0+2}}{d\Phi_{0+2}} \tag{3.30}$$

where, in essence, each generator contributes to a different jet bin, whose size if controlled by the merging scale. As we have already commented in Section 2.3.2, the presence of the unphysical scale Q_{MS} is one of the main source of issues in all merging algorithms.

First of all, the choice of a value for Q_{MS} can be quite delicate. Too large values of Q_{MS} would force describing relatively hard radiation only with tree-level accuracy, losing the benefit of merging. On the other hand, if the value of Q_{MS} is too low, the inclusive event sample is dominated by higher multiplicity events. The latter situation is also known to lead to unitarity violation problems, where the NLO accuracy of the inclusive cross section is lost. That can be easily understood by considering that for $Q_{\text{MS}} \to 0$ higher multiplicity generators like `BJ` and `BJJ` are forced to approach singular regions of phase space. When integrating over the radiation kinematics from the hard scale Q down to the merging scale to compute the inclusive cross section, the contribution of these regions manifests in large logarithmic terms of the form $L \equiv \log(Q/Q_{\text{MS}})$. As we also mentioned in Section 2.2, when $\alpha_s L^2 \sim 1$, then perturbative predictions become completely unreliable.

Besides using a properly tuned value of Q_{MS} to avoid the two extreme cases described above, as done in Ref. [170] in the context of the FxFx method, other merging solutions have been proposed to circumvent the problem. In the UNLOPS method [171], as we have already discussed, unitarity is enforced at the level of the inclusive cross section with the 'subtract-what-you-have-added' phylosophy. That allows to preserve unitarity for sufficiently inclusive quantities, but does not completely remove the merging scale dependence from all observables. Moreover, one relevant drawback of the subtraction approach used to restore unitarity is the introduction of a large amount of negative-weighted events. Alternatively, within the GENEVA framework [85, 173], unitarity is obtained for free by accounting for enhanced logarithmic terms with a sufficiently accurate resummation, which is used for merging different multiplicity samples. Despite the potential of this approach, technical issues are for instance related to the missing inclusion of power corrections (coming from the non-singular part of the resummation formula), or to the PS matching, which should not spoil the accuracy of the resummation itself.

In Ref. [174] a different merging approach was proposed, dubbed MINLO, for 'multi-scale improved NLO'. By starting from the CKKW method for tree-level merging [161], that we presented in Section 2.3.2, NLO merging can be obtained with just few adjustments. Therefore, if we want to merge the two generators `B` and `BJ`, the MINLO prescrition just requires to adapt the NLO accurate fixed-order cross section of `BJ` obtained from Eq. (2.64) as

$$\bar{B}(\mathbf{\Phi}_{\text{BJ}}) = \frac{\alpha_s(q)}{\alpha_s(Q)} \Delta^2_{\text{MiNLO}}(Q,\, q) \left[B(\mathbf{\Phi}_{\text{BJ}}) \left(1 - 2\Delta^{(1)}_{\text{MiNLO}}(Q,\, q)\right) + V(\mathbf{\Phi}_{\text{BJ}}) + \int d\Phi_{\text{rad}} R_{\text{S}}(\mathbf{\Phi}_{\text{BJJ}}) \right] \tag{3.31}$$

where we assume here for simplicity that the subtraction terms C^{α_r} have already been accounted for, and where the collinear remnant terms $G_{\ominus}^{\alpha_\ominus}$ and $G_{\oplus}^{\alpha_\oplus}$, which just marginally enter our discussion, are understood. The strong coupling and Sudakov reweighting factors in front of the squared brackets are standard Ckkw ingredients: once a shower history for the fixed-order configuration is reconstructed with a proper jet clustering algorithm, the clustering scale q of the first emission is used for evaluating the strong coupling. The same clustering scale is also chosen to fix the factorization scale in the PDFs and the collinear remnant terms. The analytic Sudakov factors associated to parton lines in the history allow for a smooth transition between the matrix element and the PS regimes in the Ckkw approach. Specifically, here Δ_{MiNLO} refers to the NLL accurate Sudakov factor associated to a quark line:

$$\Delta_{\text{MiNLO}}(Q,q) = \exp\left\{-\int_q^Q \frac{dt}{t}\left[\left(A^{(1)}\frac{\alpha_s(t)}{2\pi} + A^{(2)}\frac{\alpha_s^2(t)}{4\pi^2}\right)\log\frac{Q^2}{t^2} + B^{(1)}\frac{\alpha_s(t)}{2\pi}\right]\right\} \tag{3.32}$$

where $A^{(1)}$, $A^{(2)}$ and $B^{(1)}$ are the NLL resummation coefficients arising from the integration of the Altarelli-Parisi kernels:

$$A^{(1)} = 2\,C_F, \quad A^{(2)} = 2\,C_F K_g, \quad B^{(1)} = -3\,C_F, \tag{3.33}$$

with K_g given by Eq. (2.56). A key observation is that the presence of Sudakov reweighting terms in Eq. (3.31) suppresses the singular behaviour in the jet unresolved region of the `BJ` generator, providing a physical picture for arbitrary low values of q, as soon as the QCD cutoff Λ_{QCD} is met. That in principle allows to merge the two NLO `B` and `BJ` generators without relying on any unphysical merging scale.

In MiNLO the NLO accuracy of `BJ`, which is spoilt by the Sudakov reweighting, is restored by subtracting back the first order expansion $\Delta^{(1)}_{\text{MiNLO}}$ of the Sudakov form factor (with extra terms introduced by the reweighting being beyond the generator accuracy, i.e $\mathcal{O}(\alpha_s^3)$). For the same reason, the renormalization scale explicitly appearing in the virtual term is set to q, which ensures NLO scale compensation after the inclusion of the reweighting factor modifying the scale of α_s associated to the primary emission. But in order to guarantee NLO accuracy also for observables inclusive over QCD radiation, arbitrary Sudakov damping factors regulating the singularity of `BJ` are not necessarily sufficient. They have to be chosen in such a way not to violate unitarity, i.e to recover the NLO inclusive cross section when the radiation kinematics is integrated out. It was shown in Ref. [175] that Eq. (3.32) does actually violate unitary by terms of order $\mathcal{O}(\alpha_s^{3/2})$. Despite that, with very small but relevant modifications, unitarity can be restored, so as to promote the MiNLO prescription into a proper NLO merging algorithm, commonly referred to as MiNLO′. In order to have a better understanding of the missing ingredients to move from MiNLO to MiNLO′, we introduce in the next section some basic ideas of resummation for a specific observable, namely the transverse momentum of the colourless particle B. Actually, the formalism is more general and valid for a generic colour-singlet system of particles, as we will remark again later.

3.2.1 Ingredients for transverse-momentum resummation

We start by considering a hadronic scattering process like $h_1\, h_2 \to B + X$, where X refers to an arbitrary amount of QCD radiation accompanying the process. If B is a colour-singlet particle, there are just two relevant processes to be considered within the SM: the Higgs production, which is initiated at LO by a gluon fusion channel $gg \to H$, and the Drell-Yan process, with B either a Z or W boson (that we assume to be on-shell at the beginning of this discussion) produced via $\bar{q}'q$ annihilation. In view of the results presented in Chapter 4, we specialize the notation in our discussion to the second case in what follows. The Born kinematics $\Phi_{\rm B}$ of such processes is trivial and it is parametrized by just one variable (e.g the total invariant mass or the boson rapidity). If we want to access the transverse momentum $p_{\rm T}$ of the boson, we already know that, as soon as $p_{\rm T} \to 0$, only soft or collinear QCD radiation is kinematically allowed to contribute to the observable. Consequently, a FO description fails in such regions of phase space due to the appearance of large logarithmic contributions $\log(p_{\rm T}^2/Q^2)$ which spoil the convergence of the perturbative series. Close to these regions, all-order resummation is needed. Specifically, we can first write the cross section differential in the Born kinematics and in $p_{\rm T}$ as:

$$\frac{d\sigma}{d\Phi_{\rm B} dp_{\rm T}} \equiv \frac{d\sigma^{\rm sing}}{d\Phi_{\rm B} dp_{\rm T}} + \frac{d\sigma^{\rm reg}}{d\Phi_{\rm B} dp_{\rm T}} \tag{3.34}$$

with a separation between the singular part of the spectrum, encoded in the first term, and the regular part of it. The Born phase-space kinematics in $\Phi_{\rm B}$ should be defined by fixing a proper mapping $\mathbf{\Phi}_{\rm B,res}$ to project the full phase space with multiple emissions into the non-emission phase space. The singular contribution contains all logarithmically-enhanced terms $\alpha_s^m \log^n(p_{\rm T}^2/Q^2)$:

$$\frac{d\sigma^{\rm sing}}{d\Phi_{\rm B} dp_{\rm T}} = D_{-1}(\Phi_{\rm B}, \alpha_s)\delta(p_{\rm T}) + \sum_{n\geq 0} D_n(\Phi_{\rm B}, \alpha_s)\, \theta(p_{\rm T}) \left(\frac{\log p_{\rm T}^2}{p_{\rm T}^2}\right)_+^n \tag{3.35}$$

where the coefficients D_n admit a perturbative expansion in α_s. The Born term, together with all loop corrections, contributes to the first divergent part of Eq. (3.35), proportional to the δ distribution, since no radiation is present and the boson can only have a null $p_{\rm T}$. The second term of Eq. (3.35) accounts for the logarithmic contributions to be resummed, where the plus prescription, defined as in Eq. (2.23), encodes as usual the cancellation of IR divergences between the real and virtual terms.

The regular part of the spectrum in Eq. (3.34) just contains terms which are finite for $p_{\rm T} \to 0$ to all orders and therefore it can be safely described within a perturbative framework.

A standard approach to account for the logarithmic terms in the singular part of the spectrum to all orders is to carry out the resummation in the impact parameter or b space: in this formalism, also known as Collins-Soper-Sterman (CSS) resummation [65, 186], the two-component vector $\vec{b}$ usually denotes the Fourier conjugate variable of the projection of the transverse momentum over the plane perpendicular to the beam axis. Moving

to the b space through a Fourier transformation has the advantage of simplifying the handling of the momentum conservation constraint. In this approach, the resummed $p_{\rm T}$ spectrum can be written in a physically meaningful way as [187] (assuming $\overline{\rm MS}$ both as a factorization and renormalization scheme)

$$\frac{d\sigma^{\rm sing}}{d\Phi_{\rm B}dp_{\rm T}} = \sum_{q,q'}\int_0^1 dx_1\int_0^1 dx_2\int_0^{+\infty} db\frac{b}{2}J_0(bp_{\rm T})\sigma^{\rm (LO)}_{\bar{q}'q}(\mathbf{\Phi}_{\rm B})\,H_{\bar{q}'q}(\mathbf{\Phi}_{\rm B},\alpha_s(Q)) \\ \cdot[C_{\bar{q}'a}(\alpha_s)\otimes f_{a/h_1}](x_1,b_0/b)\,[C_{qb}(\alpha_s)\otimes f_{b/h_2}](x_2,b_0/b)\,\exp S(Q,b)\,, \quad (3.36)$$

where $b=|\vec{b}|$. The Bessel function J_0 and the constant $b_0\equiv 2e^{-\gamma_E}$ (with $\gamma_E=0.5772\ldots$ the Euler constant) originate from the integration over the polar angle θ formed by the transverse momentum and impact parameter vectors, i.e $\vec{b}\cdot\vec{p_{\rm T}}=|\vec{b}||\vec{p_{\rm T}}|\cos\theta$. In b space, the large logarithms $\log(p_{\rm T}^2/Q^2)$ for $p_{\rm T}\ll 1$ turn into $\log(Q^2b^2)$ contributions, which are enhanced for $b\gg 1$.

The differential LO cross section for the partonic process is encoded in $\sigma^{\rm (LO)}_{\bar{q}'q}(\mathbf{\Phi}_{\rm B})$, which can also be written in terms of the Born matrix element squared $|M^{\rm B}|^2_{\bar{q}'q}$ as $\sigma^{\rm (LO)}_{\bar{q}'q}(\mathbf{\Phi}_{\rm B})=d|M^{\rm B}|^2_{\bar{q}'q}/d\mathbf{\Phi}_{\rm B}$. Higher order QCD corrections are enclosed in the *hard function* $H_{\bar{q}'q}(\mathbf{\Phi}_{\rm B},\alpha_s(Q))$, which has the perturbative expansion

$$H_{\bar{q}'q}(\mathbf{\Phi}_{\rm B},\alpha_s)=1+\sum_{n=1}^{+\infty}\left(\frac{\alpha_s}{2\pi}\right)^n H^{(n)}_{\bar{q}'q}(\mathbf{\Phi}_{\rm B})\,. \quad (3.37)$$

As usual, we have denoted by $\mathbf{\Phi}_{\rm B}$ the Born phase space also including the parton energy fraction dependence, i.e $\mathbf{\Phi}_{\rm B}=\{x_1,\,x_2,\,\Phi_{\rm B}\}$. Then, the integration over x_1 and x_2 is weighted by the hadronic PDFs, in line with collinear factorization. Together with the so called *coefficient functions* $C_{ab}(\alpha_s,x)$, having the expansion

$$C_{ab}(\alpha_s,x)=\delta_{ab}\delta(1-x)+\sum_{n=1}^{+\infty}\left(\frac{\alpha_s}{2\pi}\right)^n C^{(n)}_{ab}(x)\,, \quad (3.38)$$

PDFs enter the convolution

$$[C_{qc}(\alpha_s)\otimes f_{c/h}](x,b_0/b)\equiv\sum_c\int_x^1\frac{dz}{z}C_{qc}(\alpha_s(b_0/b),z)f_{c/h}(x/z,b_0/b)\,, \quad (3.39)$$

where a sum over the flavours c is also carried out. Finally, logarithmic terms are resummed via exponentiation in the Sudakov form factor

$$\exp S(Q,b)=\exp\left\{-\int_{b_0/b}^{Q}\frac{dt}{t}\left[A(\alpha_s(t))\log\frac{Q^2}{t^2}+B(\alpha_s(t))\right]\right\}, \quad (3.40)$$

with resummation coefficients

$$A(\alpha_s)=\sum_{n=1}^{+\infty}\left(\frac{\alpha_s}{2\pi}\right)^n A^{(n)}\quad B(\alpha_s)=\sum_{n=1}^{+\infty}\left(\frac{\alpha_s}{2\pi}\right)^n B^{(n)}\,. \quad (3.41)$$

It can be easily seen that LL accuracy can simply be achieved with the inclusion of the $A^{(1)}$ coefficient, while $A^{(2)}$, $B^{(1)}$, $C^{(1)}$ and $H^{(1)}$ are required to reach NLL resummation, $A^{(3)}$, $B^{(2)}$, $C^{(2)}$ and $H^{(2)}$ for NNLL and so on.

The benefit from expressing resummation in the form of Eq. (3.36) is that all process dependence can be enclosed in the hard function, while the collinear functions and the Sudakov factor are process independent, since they only originate from the universal structure of QCD radiation in the IR regime. That also suggests a pretty intuitive physical picture, as described in Ref. [187]. The hard function accounts for contributions to the transverse-momentum spectrum of the boson at hard scales $p_{\mathrm{T}} \sim Q$ via virtual corrections: the function $H_{\bar{q}'q}$ can be obtained by the ratio of the loop contributions at a given perturbative order over the Born term. For p_{T} ranging from hard Q to b_0/b scales, the incomplete cancellation between real and virtual corrections giving rise to large logarithms is described by the Sudakov exponent. For even lower scales $p_{\mathrm{T}} < b_0/b$, the cancellation between real and virtual terms is exact for initial-state soft radiation (due to infrared safety), while incomplete for initial-state collinear radiation, which is indeed accounted for by the coefficient functions.

Therefore, despite the physical understanding provided by Eq. (3.36), Ref. [187] makes clear that the decomposition in hard, collinear and Sudakov contributions is to some extent ambiguous, due to the fact that the transverse-momentum spectrum is not a collinear safe observable. The exact definition of the coefficient functions depends on how initial-state singularities are subtracted, and this freedom propagates to $H_{\bar{q}'q}$ and S. This ambiguity manifests in the *renormalization group symmetry* of Eq. (3.36), which is expressed by the set of transformations:

$$\tilde{H}_{\bar{q}'q}(\alpha_s(Q)) = H_{\bar{q}'q}(\alpha_s(Q)) \cdot [g(\alpha_s(Q))]^{-1} \tag{3.42}$$

$$\exp \tilde{S}(Q,b) = \exp S(Q,b) \exp\left\{ - \int_{b_0/b}^{Q} \frac{dt}{t} \beta(\alpha_s(t)) \frac{d \log g(\alpha_s(t))}{d \log(\alpha_s(t))} \right\} \tag{3.43}$$

$$\tilde{C}_{ab}(\alpha_s(b_0/b), z) = C_{ab}(\alpha_s(b_0/b), z) \cdot [g(\alpha_s(b_0/b))]^{1/2} \tag{3.44}$$

with g an arbitrary function and where $\beta(\alpha_s)$ is the QCD β-function introduced in Eq. (2.7).

A first key observation for the following discussion is that by setting $g = H_{\bar{q}'q}$, the hard function contribution disappears and can be absorbed into a redefinition of the collinear functions and the Sudakov exponent. By looking at Eq. (3.44) it is clear that such a transformation is intimately related to the precise choice of the renormalization scale μ_{R} entering the strong coupling constant, which is set to b_0/b, i.e the Fourier conjugate of the p_{T} of the colour singlet. This is the same choice also done for the factorization scale μ_{F} (see Eq. (3.39)). It is worth noticing that, by exponentiating the hard function contribution with Eq. (3.43), the definition of the $A^{(n)}$ coefficients is unchanged. Moreover, since the β-function expansion starts at $\mathcal{O}(\alpha_s^2)$, the coefficient $B^{(1)}$ is also not affected by the symmetry transformation, while $B^{(n)}$ with $n > 1$ are changed according to:

$$\tilde{B}(\alpha_s) = B(\alpha_s) - \beta(\alpha_s) \frac{d \log H_{\bar{q}'q}(\alpha_s)}{d \log \alpha_s} \,. \tag{3.45}$$

Despite the exact control on the logarithmic accuracy provided by the resummation formula of Eq. (3.36) in its b space formulation, for any practical usage the transverse-momentum spectrum should be recovered by performing the inverse Fourier transformation, which is in general computationally very expensive [4]. Nevertheless, Eq. (3.36) is still extremely informative if we are not interested in the exact logarithmic accuracy of the $p_{\rm T}$ spectrum, but we just want to access its singular structure to a given perturbative order, which is precisely what we will need for the discussion of the following subsection. In this sense, a resummation formula in transverse-momentum space can be obtained, evaluating the b integral through an expansion of b_0/b around $p_{\rm T}$ at the level of the integrand:

$$\frac{d\sigma^{\rm sing}}{d\mathbf{\Phi}_{\rm B} dp_{\rm T}} = \sum_{q,q'} \frac{{\rm d}|M^{\rm B}|^2_{\bar{q}'q}}{{\rm d}\mathbf{\Phi}_{\rm B}} \frac{d}{dp_{\rm T}} \Big\{ [C_{\bar{q}'a}(\alpha_s) \otimes f_{a/h_1}](x_1, p_{\rm T}) \, [C_{qb}(\alpha_s) \otimes f_{b/h_2}](x_2, p_{\rm T}) \\ \cdot H_{\bar{q}'q}(\mathbf{\Phi}_{\rm B}, \alpha_s(p_{\rm T})) e^{S(Q, p_{\rm T})} \Big\} \tag{3.46}$$

where some of the resummation coefficients are properly redefined to retain the relevant perturbative accuracy (for the detailed derivation of the above equation see Ref. [175]).

As a final remark, it should be noted that for basic Drell-Yan processes as the ones considered in this section loop corrections have a particularly simple structure. All kinematic dependence simplifies in the virtual over Born ratio, which means that, once a particular renormalization scale choice is done, the coefficients $H^{(n)}_{\bar{q}'q}$ evaluate to constant factors. Moreover, no dependence on the flavour of the partonic channel survives in the ratio, which means that flavour indices can be safely dropped for these simple processes. Therefore, for the time being we will set $H_{\bar{q}'q}(\mathbf{\Phi}_{\rm B}, \alpha_s(p_{\rm T})) = H(\alpha_s(p_{\rm T}))$. The same simplified structure holds also when dropping the on-shell assumption for the gauge bosons and including the full kinematics of their decay products, as discussed in Ref. [175].

3.2.2 Merging with no merging scale

As we mentioned at the beginning of this section, the original MiNLO merging was proven not to preserve unitarity. That can also be shown following the reasoning of Ref. [175], which allows to identify the missing ingredients to reach NLO accuracy in the context of a `B` generator for observables inclusive over QCD radiation. The key idea is to connect the MiNLO formalism with transverse-momentum resummation.

By construction, Eq. (3.31) for $\bar{B}(\mathbf{\Phi}_{\rm BJ})$ provides the fully differential cross section which is NLO accurate for `BJ`, when the kinematics of the radiation is resolved. When integrating over the full $\mathbf{\Phi}_{\rm BJ}$ phase space, but keeping fixed the Born $\mathbf{\Phi}_{\rm B}$ kinematics and the boson $p_{\rm T}$, the singular structure of the process for $p_{\rm T} \to 0$ is captured. But the singular part of $pp \to B + j$ is also encoded in the resummation formula of Eq. (3.34),

[4] The ResBos2 code [188], successor of the ResBos code [189, 190] (which has recently received particular attention due to the CDF W mass measurement [191, 192]) can achieve N^{3}LL+NNLO accuracy in the transverse momentum of colour-singlet particles by making use of CSS resummation. In that case, for instance, an efficient Fourier transformation is achieved via sophisticated numerical integration methods [193].

with the singular spectrum given by Eq. (3.46). The singular contributions accounted for by the resummed spectrum are given in Eq. (3.35). Therefore, if the singular spectrum at a sufficient logarithmic accuracy is matched with the LO results for $pp \to B + j$ (encoded in the regular part of the spectrum), then NLO accurate results for inclusive observables can be obtained when the integration over p_{T} is performed. Since Eq. (3.31) already accounts for the non-singular part of the process at NLO, we just need to compare the MiNLO formula with the singular part of the resummed spectrum to see which terms are missing in order to achieve NLO accuracy in the jet-unresolved phase space.

We can define the luminosity $\mathcal{L}$ and rename the regular part of the resummed spectrum as follows:

$$\mathcal{L}(p_{\mathrm{T}}) \equiv \sum_{\bar{q}'q} \frac{\mathrm{d}|M^{\mathrm{B}}|^2_{\bar{q}'q}}{\mathrm{d}\mathbf{\Phi}_{\mathrm{B}}} [C_{\bar{q}'a}(\alpha_s) \otimes f_{a/h_1}](x_1, p_{\mathrm{T}}) \, [C_{qb}(\alpha_s) \otimes f_{b/h_2}](x_2, p_{\mathrm{T}}) H(\alpha_s(p_{\mathrm{T}})) \tag{3.47}$$

$$R_f(p_{\mathrm{T}}) \equiv \frac{d\sigma^{\mathrm{reg}}}{d\mathbf{\Phi}_{\mathrm{B}} dp_{\mathrm{T}}} \,, \tag{3.48}$$

where the dependence on $\mathbf{\Phi}_{\mathrm{B}}$ in the luminosity is understood. Then, if we perform the p_{T} integration over the spectrum in Eq. (3.34) down to the hadronization scale, we clearly obtain:

$$\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}} = \int_{\Lambda_{\mathrm{QCD}}}^{Q} dp_{\mathrm{T}} \frac{d}{dp_{\mathrm{T}}} \Big\{ \mathcal{L}(p_{\mathrm{T}}) e^{S(Q,p_{\mathrm{T}})} \Big\} + \int_{\Lambda_{\mathrm{QCD}}}^{Q} dp_{\mathrm{T}} R_f(p_{\mathrm{T}}) \approx \mathcal{L}(Q) + \int_{\Lambda_{\mathrm{QCD}}}^{Q} dp_{\mathrm{T}} R_f(p_{\mathrm{T}}) \tag{3.49}$$

where we have used $\exp S(Q, \Lambda_{\mathrm{QCD}}) \approx 0$, since the non-emission probability is highly suppressed over such a large energy range. Therefore, NLO accuracy for inclusive observables is retained if we have NLO accuracy on the luminosity $\mathcal{L}(Q)$, which means the $C^{(1)}_{ab}$ coefficients are included, and if we just have LO control on the regular part of the spectrum (contributing from order $\mathcal{O}(\alpha_s)$ with respect to the Born).

In order to make contact with the MiNLO formalism, we have to rewrite Eq. (3.46) by expanding the total p_{T} derivative. We can start by writing

$$\frac{d\sigma^{\mathrm{sing}}}{d\mathbf{\Phi}_{\mathrm{B}} dp_{\mathrm{T}}} = e^{S(Q,p_{\mathrm{T}})} D(p_{\mathrm{T}}) \quad \text{with} \quad D(p_{\mathrm{T}}) \equiv \frac{dS(Q,p_{\mathrm{T}})}{dp_{\mathrm{T}}} \mathcal{L}(p_{\mathrm{T}}) + \frac{d\mathcal{L}(p_{\mathrm{T}})}{dp_{\mathrm{T}}} \,, \tag{3.50}$$

where

$$\frac{dS(Q,p_{\mathrm{T}})}{dp_{\mathrm{T}}} = -\frac{1}{p_{\mathrm{T}}} \bigg(A(\alpha_s(p_{\mathrm{T}})) \log \frac{Q^2}{p^2_{\mathrm{T}}} + B(\alpha_s(p_{\mathrm{T}})) \bigg) \,. \tag{3.51}$$

Then, Eq. (3.34) can be recast as

$$\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}} dp_{\mathrm{T}}} = e^{S(Q,p_{\mathrm{T}})} \bigg\{ D(p_{\mathrm{T}}) + \frac{R_f(p_{\mathrm{T}})}{e^{S(Q,p_{\mathrm{T}})}} \bigg\} \,. \tag{3.52}$$

The term in curly brackets can now be expanded just by keeping the perturbative contributions that are needed to preserve NLO accuracy also after an integration over

the $p_{\rm T}$ of the boson, or, to say that differently, in such a way that the total $p_{\rm T}$ derivative can still be reconstructed to the desire perturbative accuracy. We know that for small transverse momenta $p_{\rm T} \sim \Lambda_{\rm QCD}$ we have large logarithmic contributions that invalidate the perturbative convergence. The dominant contributions are double logarithmic terms of order one like $\alpha_s L^2 \sim \mathcal{O}(1)$, so that, each logarithmic power parametrically counts as $L \sim \alpha_s^{-1/2}$. The modification of the α_s power counting close to the so called *Sudakov region* has to be taken into account in the expasion of Eq. (3.52): missing terms that would have been neglected by naive power counting might spoil the desired accuracy upon integration. More precisely, the following counting rule (proven in details in Appendix C of Ref. [175]) holds:

$$\int_{\Lambda_{\rm QCD}}^{Q} \frac{dp_{\rm T}}{p_{\rm T}} \alpha_s^n(p_{\rm T}) \log^m \frac{Q}{p_{\rm T}} e^{S(Q,p_{\rm T})} \approx [\alpha_s(Q)]^{n-\frac{m+1}{2}} , \tag{3.53}$$

which precisely applies to our case, where only one inverse power of $p_{\rm T}$ can appear, specifically in the derivative of the Sudakov in Eq. (3.50). According to Eq. (3.53), to preserve the NLO accuracy of the integral, i.e to account for all terms up to order α_s, the integrand involving the Sudakov derivative should contain all terms up to order α_s^2. Since powers m accompanying each logarithmic term can just take values in the finite set $\{0, 1\}$, that is sufficient to guarantee that neglected contributions are genuinely $\mathcal{O}(\alpha_s^2)$ after integration. On the other hand, since no logarithmically-enhanced contribution is present in the derivative of the luminosity and in R_f, in these cases the α_s counting proceeds as usual and NLO accuracy of the integral is maintained just by keeping all α_s terms. Therefore, in Eq. (3.50) the Sudakov factor should include the resummation coefficients $A^{(1)}$, $A^{(2)}$, $B^{(1)}$ and $B^{(2)}$, the perturbative expansions of the coefficient functions and the hard function should account for the coefficients $C^{(1)}$ and $H^{(1)}$, respectively, and R_f should be at least LO accurate.

The previous observation is crucial, since it tells us that the resummed expression of the singular spectrum in transverse-momentum space of Eq. (3.46) can be obtained from the corresponding b-space formulation at a starting logarithmic accuracy by keeping all terms up to order α_s^2. Before moving to transverse-momentum space, the b-space resummation formula has to be NNLL accurate, since the NNLL $B^{(2)}$ coefficient is required. Then, when moving from b to transverse-momentum space, the α_s^2 contributions emerging from the truncation of the resummed expression can be simply incorporated in a redefinition of some resummation coefficients. Within the set of coefficients relevant at this level, the only one that gets redefined is $B^{(2)}$:

$$B^{(2)\prime} = B^{(2)} + 2\zeta_3 (A^{(1)})^2 \tag{3.54}$$

where $\zeta_3 = 1.2020569\ldots$ is the Riemann zeta function $\zeta(z)$ evaluated at $z = 3$.

To make contact with MiNLO, it is crucial for the renormalization scale of the α_s coupling to be evaluated at $p_{\rm T}$, so that, only upon integrating over it, the FO value $\mu_{\rm R} = Q$ is recovered. As we have shown in the previous section, that corresponds to a precise resummation scheme choice. Indeed, by setting $\mu_{\rm R} = p_{\rm T}$ in $H(\alpha_s(p_{\rm T}))$ entering the luminosity, the hard function contribution can be entirely included in the Sudakov

exponent and in the coefficient function definitions, according to the transformations in Eq. (3.42) and Eq. (3.43). At the perturbative order we are considering, that amounts to a further redefinition of the $B^{(2)\prime}$ in Eq. (3.54) and $C^{(1)}$

$$\begin{aligned}\tilde{B}^{(2)} &= B^{(2)\prime} + 2\pi\beta_0 H^{(1)}\,,\\ \tilde{C}^{(1)}_{ab}(z) &= C^{(1)}_{ab} + \delta_{ab}\delta(1-z)\frac{1}{2}H^{(1)}\,,\end{aligned} \tag{3.55}$$

where we used Eq. (3.45). Clearly, the factorization scale in the PDFs should also be set to p_{T}, which is required by the exact resummation formula in Eq. (3.36) in b-space. As emerging from the running coupling equation in Eq. (2.9) and the DGLAP evolution formula of Eq. (2.21), different values for μ_{R} and μ_{F} induce differences with respect to the correct scale choice behaving as $\sim \alpha_s^2(\mu_{\mathrm{R/F}})\log(\mu_{\mathrm{R/F}}/p_{\mathrm{T}})$, which would spoil the NLO accuracy of the inclusive cross section after integration.

Even though R_f should just be known at LO to get NLO accuracy like in a `B` generator, MiNLO is a NLO merging approach based on a NLO `BJ` generator. So, we can identify the regular spectrum at second order with the NLO-accurate fixed-order result

$$R_f(p_{\mathrm{T}}) = \frac{d\sigma^{(\mathrm{NLO})}_{\mathrm{BJ}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}} - \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\left[\frac{d\sigma^{\mathrm{sing}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(1)} - \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2\left[\frac{d\sigma^{\mathrm{sing}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(2)}\,, \tag{3.56}$$

where the FO expansion up to $\mathcal{O}(\alpha_s^2)$ of the singular spectrum has been subtracted to avoid double counting in the matching with the resummed calculation. Indeed, $[X]^{(i)}$ denotes the coefficient of the i^{th} term in the perturbative expansion of the quantity X. For the NLO differential cross section, we clearly have:

$$\frac{d\sigma^{(\mathrm{NLO})}_{\mathrm{BJ}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}} = \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(1)} + \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(2)}\,. \tag{3.57}$$

Then, we can expand up to second order the $D(p_{\mathrm{T}})$ term in Eq. (3.52) and write

$$\begin{aligned}\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}} &= e^{S(Q,p_{\mathrm{T}})}\Bigg\{D^{(1)}(p_{\mathrm{T}}) + D^{(2)}(p_{\mathrm{T}}) + \frac{1}{e^{S(Q,p_{\mathrm{T}})}}\Bigg(\frac{d\sigma^{(\mathrm{NLO})}_{\mathrm{BJ}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}} - \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\left[\frac{d\sigma^{\mathrm{sing}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(1)}\\ &\quad - \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2\left[\frac{d\sigma^{\mathrm{sing}}}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(2)}\Bigg)\Bigg\} = \\ &= e^{S(Q,p_{\mathrm{T}})}\Bigg\{\frac{1}{e^{S(Q,p_{\mathrm{T}})}}\Bigg(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(1)} + \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(2)}\Bigg)\Bigg\} = \\ &\approx e^{S(Q,p_{\mathrm{T}})}\Bigg\{\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(1)}\left(1 - \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}S^{(1)}(Q,p_{\mathrm{T}})\right) + \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2\left[\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}dp_{\mathrm{T}}}\right]^{(2)}\Bigg)\Bigg\}\,,\end{aligned} \tag{3.58}$$

where in the last line again only terms up to $\mathcal{O}(\alpha_s^2)$ have been retained. If we use MiNLO for the calculation of an observable, as described by Eq. (2.44), we obtain

$$\begin{aligned}\langle O\rangle = &\int d\mathbf{\Phi}_{\mathrm{BJ}}\,\Delta^2_{\mathrm{MiNLO}}(Q,\,q)[B(\mathbf{\Phi}_{\mathrm{BJ}})(1 - 2\Delta^{(1)}_{\mathrm{MiNLO}}(Q,q)) + V(\mathbf{\Phi}_{\mathrm{BJ}})]\,O(\mathbf{\Phi}_{\mathrm{BJ}}) +\\ &+ \int d\mathbf{\Phi}_{\mathrm{BJJ}}\Delta^2_{\mathrm{MiNLO}}(Q,q)R(\mathbf{\Phi}_{\mathrm{BJJ}})O(\mathbf{\Phi}_{\mathrm{BJJ}})\,,\end{aligned} \tag{3.59}$$

with the reweighting factor $\alpha_s(q)/\alpha_s(Q)$ understood. Then, we recover the cross section differential in the variables $\{\mathbf{\Phi}_{\rm B}, p_{\rm T}\}$ simply by choosing

$$O(\mathbf{\Phi}) = \delta(\mathbf{\Phi}_{\rm B,res}(\mathbf{\Phi}) - \mathbf{\Phi}_{\rm B})\delta(p_{\rm T}(\mathbf{\Phi}) - p_{\rm T})\,. \tag{3.60}$$

In Eq. (3.60), we made use of the mapping $\mathbf{\Phi}_{\rm B,res}$ introduced in the previous section to project the phase space with resolved radiation such as $\mathbf{\Phi} = \{\mathbf{\Phi}_{\rm BJ}, \mathbf{\Phi}_{\rm BJJ}\}$ into the underlying Born phase space $\mathbf{\Phi}_{\rm B}$. With the notation $p_{\rm T}(\mathbf{\Phi})$ we also make clear that the boson transverse momentum defined in the initial phase space should be projected to a value $p_{\rm T}$ which depends on the treatment of radiation recoil.

By finally identifying $\exp S(Q,q) \equiv \Delta^2_{\rm MiNLO}(Q,q)$, we can effectively say that the last line of Eq. (3.58) is equivalent to the MiNLO formalism, where the resummation coefficients in the MiNLO Sudakov of Eq. (3.32) compare to the ones in $\exp S(Q,q)$ up to a rescaling by a factor of two [5].

At this level of the discussion, it is clear that, in order for MiNLO to respect unitarity, two main adjustments to Eq. (3.31) should be done. The factorization and renormalization scales should be set to the transverse momentum $p_{\rm T}$ of the boson rather than the scale of the first clustering. Moreover, the NNLL resummation coefficient $B^{(2)}$, which was missing in Eq. (3.32), should also be included. In Ref. [175], these modifications allowed to correctly account for the terms of order $\alpha_s^{3/2}$, which were not correctly included in MiNLO. The modified algorithm, often dubbed MiNLO′, allows to reach NLO accuracy for observables inclusive over QCD radiation by starting from a NLO `BJ` generator involving one resolved radiation, and to easily match to PS programs the second-to-hardest radiation via the Powheg method. Indeed, MiNLO′ is naturally embedded into the Powheg framework: making use of MiNLO′ in the context of Powheg simply amounts to modifying the $\bar{B}$ function of Eq. (3.12) as

$$\bar{B}(\mathbf{\Phi}_{\rm BJ}) = e^{\tilde{S}(p_{\rm T}(\mathbf{\Phi}_{\rm BJ}))}(B(\mathbf{\Phi}_{\rm BJ})\,(1-\tilde{S}^{(1)}(p_{\rm T})) + V(\mathbf{\Phi}_{\rm BJ})) + \int d\Phi_{\rm rad} e^{\tilde{S}(p_{\rm T}(\mathbf{\Phi}_{\rm BJJ}))} R_{\rm S}(\mathbf{\Phi}_{\rm BJ}, \Phi_{\rm rad}) \tag{3.61}$$

where the subtraction counterterms and the collinear remnants are still understood, together with the choice of $\mu_{\rm R} = \mu_{\rm F} = p_{\rm T}$. In the previous equation, we used $\tilde{S}$ to denote the new MiNLO′ Sudakov, where the first argument referring to the upper integration extreme Q has been dropped. It is important to stress that the transverse momentum of the boson entering the Sudakov factor multiplying the real term should be computed from the real phase space $\mathbf{\Phi}_{\rm BJJ}$, rather than $\mathbf{\Phi}_{\rm BJ}$, used for the Born and virtual contributions. It is also worth noticing that the mismatch in the two Sudakov factors does not spoil the IR cancellation of singularities, since the two definitions of $p_{\rm T}$ computed in $\mathbf{\Phi}_{\rm BJ}$ and $\mathbf{\Phi}_{\rm BJJ}$ agree in the singular regions of the secondary emission.

The benefit of this approach is that the accuracy of the NLO `B` generator is recovered without any need to combine `B` and `BJ` samples with an unphysical merging scale. Indeed,

[5] Notice that, for ease of reference and for consistency with the upcoming discussion carried out in Section 3.3, the correct MiNLO′ normalization for the resummation coefficients has already been implicitly assumed in Eq. (3.54), as well as in Eq. (3.55).

the presence of the MiNLO′ Sudakov allows to regulate the BJ divergences in the unresolved regions of the primary emission and, therefore, to push the merging scale down to the hadronization scale. In this respect, MiNLO′ differs from all other merging approaches. But to properly achieve merging, the resummation ingredients for the resolution variable are required as an input. In principle, any observable $\mathcal{X}$ such that $\mathcal{X} \to 0$ in the IR singular region of the radiation can be used as a resolution variable to separate different jet multiplicities, as soon as an accurate enough resummation for $\mathcal{X}$ is available. In our discussion we employed the p_{T} of the colour-singlet as a resolution variable, while in Ref. [194] the three-jet resolution parameter was used (with some of the resummation ingredients provided numerically within the CAESAR framework [104]). Anyway, contrary to other merging methods, where arbitrary jet muplitplicities can in principle be combined, MiNLO′ is applicable just to the case of 0- and 1-jet merging. An extension of the MiNLO method to more complex situations was suggested in Ref. [176], and it was used to merge up to three jet samples obtained from the generators H, HJ and HJJ. In such an extension, no analytic control on resummation is demanded, but unitarity is used to estimate the required resummation ingredients.

In order to introduce the main idea on which MiNLO′ is based, we restricted our discussion to the case of merging event generators where only one colourless particle is involved. In Ref. [195], the first application of MiNLO′ to a more general colour-singlet process was presented, where the single boson was replaced by a system of colourless particles (specifically a pair of W bosons). We will describe in the next section, in a slightly different context, the modifications that are needed to move to this more general case, which is extremely relevant for the results that we are going to present in Chapter 4.

3.3 The MiNNLO$_{\mathrm{PS}}$ method

To improve the accuracy of their predictions, event generators should rely on NNLO calculations for processes where they are available. The improvement in the normalization of the inclusive cross section is one important ingredient which is required to match the increasing accuracy of the experimental measurements. That poses the problem of reaching NNLO+PS accuracy, which means to match a NNLO calculation to PS. A NNLO+PS event generator should meet at least the following requirements: NNLO accuracy is recovered for observables inclusive over QCD radiation, NLO and LO accuracy is reached for quantities involving one or two resolved jets, respectively, and, finally, a proper shower matching which does not spoil the logarithmic accuracy of the PS itself has to be guaranteed. If NLO matching is nowadays standard and different NLO merging schemes are available, no general approach to reach NNLO+PS has been formulated yet. Just very few methods have been proposed in the last decades, which mainly address the issue for colour-singlet processes.

One of these methods is UNNLOPS, which has by now been applied to simple processes like Drell-Yan [196] or single Higgs production [197]. Starting from a BJ generator matched to PS with the MC@NLO approach, divergences associated to the primary emission are regulated putting a cutoff on the boson transverse momentum p_{T}. Then,

the NNLO calculation is added to the zero-$p_{\rm T}$ bin. In order to preserve unitarity and the desired accuracy for the different jet samples, the UNLOPS 'subtract-what-you-have-added' philosophy is invoked. This approach was the only one to be further extended to N^{3}LO+PS accuracy in Ref. [198] in the UN^{3}LOPS method. One of the main drawback of this family of approaches resides in the fact that they only include the two-loop virtual corrections in the zero-jet bin, contrary to the physical behaviour dictated by resummation.

In the GENEVA framework [85, 173], NNLO+PS accuracy has already been obtained for many colour-singlet processes, like Drell-Yan [199], Higgs-strahlung [200], photon pair production [201], hadronic Higgs decays [202], ZZ production [203], and Wγ production [204]. As already mentioned, merging is achieved by separating different jet samples using appropriate jet resolution variables, which are resummed to a logarithmic accuracy consistent with the fixed-order calculation. As soon as the required resummation ingredients are available, different observables can be used to partition the phase space. Therefore, even though the 0-jettiness τ_0, resummed in the SCET formalism, was mainly employed for hadronic processes, more recently [205] results using the $p_{\rm T}$ of the color-singlet system were presented. To obtain NNLO accuracy for the inclusive sample, the two-loop contributions are provided by increasing the accuracy of the resummation for the 0-jet resolution variable to NNLL$'$ (where the prime here indicates how the NNLL resummation also includes some N^{3}LL terms in such a way to incorporate the full singular structure in the resolution variable of the NNLO calculation). Similarly to UNNLOPS, divergences for the primary emission are regulated by a $\tau_{0,\rm cut}$ cutoff, with a non-local counterterm used to compensate the divergent behaviour of the real emission term (in line with phase space slicing subtraction, presented in Section 2.1.3). Therefore, contrary to the MiNLO$'$ approach, the merging scale $\tau_{0,\rm cut}$ can not be made arbitrary small without introducing large numerical instabilities. On the other hand, moving the $\tau_{0,\rm cut}$ scale to smaller and smaller values has the benefit of reducing the impact of missing power corrections arising from the region $\tau_0 < \tau_{0,\rm cut}$, which are not captured by the resummation. Indeed, as commented in Ref. [201], their correct inclusion into a term that matches resummation with the full NNLO calculation would require a fully-local NNLO subtraction procedure. Therefore, despite the fact that the usage of resummation correctly spreads two-loop terms also for $\tau_0 \neq 0$ and so cures the unphysical UNNLOPS behaviour, NNLO accuracy can just be claimed up to power corrections.

Very recently, a NNLO+PS matching scheme was proposed in the context of the VINCIA antenna-sector shower, where for the first time NNLO divergences are handled with the fully local antenna subtraction scheme [206]. Despite the novelty of the approach, its applicability is currently limited to leptonic collisions.

As we have already discussed, in MiNLO$'$ the problems with the merging scale are circumvented by using a Sudakov suppression factor. As proposed in Ref. [175], the MiNLO$'$ method can be easily upgraded into a NNLO+PS generator. Once the BJ-MiNLO output is available as a set of events which are NLO accurate both in presence of one resolved final-state parton and for unresolved radiation, a reweighting factor can

be applied to correct the normalization of the sample:

$$w(\mathbf{\Phi}_{\mathrm{B}}) = \frac{\left(\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}}\right)_{\mathrm{NNLO}}}{\left(\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}}\right)_{\mathrm{MiNLO'}}} \,. \tag{3.62}$$

By reweighting each event for $w(\mathbf{\Phi}_{\mathrm{B}})$, which is fully differential in the Born phase-space kinematics, NNLO accuracy can be correctly reached for inclusive observables. It is easy to show that the NLO accuracy of the original `BJ` generator is not spoilt. Indeed we can write:

$$\frac{\left(\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}}\right)_{\mathrm{NNLO}}}{\left(\frac{d\sigma}{d\mathbf{\Phi}_{\mathrm{B}}}\right)_{\mathrm{MiNLO'}}} = \frac{c_0(\mathbf{\Phi}_{\mathrm{B}}) + c_1(\mathbf{\Phi}_{\mathrm{B}})\alpha_s + c_2(\mathbf{\Phi}_{\mathrm{B}})\alpha_s^2}{c_0(\mathbf{\Phi}_{\mathrm{B}}) + c_1(\mathbf{\Phi}_{\mathrm{B}})\alpha_s + d_2(\mathbf{\Phi}_{\mathrm{B}})\alpha_s^2} \approx 1 + \frac{c_2(\mathbf{\Phi}_{\mathrm{B}}) - d_2(\mathbf{\Phi}_{\mathrm{B}})}{c_0(\mathbf{\Phi}_{\mathrm{B}})}\alpha_s^2 + \mathcal{O}(\alpha_s^3) \tag{3.63}$$

since MiNLO′ agrees with the NNLO result up to α_s terms. The wrong α_s^2 term $d_2(\mathbf{\Phi}_{\mathrm{B}})$ is then subtracted by the reweighting procedure and corrected with the right NNLO coefficient $c_2(\mathbf{\Phi}_{\mathrm{B}})$. Since events with resolved radiation are correct up to order $\mathcal{O}(\alpha_s^2)$, with dominant contributions starting at order α_s, the reweighting of these events includes at worse spurious terms of order α_s^3 or higher, which are beyond the original accuracy. This method has the advantage of being conceptually very simple, and it has been applied to many $2 \to 1$ LHC processes, such as Higgs-boson production [207], Drell-Yan [208], Higgs-strahlung [209, 210], which in terms of QCD corrections is still a $2 \to 1$ process, and to the $1 \to 2$ process $H \to b\bar{b}$ decay [194]. Anyway, when moving to more complex processes, the method becomes quickly unpractical. The main issue with this approach is that the reweighting factor of Eq. (3.63) has to be fully differential in the Born kinematics, so that the more complex the final state structure of the event is, the more computationally challenging the reweighting procedure becomes. In particular, in Ref. [211] this approach had been taken to its extreme by applying it to a genuine $2 \to 2$ process like W^+W^-. To this end, a number of features of the W-boson decays had to be resorted, in order to simplify the parametrization of the nine dimensional $W^+W^- \to e^+e^-\mu^-\bar{\nu}_\mu$ Born phase space, such as the fact that the full angular dependence of each vector-boson decay can be parametrized through eight spherical harmonic functions [212] or the fact that QCD corrections are largely independent of the off-shellness of the vector bosons. Moreover, the discretization of the residual variables in the parametrization of the Born phase space for the reweighting limits the numerical accuracy in phase-space regions sensitive to coarse bins. Not rarely, such regions can be relevant for BSM searches, especially when situated in the tails of kinematic distributions.

More recently, the latter method has been superseeded by a new extension of MiNLO′, dubbed MiNNLO$_{\mathrm{PS}}$. The MiNNLO$_{\mathrm{PS}}$ method is a very powerful approach that inherits all of MiNLO′ positive features and allows to reach NNLO+PS accuracy. The general idea was first introduced in Ref. [213] for $2 \to 1$ processes, with some optimizations presented in Ref. [214]. But the underlying idea of the approach applies even beyond Higgs or Drell-Yan production. By now, this method has been validated against many processes. In Ref. [215], NNLO+PS accuracy was achieved for the first time both in

production and decay for the Higgs-strahlung process, with the Higgs boson decaying into a bottom pair, where effects of SM effective field theory (EFT) operators have also been included in a more recent work [216]. In Ref. [217], the MiNNLO$_{\rm PS}$ methods was generalized to any colour-singlet process, paving the way for NNLO+PS simulations for all diboson processes. Indeed, the approach has already been successfully applied to Zγ production [217], also including the effect of anomalous triple gauge couplings [218], W^+W^- [219] and ZZ [220] processes. An even more recent application to $\gamma\gamma$ production has been presented in Ref. [221]. It is worth mentioning that MiNNLO$_{\rm PS}$ has been the first method which allowed to reach NNLO+PS accuracy even for colourful final states in Ref. [222, 223], with the relevant application to hadronic $t\bar{t}$ production. That has been achieved by properly extending the core idea of the method in order to account for QCD radiation from the final states, as well.

Despite the wide variety of applications, all MiNNLO$_{\rm PS}$-based methods share the following important features:

- NNLO corrections are calculated directly during the generation of the events, with no need for further reweighting or event post-processing: that allows to reach NNLO accuracy in a more efficient way compared to the MiNLO$'$ approach corrected by a NNLO reweighting factor;

- no merging scale is required to separate different jet multiplicities in the generated event samples, just as with MiNLO$'$: the absence of any slicing cutoff bypasses the problem of missing power corrections, which can threaten the NNLO accuracy of the inclusive cross section, and removes any dependence on unphysical parameters from the final results;

- when combined with transverse-momentum ordered parton showers, the matching preserves the leading logarithmic structure of the shower simulation [6]: this feature is a natural consequence of the fact that MiNNLO$_{\rm PS}$ ordering variable is the $p_{\rm T}$ of the colour-singlet system, which agrees, at LL level, with the $p_{\rm T}$ of the leading jet of Powheg, which is employed for the PS matching. As mentioned in different occasions, if the PS is ordered in a variable different than the transverse momentum of the radiation, preserving the LL accuracy of the latter can require the introduction of vetos to the shower radiation and of additional contributions, such as truncated showers [135].

The main goal of this section is to introduce the MiNNLO$_{\rm PS}$ method, with a special focus on its extension to any colur-singlet process, as presented in Ref. [217]. To this end, we will resume the discussion done for MiNLO$'$ starting from a more general perspective and then we will introduce the missing ingredients that are needed to promote MiNLO$'$ to a NNLO+PS generator.

[6] As explicitly mentioned in Ref. [223], the class of NLL corrections captured by the CMW scheme (see discussion in the Section 2.2) should also be preserved by the matching procedure. This is simply achieved in MiNLO$'$ by correctly including the terms $\alpha_s^2 \log^2(Q^2/p_{\rm T}^2)$ in the Sudakov exponent of Eq. (3.32), with $q = p_{\rm T}$.

3.3.1 Reaching NNLO+PS for colour-singlet processes

Similarly to MiNLO′, the key idea of MiNNLO$_{\rm PS}$ is to reach NNLO accuracy for inclusive observables by a careful comparison with the resummation formalism. Then, the $\bar{B}$ function of Eq. (3.61) is modified in such a way to account for all logarithmically-enhanced contributions at low transverse momentum to the desired perturbative accuracy. Finally, just as with MiNLO′, MiNNLO$_{\rm PS}$ is naturally embedded into the Powheg framework, which makes it straightforward to match with PS. In the following discussion, we will refer from the very beginning to a general colour-singlet system, denoted by F, with invariant mass Q.

Therefore, the starting point of MiNNLO$_{\rm PS}$ is a Powheg implementation of colour-singlet production in association with one jet within a `FJ` generator, with phase space $\mathbf{\Phi}_{\rm FJ}$. We can specialize to our case the Powheg master equation in Eq. (3.14) in a simplified notation as follows:

$$d\sigma_{\rm FJ} = \sum_{\ell_{\rm FJ}} \bar{B}_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})\, d\mathbf{\Phi}_{\rm FJ} \left\{ \Delta^{\ell_{\rm F}}_{\rm pow}(\mathbf{\Phi}_{\rm FJ}, p_{\rm T,min}) + \right. \tag{3.64}$$

$$\left. \sum_{\alpha_r \in \{\alpha_r | \ell_{\rm FJ}\}} \frac{\left[d\Phi_{\rm rad} \Delta^{\ell_{\rm F}}_{\rm pow}(\mathbf{\Phi}_{\rm FJ}, k_{\rm T}(\mathbf{\Phi}_{\rm FJJ})) R^{\rm sing}_{\ell_{\rm FJJ}}(\mathbf{\Phi}_{\rm FJJ}) \theta(k_{\rm T}(\mathbf{\Phi}_{\rm FJJ}) - p_{\rm T,min}) \right]^{\bar{\mathbf{\Phi}}^{(\alpha_r)}_{\rm FJ} = \mathbf{\Phi}_{\rm FJ}}_{\alpha_r}}{B_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})} \right\},$$

where the remnant and regular contributions have been dropped, since they do not have an active role in this discussion. As usual, $\bar{B}_{\ell_{\rm FJ}}$ describes the FJ process, which includes the first radiation, using the full NLO cross section. On the other hand, the content of the curly brackets accounts for the second QCD emission through the Powheg mechanism. In this context, $k_{\rm T}(\mathbf{\Phi}_{\rm FJJ})$ refers to the transverse momentum of the second radiation, computed in the full real phase space. Radiation beyond the second one is generated by the parton shower, which adds corrections of order $\mathcal{O}(\alpha_s^3(Q))$ and higher at all orders in perturbation theory. The shower cutoff $p_{\rm T,min}$ is set to the non-perturbative scale $\Lambda_{\rm QCD}$.

The derivation of the $\bar{B}_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})$ function in MiNNLO$_{\rm PS}$ [213] has once again as a starting point Eq. (3.46), just as with MiNLO′. In this context, the same equation is required to describe the singular part of the cumulative transverse-momentum spectrum $p_{\rm T}$ of a colour-singlet system up to NNLO (see coming discussion in this subsection), fully differentially in the Born phase space $\mathbf{\Phi}_{\rm F}$. Here, we report the same formula in a flavour-aware notation [217]:

$$\begin{aligned} \frac{{\rm d}\sigma}{{\rm d}\mathbf{\Phi}_{\rm F}{\rm d}p_{\rm T}} &= \frac{\rm d}{{\rm d}p_{\rm T}} \left\{ \sum_{\ell_{\rm F}} \exp[\tilde{S}_{\ell_{\rm F}}(p_{\rm T})] \mathcal{L}_{\ell_{\rm F}}(p_{\rm T}) \right\} + R_f(p_{\rm T}) \\ &= \sum_{\ell_{\rm F}} \exp[\tilde{S}_{\ell_{\rm F}}(p_{\rm T})]\, D_{\ell_{\rm F}}(p_{\rm T}) + R_f(p_{\rm T})\,, \end{aligned} \tag{3.65}$$

with $\ell_{\rm F}$ referring to the flavour structure of the underlying process $\bar{q}'q \to {\rm F}$ and where now

$$D_{\ell_{\rm F}}(p_{\rm T}) \equiv \frac{{\rm d}\tilde{S}_{\ell_{\rm F}}(p_{\rm T})}{{\rm d}p_{\rm T}} \mathcal{L}_{\ell_{\rm F}}(p_{\rm T}) + \frac{{\rm d}\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})}{{\rm d}p_{\rm T}}\,. \tag{3.66}$$

Introducing the flavour sum in Eq. (3.66) becomes necessary for a general colour-singlet process. Indeed, depending on the choice of resummation scheme, not only the luminosity function $\mathcal{L}_{\ell_{\rm F}}$, but also the Sudakov form factor $\tilde{S}_{\ell_{\rm F}}(p_{\rm T})$ can become in general flavour dependent. Therefore, we can write:

$$\mathcal{L}_{\ell_{\rm F}=\bar{q}'q}(p_{\rm T}) = \left(\tilde{C}_{\bar{q}'a}\otimes f_{a/h_1}\right)(x_1,p_{\rm T})\frac{\mathrm{d}\left[|M^{\rm F}|^2_{\bar{q}'q}\,\tilde{H}_{\ell_{\rm F}=\bar{q}'q}(p_{\rm T})\right]}{\mathrm{d}\mathbf{\Phi}_{\rm F}}\left(\tilde{C}_{qb}\otimes f_{b/h_2}\right)(x_2,p_{\rm T}) \quad (3.67)$$

$$\tilde{S}_{\ell_{\rm F}}(p_{\rm T}) = -2\int_{p_{\rm T}}^{Q}\frac{\mathrm{d}q}{q}\left(A(\alpha_s(q))\ln\frac{Q^2}{q^2}+\tilde{B}_{\ell_{\rm F}}(\alpha_s(q))\right), \quad (3.68)$$

where in Eq. (3.67) the sum over the a and b indices is understood. In Eq. (3.68), a factor of two has been made explicit compared to Eq. (3.40) to match the standard MiNNLO$_{\rm PS}$ convention for the definition of the resummation coefficients. With the notation $\ell_{\rm F}=\bar{q}'q$ we indicate the Born flavour structure $\ell_{\rm F}$ having $\bar{q}'q$ as initial-state partons. We clearly see from this equation that the main difference compared to $2\to 1$ hadronic processes, such as Higgs and Drell-Yan production, is that the hard function $H_{\ell_{\rm F}=\bar{q}'q}(p_{\rm T})$ for a general process can not be written as a simple form factor times the Born amplitude, and thereby it inherits a dependence on the respective flavour and phase-space configuration:

$$H_{\ell_{\rm F}}(\mathbf{\Phi}_{\rm F},p_{\rm T}) = 1+\left(\frac{\alpha_s(p_{\rm T})}{2\pi}\right)H^{(1)}_{\ell_{\rm F}}(\mathbf{\Phi}_{\rm F})+\left(\frac{\alpha_s(p_{\rm T})}{2\pi}\right)^2 H^{(2)}_{\ell_{\rm F}}(\mathbf{\Phi}_{\rm F})+\mathcal{O}(\alpha_s^3(p_{\rm T}))\,. \quad (3.69)$$

This dependence propagates throughout the formalism by means of the transformations in Eq. (3.43) and Eq. (3.44), which define a specific resummation scheme. We have already discussed that the MiNLO$'$ formalism is recovered by setting $\mu_{\rm R}=p_{\rm T}$, which is equivalent to a resummation scheme where the hard function contribution is absorbed in the Sudakov exponent and in the coefficient functions. This scheme choice has been signaled from the beginning of this section by redefining $S_{\ell_{\rm F}}\to\tilde{S}_{\ell_{\rm F}}$, $C_{\ell_{\rm F}}\to\tilde{C}_{\ell_{\rm F}}$, and $H_{\ell_{\rm F}}\to\tilde{H}_{\ell_{\rm F}}$. In particular, Eq. (3.45) tells us that the flavour and kinematic dependence in the Sudakov of Eq. (3.68) is entirely encoded in the B coefficient:

$$\tilde{B}_{\ell_{\rm F}}(\alpha_s) = \left(\frac{\alpha_s}{2\pi}\right)B^{(1)}+\left(\frac{\alpha_s}{2\pi}\right)^2\tilde{B}^{(2)}_{\ell_{\rm F}}+\mathcal{O}(\alpha_s^3)\,. \quad (3.70)$$

Following now the same reasoning outlined in the Subsection 3.2.2, we can rewrite Eq. (3.65) in a form where the Sudakov exponent is factored out:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\mathbf{\Phi}_{\rm F}\mathrm{d}p_{\rm T}} = \sum_{\ell_{\rm F}}\exp[\tilde{S}_{\ell_{\rm F}}(p_{\rm T})]\left\{D_{\ell_{\rm F}}(p_{\rm T})+\frac{R_f(p_{\rm T})}{e^{\tilde{S}_{\ell_{\rm F}}(p_{\rm T})}}\right\} \quad (3.71)$$

making clear that, if we want to reach NNLO accuracy and so to account for all α_s^2 terms after a $p_{\rm T}$ integration, all α_s^3 contributions in the curly brackets of Eq. (3.71) should be retained, according to the power counting rule of Eq. (3.53). Therefore, all required terms come from a N^{3}LL accurate resummation formula, which is properly expanded about small $p_{\rm T}$ values to just retain the full singular structure of the integrated spectrum at

order α_s^2. The ingredients that we need in this case are the coefficients $\{A^{(1)}, A^{(2)}, A^{(3)}\}$ and $\{B^{(1)}, B^{(2)}\}$ for the Sudakov exponent, together with the knowledge of the hard and coefficient functions up to second order. The expansion of the b-space resummation formula up to $\mathcal{O}(\alpha_s^3)$ (or to $\mathcal{O}(\alpha_s^2)$ at the level of the cumulative spectrum, where the p_{T} integration is carried out) leads to the redefinition in Eq. (3.54), accompanied by

$$H^{(2)\prime} = H^{(2)} - 2\zeta_3 A^{(1)} B^{(1)} \tag{3.72}$$

$$C_{ab}^{(2)\prime}(z) = C_{ab}^{(2)}(z) - 2\zeta_3 A^{(1)} P_{ab}^{(0)}(z) \tag{3.73}$$

with $P^{(0)}$ the leading order term of the regularized AP splitting function. When the required resummation scheme change is performed, we further have to apply the scheme conversion described in Eq. (3.55), supplemented by

$$\tilde{C}_{ab}^{(2)}(z) = C_{ab}^{(2)\prime}(z) + \frac{1}{2} H^{(1)} C_{ab}^{(1)} + \delta_{ab}\delta(1-z)\frac{1}{2}\left[H^{(2)\prime} - \frac{1}{4}(H^{(1)})^2\right], \tag{3.74}$$

where we see that the coefficients $\tilde{B}^{(2)}$, $\tilde{C}_{ab}^{(1)}$, and $\tilde{C}_{ab}^{(2)}$ get a phase-space and kinematic dependence through $H^{(1)}$.

Writing also the differential NLO cross section for FJ production as a sum over its flavour structures ℓ_{FJ}

$$\frac{\mathrm{d}\sigma_{\mathrm{FJ}}^{(\mathrm{NLO})}}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}} = \sum_{\ell_{\mathrm{FJ}}} \left\{ \frac{\alpha_s(p_{\mathrm{T}})}{2\pi} \left[\frac{\mathrm{d}\sigma_{\mathrm{FJ}}}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}}\right]_{\ell_{\mathrm{FJ}}}^{(1)} + \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2 \left[\frac{\mathrm{d}\sigma_{\mathrm{FJ}}}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}}\right]_{\ell_{\mathrm{FJ}}}^{(2)} \right\} \tag{3.75}$$

allows us to recast our starting formula in Eq. (3.71) in a form similar to Eq. (3.58), where terms up to order α_s^3 are properly included:

$$\begin{aligned}\frac{\mathrm{d}\sigma}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}} = \sum_{\ell_{\mathrm{FJ}}} \Bigg\{ \exp[\tilde{S}_{\ell_{\mathrm{F}}\leftarrow\ell_{\mathrm{FJ}}}(p_{\mathrm{T}})] \Bigg\{ \frac{\alpha_s(p_{\mathrm{T}})}{2\pi} \left[\frac{\mathrm{d}\sigma_{\mathrm{FJ}}}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}}\right]_{\ell_{\mathrm{FJ}}}^{(1)} \left(1 - \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}[\tilde{S}_{\ell_{\mathrm{F}}\leftarrow\ell_{\mathrm{FJ}}}(p_{\mathrm{T}})]^{(1)}\right) \\ + \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2 \left[\frac{\mathrm{d}\sigma_{\mathrm{FJ}}}{\mathrm{d}\boldsymbol{\Phi}_{\mathrm{F}}\mathrm{d}p_{\mathrm{T}}}\right]_{\ell_{\mathrm{FJ}}}^{(2)} \Bigg\} \Bigg\} + \sum_{\ell_{\mathrm{F}}} \exp[\tilde{S}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})]\, \mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}}) + \text{regular terms of } \mathcal{O}(\alpha_s^3)\,.\end{aligned} \tag{3.76}$$

In the previous expression $\ell_{\mathrm{F}} \leftarrow \ell_{\mathrm{FJ}}$ denotes the projection from the flavour structure of FJ production to the one of the Born process F. Note that the evaluation of $\exp[\tilde{S}_{\ell_{\mathrm{F}}\leftarrow\ell_{\mathrm{FJ}}}(p_{\mathrm{T}})]$ also requires a projection for $\boldsymbol{\Phi}_{\mathrm{FJ}} \to \boldsymbol{\Phi}_{\mathrm{F}}$ as discussed below. The regular terms originate from the expansion of $R_f(p_{\mathrm{T}})/e^{\tilde{S}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})}$, which can be safely omitted. Indeed, the latter terms are of order α_s^3 but they do not contain any singular contribution: after a p_{T} integration, they contribute beyond NNLO accuracy. The terms containing the relevant logarithmically-enhanced behaviour for low p_{T} values are contained in $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$. A minimal choice of $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ is such that it includes all α_s^3 terms of the $D_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ perturbative

expansion [213] (therefore neglecting contributions of order $\mathcal{O}(\alpha_s^4)$):

$$\begin{aligned}
\mathcal{D}_{\ell_{\rm F}}(p_{\rm T}) &\equiv \left(\frac{\alpha_s(p_{\rm T})}{2\pi}\right)^3 [D_{\ell_{\rm F}}(p_{\rm T})]^{(3)} = \left(\frac{\alpha_s(p_{\rm T})}{2\pi}\right)^3 \left\{ \left[\frac{d\tilde{S}_{\ell_{\rm F}\leftarrow\ell_{\rm FJ}}(p_{\rm T})}{dp_{\rm T}}\right]^{(1)} [\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(2)} \right. \\
&+ \left[\frac{d\tilde{S}_{\ell_{\rm F}\leftarrow\ell_{\rm FJ}}(p_{\rm T})}{dp_{\rm T}}\right]^{(2)} [\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(1)} + \left[\frac{d\tilde{S}_{\ell_{\rm F}\leftarrow\ell_{\rm FJ}}(p_{\rm T})}{dp_{\rm T}}\right]^{(3)} [\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(0)} + \left.\left[\frac{d\mathcal{L}_{\ell_{\rm F}\leftarrow\ell_{\rm FJ}}(p_{\rm T})}{dp_{\rm T}}\right]^{(3)} \right\} \\
&= \left(\frac{\alpha_s(p_{\rm T})}{2\pi}\right)^3 \left\{ \frac{2}{p_{\rm T}}\left(A^{(1)}\log\frac{Q^2}{p_{\rm T}^2} + B^{(1)}\right)[\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(2)} + \frac{2}{p_{\rm T}}\left(A^{(2)}\log\frac{Q^2}{p_{\rm T}^2} + \tilde{B}^{(2)}_{\ell_{\rm F}}\right)[\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(1)} \right. \\
&\left. + \frac{2}{p_{\rm T}}A^{(3)}\log\frac{Q^2}{p_{\rm T}^2}[\mathcal{L}_{\ell_{\rm F}}(p_{\rm T})]^{(0)} + \left[\frac{d\mathcal{L}_{\ell_{\rm F}\leftarrow\ell_{\rm FJ}}(p_{\rm T})}{dp_{\rm T}}\right]^{(3)} \right\}. \qquad (3.77)
\end{aligned}$$

We refer to appendix C and D of Ref. [213] and appendix A of Ref. [214] for explicit expressions relevant to the evaluation of the previous equation, where the flavour dependence can be simply included through the replacements $H^{(1)} \to H^{(1)}_{\ell_{\rm F}}$, $H^{(2)} \to H^{(2)}_{\ell_{\rm F}}$, and $\tilde{B}^{(2)} \to \tilde{B}^{(2)}_{\ell_{\rm F}}$.

To connect the MiNNLO$_{\rm PS}$ master equation in Eq. (3.76) to a Powheg `FJ` generator as described by Eq. (3.64), there is still one missing ingredient. Indeed, the $\bar{B}_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})$ contains information of the full $\mathbf{\Phi}_{\rm FJ}$ and $\mathbf{\Phi}_{\rm FJJ}$ kinematics, while in the $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ term, which comes from a resummation formalism, only carries information of the underlying Born kinematics $\mathbf{\Phi}_{\rm F}$ and of the $p_{\rm T}$ of the colourless system: all extra details on the kinematics of the radiation have been integrated out in Eq. (3.76). Therefore, in order to embed the $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ term into a `FJ` generator, the $\{\mathbf{\Phi}_{\rm F}, p_{\rm T}\}$ kinematics should be completed to the full $\mathbf{\Phi}_{\rm FJ}$ one via a proper mapping. The exact details of how that is done are beyond the formal generator accuracy, since at high $p_{\rm T}$ the $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ term adds contributions which are of α_s^3 order. Consequently, we just need the remapping to smoothly project $\mathbf{\Phi}_{\rm FJ}$ to $\mathbf{\Phi}_{\rm F}$ for $p_{\rm T} \to 0$. Once a projection $\mathbf{\Phi}_{\rm FJ} \to \mathbf{\Phi}_{\rm F}$ is provided, we have to distribute the NNLO corrections of $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ in the full $\mathbf{\Phi}_{\rm FJ}$ phase space with a proper spreading factor $F^{\rm corr}(\mathbf{\Phi}_{\rm FJ})$, defined as:

$$F^{\rm corr}_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ}) = \frac{J_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})}{\sum_{\ell'_{\rm F}} \int {\rm d}\mathbf{\Phi}'_{\rm FJ} J_{\ell'_{\rm FJ}}(\mathbf{\Phi}'_{\rm FJ})\delta(p_{\rm T} - p_{\rm T}{}')\delta(\mathbf{\Phi}_{\rm F} - \mathbf{\Phi}'_{\rm F})}\,, \qquad (3.78)$$

where $\mathbf{\Phi}_{\rm F}$ and $\mathbf{\Phi}'_{\rm F}$ result from the projection we mentioned (a possible mapping preserving the invariant mass and the rapidity of the colour-singlet system is described in Ref. [213]). In Eq. (3.78), the $p_{\rm T}$ is computed in the full space space $\mathbf{\Phi}_{\rm FJ}$. The arbitrariness in the spreading of the NNLO corrections is encoded in the choice of the function $J_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})$. Indeed, the only requirement that $F^{\rm corr}_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})$ should satisfy is that its integral over $\mathbf{\Phi}_{\rm FJ}$ in combination with a function that just depends on $\mathbf{\Phi}_{\rm F}$ and $p_{\rm T}$ has to reproduce the integral of that function over $\mathbf{\Phi}_{\rm F}$ and $p_{\rm T}$, consistently with the resummation formula encoded in Eq. (3.76). This requirement is trivially met by Eq. (3.78).

We are now in the position to connect Powheg and the parton shower matching formula in Eq. (3.64) with the MiNNLO$_{\rm PS}$ formalism. By making use of the spreading

factor $F^{\text{corr}}_{\ell_{\text{FJ}}}(\mathbf{\Phi}_{\text{FJ}})$, we can write the MiNNLO$_{\text{PS}}$ $\bar{B}(\mathbf{\Phi}_{\text{FJ}})$ function as

$$\bar{B}(\mathbf{\Phi}_{\text{FJ}}) \equiv \sum_{\ell_{\text{FJ}}} \Bigg\{ \exp[\tilde{S}_{\ell_{\text{F}}\leftarrow\ell_{\text{FJ}}}(p_{\text{T}})] \Bigg\{ \frac{\alpha_s(p_{\text{T}})}{2\pi} \left[\frac{\text{d}\sigma_{\text{FJ}}}{\text{d}\mathbf{\Phi}_{\text{FJ}}}\right]^{(1)}_{\ell_{\text{FJ}}} \left(1 - \frac{\alpha_s(p_{\text{T}})}{2\pi}[\tilde{S}_{\ell_{\text{F}}\leftarrow\ell_{\text{FJ}}}(p_{\text{T}})]^{(1)}\right) + \left(\frac{\alpha_s(p_{\text{T}})}{2\pi}\right)^2 \left[\frac{\text{d}\sigma_{\text{FJ}}}{\text{d}\mathbf{\Phi}_{\text{FJ}}}\right]^{(2)}_{\ell_{\text{FJ}}} \Bigg\} + \Bigg\{ \sum_{\ell_{\text{F}}} \exp[\tilde{S}_{\ell_{\text{F}}}(p_{\text{T}})]\, \mathcal{D}_{\ell_{\text{F}}}(p_{\text{T}}) \Bigg\} F^{\text{corr}}_{\ell_{\text{FJ}}}(\mathbf{\Phi}_{\text{FJ}}) \Bigg\}. \tag{3.79}$$

One should bear in mind that there is an implicit dependence on $\mathbf{\Phi}_{\text{F}}$ for all quantities that depend on ℓ_{F}. Therefore, obtaining the respective kinematics requires once again to make use of the projection $\mathbf{\Phi}_{\text{FJ}} \to \mathbf{\Phi}_{\text{F}}$.

To go one step further and embed the MiNNLO$_{\text{PS}}$ corrections within the Powheg-Box-Res framework, one needs some extra care to properly handle its resonance-aware features. Since the $\bar{B}(\mathbf{\Phi}_{\text{FJ}})$ function is modified in an additive way in Eq. (3.79), such that the NNLO corrections are treated on the same footing as all other contributions with $\mathbf{\Phi}_{\text{FJ}}$ kinematics, we can simply decompose the term involving $\mathcal{D}_{\ell_{\text{F}}}(p_{\text{T}})$ as a weighted sum over $\hat{\ell}_{\text{FJ}}$ using the weight functions $\mathcal{P}_{\hat{\ell}_{\text{FJ}}}$ of Eq. (3.18) just like in Eq. (3.16). With this simple adjustment, the resonance-aware MiNNLO$_{\text{PS}}$ master formula reads

$$\bar{B}(\mathbf{\Phi}_{\text{FJ}}) \equiv \sum_{\ell_{\text{FJ}}} \Bigg\{ \exp[\tilde{S}_{\ell_{\text{F}}\leftarrow\ell_{\text{FJ}}}(p_{\text{T}})] \Bigg\{ \frac{\alpha_s(p_{\text{T}})}{2\pi} \left[\frac{\text{d}\sigma_{\text{FJ}}}{\text{d}\mathbf{\Phi}_{\text{FJ}}}\right]^{(1)}_{\ell_{\text{FJ}}} \left(1 - \frac{\alpha_s(p_{\text{T}})}{2\pi}[\tilde{S}_{\ell_{\text{F}}\leftarrow\ell_{\text{FJ}}}(p_{\text{T}})]^{(1)}\right) + \left(\frac{\alpha_s(p_{\text{T}})}{2\pi}\right)^2 \left[\frac{\text{d}\sigma_{\text{FJ}}}{\text{d}\mathbf{\Phi}_{\text{FJ}}}\right]^{(2)}_{\ell_{\text{FJ}}} \Bigg\} + \sum_{\hat{\ell}_{\text{FJ}}\in T(\ell_{\text{FJ}})} \mathcal{P}_{\hat{\ell}_{\text{FJ}}} \Bigg\{ \sum_{\ell_{\text{F}}} \exp[\tilde{S}_{\ell_{\text{F}}}(p_{\text{T}})]\, \mathcal{D}_{\ell_{\text{F}}}(p_{\text{T}}) \Bigg\} F^{\text{corr}}_{\ell_{\text{FJ}}}(\mathbf{\Phi}_{\text{FJ}}) \Bigg\}, \tag{3.80}$$

where we have just highlighted the resonance decomposition of the NNLO corrections, which are the new ingredients of this subsection.

3.3.2 Some practical details

In this last subsection we present some details that are relevant for the actual implementation of MiNNLO$_{\text{PS}}$. The MiNNLO$_{\text{PS}}$ algorithm is nowadays available for many processes within the Powheg-Box and the Powheg-Box-Res [217] framework as a fully automatized feature. In order to use MiNNLO$_{\text{PS}}$ for a process implemented in Powheg-Box(-Res), one just needs to set the input card of the process, which contains the details on how to run the event simulation, in such a way to include the input line `minnlo 1`. The following discussion is important to better understand the setup used for the results presented in the next chapter.

As commented in the previous subsection, MiNNLO$_{\text{PS}}$ formula offers some flexibility in the precise choice of the spreading factor of Eq. (3.78) through the definition of $J_{\ell_{\text{FJ}}}(\mathbf{\Phi}_{\text{FJ}})$. A uniform spreading of the NNLO corrections in the $\mathbf{\Phi}_{\text{FJ}}$ phase space with $J_{\ell_{\text{F}}}(\mathbf{\Phi}_{\text{FJ}}) = 1$, despite being formally correct, would introduce spurious contributions in the high transverse-momentum regions of the leading jet. To have a more physical

spreading, one could choose $J_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ}) = |M_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})|^2 (f_{h_1} f_{h_2})_{\ell_{\rm FJ}}$, that is the Born matrix element for the FJ process weighted by the product of the PDFs corresponding to the initial-state flavour configuration in $\ell_{\rm FJ}$. With such a choice, NNLO corrections are correctly spread according to the pseudorapidity of the radiation. Clearly, the evaluation of the integral in the denominator of Eq. (3.78) for every phase-space point can easily become a bottleneck for complex processes. As described in Ref. [213], a suitable compromise between computational speed and physically sound results is to make use of the collinear approximation of $|M_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ})|^2$, which effectively amounts to setting $J_{\ell_{\rm FJ}}(\mathbf{\Phi}_{\rm FJ}) = P_{\ell_{\rm FJ}}(\Phi_{\rm rad})(f_{h_1} f_{h_2})_{\ell_{\rm FJ}}$, being $P_{\ell_{\rm FJ}}$ the appropriate AP splitting function. With this choice the numerical convergence of the integral in the denominator of Eq. (3.78) does not scale with the complexity of the process, which is a crucial requirement for multi-leg processes, such as diboson ones.

The Sudakov factor $\tilde{S}_{\ell_{\rm F}}(p_{\rm T})$ (together with its first order expansion $\tilde{S}^{(1)}_{\ell_{\rm F}}(p_{\rm T})$) and the NNLO $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ term modify the NLO accurate matching formula for the `FJ` generator by introducing contributions at high $p_{\rm T}$ which are beyond its formal accuracy. To reduce the impact of these spurious contributions in regions of phase space which are described at NLO accuracy, it is important to switch off the higher-order logarithmic terms for $p_{\rm T} \geq Q$. A standard choice (first presented in Ref. [213]) is to introduce some modified logarithms via the replacement

$$\log\left(\frac{Q}{p_{\rm T}}\right) \to L \equiv \frac{1}{p}\log\left(1+\left(\frac{Q}{p_{\rm T}}\right)^p\right), \tag{3.81}$$

with $p > 0$ (usually set to $p = 6$) controlling how fast the logarithms tend to zero for high $p_{\rm T}$ values. In Ref. [223] a more sophisticated choice was presented. Indeed, the details on how these logarithms [7] are turned off at high $p_{\rm T}$ do not impact on the accuracy of the simulation, as soon as the correct logarithmic structure is maintained at low $p_{\rm T}$. Moreover, as pointed out in Ref. [214], in order to preserve the total derivative of Eq. (3.65), regardless of the exact choice of the modified logarithms, that replacement requires three further adjustments. First of all, the lower integration bound of the Sudakov has to be replaced by $p_{\rm T} \to Qe^{-L}$. Secondly, $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ needs to be multiplied by a proper jacobian factor $\mathcal{J}_Q$, that depends on the exact form of L. For instance, with the choice done in Eq. (3.81) this jacobian factor reads

$$\mathcal{J}_Q = \frac{(Q/p_{\rm T})^p}{1+(Q/p_{\rm T})^p}, \tag{3.82}$$

[7] In practise, the form of the modified logarithms can be selected from the input card of the process by properly setting the input keyword `modlog_p`. By choosing a positive value, one selects the power of p entering Eq. (3.81). The values `-1d0` and `-2d0` trigger the usage of a piecewise-defined modified logarithm [223], which is the standard logarithm below $Q/2$ and exactly zero above Q. In-between those values two different functional forms may be chosen to smoothly interpolate between $Q/2$ and Q, keeping the modified logarithms and its derivative continuous. By setting `modlog_p 0`, the standard logarithm with a theta function is used, setting both the logarithmic contribution and the jacobian factor $\mathcal{J}_Q$ to zero for $p_{\rm T} > Q$. For a thourough discussion of this and other MiNNLO$_{\rm PS}$ input keywords necessary to run the algorithm, see also discussion in Appendix B.4.

which further switches off the effect of the $1/p_{\mathrm{T}}$ singularities. Finally, the perturbative scales need to be set consistently with the scale of the modified logarithms at small transverse momenta, i.e

$$\mu_{\mathrm{R},0} = Qe^{-L} \quad \mu_{\mathrm{F},0} = Qe^{-L}\,. \tag{3.83}$$

In Eq. (3.83), $\mu_{\mathrm{R},0}$ and $\mu_{\mathrm{F},0}$ refer to the central scale choices. Then, as standard, the size of missing higher order corrections is estimated by the 7-point scale variation approach of Eq. (2.26), by setting $\mu_{\mathrm{R}} = K_{\mathrm{R}}\mu_{\mathrm{R},0}$ and $\mu_{\mathrm{F}} = K_{\mathrm{F}}\mu_{\mathrm{F},0}$. But, as clarified in appendix D of Ref. [213], MiNNLO$_{\mathrm{PS}}$ formula includes extra sources of scale variations compared to a standard FO calculation, that arise from the Sudakov form factor. Scale variations within the Sudakov should be included for internal consistency of the approach, since the $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ term, which contains the derivative of $\tilde{S}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$, is evaluated at μ_{R}. Due to this feature, it is expected that MiNNLO$_{\mathrm{PS}}$ uncertainty bands arising from scale variation are generally wider that the corresponding ones from a NNLO calculation, despite the same formal accuracy for inclusive observables.

Moreover, MiNNLO$_{\mathrm{PS}}$ offers one extra possibility to further control the scale at which resummation effects are switched off, as first documented in Ref. [223]. In a way similar in spirit to resummation, the resummation scale Q can be rescaled by a factor K_{Q} by choosing the value of the input keyword `kappaq` accordingly.

In Ref. [214], a set of optimizations was introduced, going in the direction of reducing the sources of differences between a NNLO calculation and MiNNLO$_{\mathrm{PS}}$. Despite the NNLO accuracy reached by MiNNLO$_{\mathrm{PS}}$ for inclusive observables, the different treatment of higher-order corrections $\mathcal{O}(\alpha_s^3)$ can still have a sizable numerical impact when compared to FO results. Ref. [214] suggested three main adjustments to ameliorate the comparison. The first of them consists in the replacement of the expression for $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ of Eq. (3.77), which just included terms up to $\mathcal{O}(\alpha_s^3)$, with

$$\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}}) \equiv D_{\ell_{\mathrm{F}}}(p_{\mathrm{T}}) - \frac{\alpha_s(p_{\mathrm{T}})}{2\pi}[D_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})]^{(1)} - \left(\frac{\alpha_s(p_{\mathrm{T}})}{2\pi}\right)^2 [D_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})]^{(2)}\,. \tag{3.84}$$

Indeed, with the truncation carried out in Eq. (3.77) at the differential level, Eq. (3.76) does not reproduce anymore the exact total derivative of Eq. (3.65). With the new form of Eq. (3.84) (which is now the default choice within MiNNLO$_{\mathrm{PS}}$), the total derivative can be preserved, keeping into account also additional terms beyond accuracy of order $\mathcal{O}(\alpha_s^4)$ and higher. The better agreement between NNLO and MiNNLO$_{\mathrm{PS}}$ results, once the prescription of Eq. (3.84) is used, can be explained by keeping in mind that the impact of these subleading terms, which nonetheless include logarithmically enhanced contributions, can indeed be numerically important.

A second technical issue raised in Ref. [214] concerns the ability of Eq. (3.79) (or of Eq. (3.80)) to properly reproduce a NNLO result once the integration over p_{T} is carried out. In order for that to be the case, the p_{T} integration should be performed using a lower integration bound starting from sufficiently low values, where the integrand becomes vanishing due to the Sudakov suppression. This can be easily observed in Eq. (3.49), where the contribution to the integral of the total derivative evaluated at the lower bound

of integration is required to vanish. The success of this integration can be threaten by potentially too high internal PDF cutoff $\Lambda_{\rm PDF} > \Lambda_{\rm QCD}$. As discussed in Section 2.1.2, PDFs are non-perturbative objects whose values are obtained by appropriate fits of experimental measurements, whose results are saved into grids. Each set has its own infrared cutoff $\Lambda_{\rm PDF}$, which represents the lower scale at which the grids end due to lack of data points to carry out the fit procedure. That means that usual packages like LHAPDF [224], used to interface computer programs with PDF grids, would simply evaluate PDF functions to zero for scales $p_{\rm T} < \Lambda_{\rm PDF}$. To correctly perform the integration to even lower scales, Ref. [214] proposed to use a hybrid PDF set obtained as follows. As a first step, the PDF values read by LHAPDF are copied into HOPPET [225] grids. Indeed, the HOPPET package allows for an efficient manipulation of many operations involving PDFs. Up to this point, HOPPET grids exactly correspond to LHAPDF ones for scales $\mu_{\rm F} > \Lambda_{\rm PDF}$. Then, the HOPPET package is used to evolve PDFs for $\mu_{\rm F} < \Lambda_{\rm PDF}$ down to $\Lambda_{\rm QCD}$ according to DGLAP evolution equations. The possibility to perform the integration close to non-perturbative scales might require to smoothly approach the non-perturbative regime at small $p_{\rm T}$. That can be done by introducing a non-perturbative parameter Q_0 (whose value can be controlled from the input setting `Q0`) to regularize the Landau singularity. The introduction of this new parameter further modify Eq. (3.83), which now reads

$$\mu_{\rm R,0} = Qe^{-L} + Q_0\, g(p_{\rm T})\,, \qquad \mu_{\rm F,0} = Qe^{-L} + Q_0\, g(p_{\rm T})\,, \tag{3.85}$$

where $g(p_{\rm T})$ is chosen such that it suppresses the Q_0 shift at large values of $p_{\rm T}$. The analytic form of $g(p_{\rm T})$ is again to a large extend arbitrary. Once a choice is done, the scale setting of Eq. (3.85) should also be adopted in $\tilde{S}_{\ell_{\rm F}}(p_{\rm T})$ at the integrand level, and a jacobian factor $\mathcal{J}_{Q_0}$ (whose exact expression depends on $g(p_{\rm T})$) should multiply the derivative of the luminosity in Eq. (3.66).

The scale setting of Eq. (3.85) provides a consistent treatment in the small $p_{\rm T}$ region and preserves the total derivative of Eq. (3.65). However, at large $p_{\rm T}$ it yields $\mu_{\rm R,0} \sim \mu_{\rm F,0} \sim Q$, while a scale setting such that $\mu_{\rm R,0} \sim \mu_{\rm F,0} \sim p_{\rm T}$ might be preferred in that region. To this end, the scales entering the NLO F+jet cross section in Eq. (3.80) can be chosen from the input card (setting `largeptscales 1`) so that

$$\mu_{\rm R,0} = p_{\rm T} + Q_0\, g(p_{\rm T})\,, \quad \mu_{\rm F,0} = p_{\rm T} + Q_0\, g(p_{\rm T})\,. \tag{3.86}$$

It is important that Eq. (3.86) matches the scales of the Sudakov form factor and the $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ term at small $p_{\rm T}$. At the same time it ensures a dynamical scale choice of $\mu_{\rm R,0} = \mu_{\rm F,0} \sim p_{\rm T}$ at large $p_{\rm T}$.

One final observation done in Ref. [214] is strictly related to the PS matching of MiNNLO$_{\rm PS}$ formula. In order to reduce the effects of the shower on the kinematics of the colour-singlet system, a different shower recoil scheme was suggested. For the Pythia8 shower [124], that is activated by the flag `SpaceShower:dipoleRecoil 1` (see Ref. [226]). Its effect is to enable a local recoil scheme (i.e. one that does not affect the colour-singlet system) for initial-final colour dipole emissions, while keeping a global recoil (i.e. one that affects all final state particles including the colour-singlet system) for initial-initial

ones. As pointed out in Ref. [214], specific choices of different shower recoil schemes formally contribute beyond NNLO accuracy: the prescription adopted in that work had the benefit to reduce differences between the NNLO and the MiNNLO$_{\mathrm{PS}}$ descriptions of inclusive observables, especially in the colour-singlet high-rapidity tails. Nonetheless, a less global recoil can potentially affect the treatment of NLL terms entering the PS evolution and this fact should be kept in mind in view of future developments.

4 Accurate diboson phenomenology

In this chapter we present results for some relevant diboson processes obtained with the MiNNLO$_{\rm PS}$ method. Colour-singlet processes are among the LHC reactions which are measured at the most remarkable accuracy, with some measurements hitting percent-level uncertainties even for differential observables. Therefore, they offer a unique test to validate the MiNNLO$_{\rm PS}$ framework, which allows to interface the best available fixed-order predictions with full-fledged Monte Carlo event generators. In particular, vector-boson pair production processes have become an integral part of the rich precision programme at the LHC. Being measured by reconstructing the vector bosons from their leptonic decay products, those processes offer among the cleanest experimental signatures with rather small experimental uncertainties, so as to be considered *SM candle processes*. But their precise measurement is extremely relevant also in view of direct and indirect BSM searches. Indeed, small deviations of data from theory predictions in the shape of kinematic distributions might be a signal of new physics. Moreover, accurate theory predictions for vector-boson pair production are of great importance for improving the description of the background for other LHC measurements, like Higgs boson production.

Encouraged by all these strong motivations, we report in what follows a selection of results for accurate diboson simulations. We start from Section 4.1 with the description of $Z\gamma$ production, including some dedicated BSM studies in Section 4.2. Then, we discuss event generators for two massive diboson processes like W^+W^- and ZZ, respectively in Section 4.3 and 4.4.

4.1 Processes with hard photons: $Z\gamma$ production

Among vector-boson pair production processes, $Z\gamma$ is an extremely interesting LHC reaction. Indeed, especially in the $Z \to \ell^+\ell^-$ decay channel, it provides a particularly pure experimental signature as the final state can be fully reconstructed. In combination with its relatively large cross section, this process is well-suited for precision phenomenology. That explains the great interest in the measurement of this process at the LHC, which was performed at 7 TeV [227–232], 8 TeV [233–236], and even 13 TeV [237, 238], where Ref. [238] was the first diboson analysis to include the full Run II data set. On top of that, $Z\gamma$ has been extensively used as a probe for BSM physics. As we discuss in Section 4.2, measurements on anomalous triple and quartic gauge couplings are a strong test of the gauge structure of electroweak (EW) interactions and of the EW symmetry breaking mechanism. Furthermore, $Z\gamma$ final states are also relevant in direct searches for BSM resonances and in Higgs boson measurements (see for instance Refs. [239, 240]), with the SM production being an irreducible background. Although the Higgs decay into a $Z\gamma$

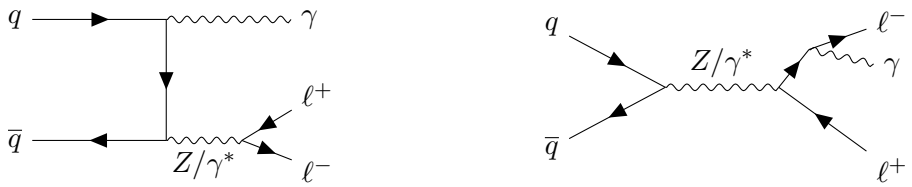

(a) q-type diagram (b) ℓ-type diagram

Figure 4.1: Sample LO diagrams for $\ell^+\ell^-\gamma$ production with two different resonance stuctures.

pair is rare, since it is loop induced in the SM, effects from new physics may significantly enhance this decay channel.

In this section we discuss the implementation of a `Zγ` and `ZγJ` generators as presented in Ref. [217]. Both generators have been implemented for the first time within the Powheg-Box-Res framework, introduced in Section 3.1.1. In particular, the `ZγJ` generator serves as starting point to reach NNLO accuracy for $Z\gamma$ production through the MiNNLO$_{\text{PS}}$ method.

We start by considering the production processes

$$pp \to \ell^+\ell^-\gamma \quad \text{and} \quad pp \to \ell^+\ell^-\gamma + \text{jet}\,, \tag{4.1}$$

where $\ell \in \{e, \mu\}$ is a massless charged lepton, which are relevant for $Z\gamma$ and $Z\gamma$+jet production, respectively.

As illustrated in Figure 4.11, $Z\gamma$ production is initiated by quark-antiquark annihilation at LO. The photon can be emitted either by the quark line (q-type diagrams) or by the lepton line (ℓ-type diagrams), each of which yields a different resonance structure of the respective amplitudes. Sample LO diagrams for $Z\gamma$+jet production are shown in Figure 4.2, with the same classification into q-type and ℓ-type diagrams. The distinction between those two resonance structures is relevant when treating them as two different resonance histories within the Powheg-Box-Res framework. In terms of the nomenclature introduced in Section 3.1.1, we can say that any bare flavour structure for $Z\gamma$ at Born level, for instance $\ell_{\text{B}} = \{u\bar{u} \to \ell^+\ell^-\gamma\}$, is associated to two full flavour structures $\hat{\ell}_{\text{B}}$, like $\hat{\ell}_{\text{B}} = \{u\bar{u} \to Z \to \ell^+\ell^-\gamma\}$ and $\hat{\ell}_{\text{B}} = \{u\bar{u} \to Z\gamma \to \ell^+\ell^-\gamma\}$. The latters are used to construct the weights $P_{\hat{\ell}_{\text{B}}}$ of Eq. (3.17), which in this specific case assume two different functional forms depending on whether $\hat{\ell}_{\text{B}}$ refers to q-type or ℓ-type diagrams, and they read

$$P_{\hat{\ell}_{\text{B}}} = \begin{cases} \frac{m_{\text{Z}}^2}{(s_{\ell\ell}-m_{\text{Z}}^2)^2+\Gamma_Z^2 m_{\text{Z}}^2} & \hat{\ell}_{\text{B}} \text{ is of } q\text{-type}\,, \\ \frac{m_{\text{Z}}^2}{(s_{\ell\ell\gamma}-m_{\text{Z}}^2)^2+\Gamma_Z^2 m_{\text{Z}}^2} & \hat{\ell}_{\text{B}} \text{ is of } \ell\text{-type}\,, \end{cases} \tag{4.2}$$

where $s_{\ell\ell}$ is the invariant mass of the lepton pair and $s_{\ell\ell\gamma}$ that of the $Z\gamma$ system.

In addition to the tree-level amplitudes at LO, the NLO calculation of the $Z\gamma$ ($Z\gamma$+jet) process requires the respective one-loop contributions as well as the tree-level real emission

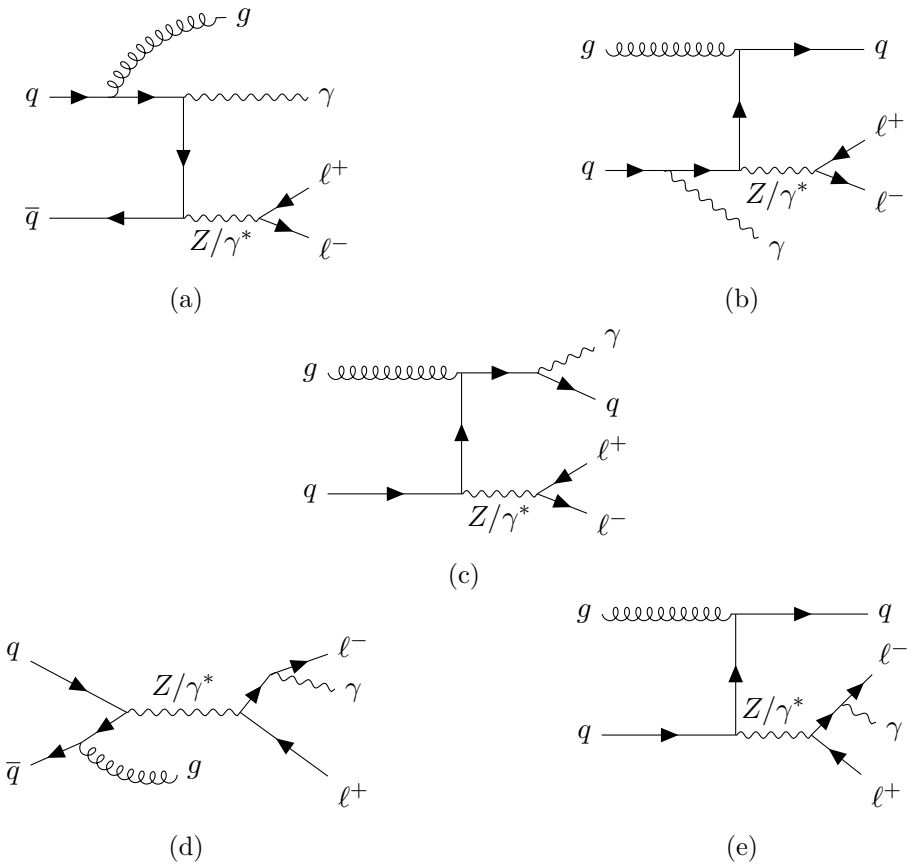

Figure 4.2: Sample LO diagrams for $\ell^+\ell^-\gamma$+jet production including q-type diagrams (a-c) and ℓ-type diagrams (d-e).

$Z\gamma$+jet ($Z\gamma$+2-jet) amplitudes. The NLO corrections to $Z\gamma$ and $Z\gamma$+jet production have been implemented within the POWHEG-BOX-RES framework. For the $Z\gamma$ generator the relevant amplitudes have been extracted from MCFM [1] [242], while for the $Z\gamma$+jet generator they have been implemented both using MCFM and via an interface to OPENLOOPS 2 [2] [245–247]. The helicity amplitudes of MCFM are implemented from the analytic expressions computed in Refs. [248, 249] for $Z\gamma$ production and in Ref. [250] for $Z\gamma$+jet production. For the MCFM amplitudes of $Z\gamma$+jet production the contribution from third generation quarks inside the loops has been entirely removed for those diagrams

[1] MCFM, that we have already recalled in Section 2.1.3, is a standalone framework for the computation of cross sections for many processes at NLO QCD, with some of them available at NNLO QCD or NLO EW. From version 10.0, N^{3}LL+NNLO can be achieved for single boson and diboson processes, as well, in the context of CUTE-MCFM [241].

[2] OPENLOOPS is a program for the fast evaluation of tree and one-loop amplitudes within the SM, allowing for predictions up to NLO in QCD and EW. Recent developments towards a two-loop automation have also been pursued in Ref. [243, 244].

where the Z boson is attached to a fermion loop through an axial vector coupling, while the massless bottom loop has been retained for those contributions where the corresponding top effects decouple as $1/m_{\text{top}}^4$ [251]. The impact of this approximation is expected to be rather small as shown in Ref. [252], where the leading heavy-quark loop contribution has been evaluated in the $1/m_{\text{top}}^2$ expansion in the context of Z+jet and Z+2-jet production. We further note that, in view of the NNLO calculation for $Z\gamma$ production discussed in Section 4.1.2, omitting the contribution of third generation quarks is in line with the fact that the heavy-quark loop contributions at two loops are currently not known, and therefore not included throughout our NNLO+PS results.

When using the OpenLoops interface for the NLO amplitudes of the $Z\gamma$+jet computation, the complex-mass scheme [3] [254] can be exploited and the full top-mass effects can be accounted for, while in the MCFM amplitudes the width is implemented only in a fixed-width scheme and heavy-quark loop effects are included only approximately, as described above. Since QED effects are included just at LO, the difference between the complex-mass and the fixed-width scheme amounts to an overall normalization, whose impact is below 0.1%. When comparing results with full top-mass effects as available in OpenLoops to approximate results as implemented in MCFM we find per mille effects for quantities inclusive over QCD radiation. This is expected, since heavy-quark effects at one loop are non-vanishing only in the presence of final-state radiation. For jet-related quantities, the differences between the results, as shown in Appendix A.3, are negligible at low transverse momentum and can range up to several tens of percent in the boosted region ($p_{\text{T}} \gtrsim m_{\text{top}}$). This is not surprising, since the process at hand involves s-channel fermion-loop contributions, which become more important in these phase-space regions. But for observables involving a jet our results are NLO accurate only, hence characterized by larger theoretical uncertainties. In summary, we find that mass effects are always much smaller than our quoted theoretical uncertainties and for the numerical studies performed in this work, which are not devoted to boosted regions, using approximate results for the heavy-quark mass effects is justified. Hence, because of the better numerical performance of MCFM, we use the MCFM amplitudes to obtain all results presented in what follows. Specifically, we find that the MCFM virtual $Z\gamma$+jet amplitudes are about ten times faster than the OpenLoops ones. On the other hand, by making use of the folding option in Powheg, that we have presented in Section 3.1.2, where the real contribution is evaluated multiple times for each virtual one, we could improve the numerical performance of the code using OpenLoops amplitudes. In fact, this is always the case when the virtual amplitudes constitute the bottleneck in the numerical evaluation. For greater flexibility, the release of the numerical code includes the option to choose between the OpenLoops and the MCFM implementations of the amplitudes, by swithing the input parameter `useOL` on or off, respectively.

[3] Compared to the fixed-width scheme, the complex-mass scheme makes use of complex masses for the vector bosons not just in their propagators, but also in the calculation of the Weinberg angle θ_W, defined by Eq. (1.13), which becomes a complex parameter in its turn. This scheme choice has the relevant theoretical feature of preserving gauge invariace in calculations involving resonances, which guarantees that matrix elements respect unitarity at high energies [253].

4.1.1 Treatment of the isolated photon

The emission of a soft or collinear photon from a quark or a charged lepton induces QED singularities. Therefore, processes with final-state photons, as the one considered here, require not only suitable criteria to isolate photons in the experimental analyses, but they also call for an IR-safe isolation procedure on the theory side. Since in the POWHEG-BOX framework fiducial cuts are usually applied at analysis level after parton showering, which modifies the kinematics of the final states, we discuss how to include photon-isolation requirements already at the event generation level in this framework to obtain IR-safe predictions.

To produce isolated photons in the final state there are two relevant mechanisms: the *direct* production in the hard process, which can be described perturbatively, and the production through *fragmentation* of a quark or a gluon, which is non-perturbative. The separation between the two production mechanisms in theoretical predictions is quite delicate, as sharply isolating the photon from the partons would spoil infrared (IR) safety. So-called *fragmentation functions* are required to absorb singularities related to collinear photon emissions in the latter production mechanism. Those functions are determined from data with relatively large uncertainties. On the other hand, Frixione's *smooth-cone isolation* of the photons [255] offers an IR-safe alternative that completely removes the fragmentation component. This substantially simplifies theoretical calculations of processes with isolated photons at higher orders in perturbation theory. Although the direct usage of smooth-cone isolation in experimental analyses is not possible due to the finite granularity of the calorimeters, data–theory comparisons are facilitated by tuning the smooth-cone parameters to mimic the fixed-cone isolation used by the experiments, as discussed for instance in Ref. [256].

So far, only few processes involving final-state photons have been implemented in the POWHEG-BOX framework: $W\gamma$ production [257] and direct photon production [258, 259]. These two generators make use of the photon fragmentation component [4]. In particular in Ref. [257] the hadron fragmentation into photons is modelled within POWHEG in combination with a QCD+QED shower. In this case the theoretical prediction can apply directly the photon isolation criteria imposed by the experiments to distinguish prompt photons taking part in the hard scattering process from possible background sources (such as photons from decay of π^0 mesons or from quark fragmentation). Those isolation criteria limit the hadronic activity in the vicinity of the photon by imposing

$$\sum_{\text{had}\in R_0} E_{\text{T}}^{\text{had}} < E_{\text{T}}^{\text{max}} \quad \text{with} \quad R_0 = \sqrt{\Delta\eta^2 + \Delta\phi^2}\,, \tag{4.3}$$

where the sum of the transverse energy $E_{\text{T}}^{\text{had}}$ of the hadrons inside a fixed cone of radius R_0 around the photon is constrained to be less than $E_{\text{T}}^{\text{max}}$.

In view of the NNLO extension of our calculation, introduced in Section 4.1.2, we instead rely on smooth-cone isolation [255] to turn off the fragmentation component

[4] Some recent results for photon pair production have been presented in Ref. [221], where the Frixione isolation prescription is used, just as in our approach.

and to deal with QED collinear singularities perturbatively in an IR-safe manner. In this case, phase-space configurations where the photon becomes collinear to a quark are removed while preserving IR safety by allowing arbitrarily soft QCD radiation within smoothly decreasing cones around the photon direction. In practice, this means that the smooth-cone isolation is implemented by restricting the hadronic (partonic) activity within every cone of radius $\delta = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < \delta_0$ around a final-state photon, where δ_0 sets the maximal cone size, by imposing the following condition

$$\sum_{\text{had/part}\in\delta} E_{\text{T}}^{\text{had/part}} \;\leq\; E_{\text{T}}^{\text{max}}(\delta) \;=\; E_{\text{T}}^{\text{ref}} \left(\frac{1-\cos\delta}{1-\cos\delta_0}\right)^n , \qquad \forall\, \delta \;\leq\; \delta_0 \,, \tag{4.4}$$

such that the total hadronic (partonic) transverse energy inside the cone is smaller than $E_{\text{T}}^{\text{max}}(\delta)$. The parameter n controls the smoothness of the isolation function and $E_{\text{T}}^{\text{ref}}$ is a reference transverse-momentum scale that can be chosen to be either a fraction ϵ_γ of the transverse momentum of the respective photon ($p_{\text{T},\gamma}$) or a fixed value (p_{T}^0),

$$E_{\text{T}}^{\text{ref}} \;=\; \epsilon_\gamma\, p_{\text{T},\gamma} \qquad \text{or} \qquad E_{\text{T}}^{\text{ref}} \;=\; p_{\text{T}}^0 \,. \tag{4.5}$$

In our calculations we impose smooth-cone isolation on the phase space of all $Z\gamma$+jet and $Z\gamma$+2-jet configurations. Furthermore, various technical phase-space cuts at event-generation level are necessary in order to obtain IR safe results. Those generation cuts and parameters of the smooth-cone isolation are given in Appendix A.1. They are chosen to be much looser than the ones eventually applied at analysis level after parton showering. We stress that, since we also employ suppression factors for the NLO squared amplitudes (as discussed in detail below), the resulting differential cross section times suppression factors vanishes in the singular regions, which will not pass fiducial cuts.

As commonly used in many POWHEG-BOX generators, we exploit the possibility to split the real squared matrix element R into a singular and a finite (remnant) contribution, that we discussed in Section 3.1. Such splitting improves the numerical performance of the code, especially as far as the efficiency of the event generation is concerned, in cases where the ratio R/B departs from its corresponding soft/collinear approximation, for instance in presence of Born zeros [260]. According to the definition of singular and remnant contributions of Eq. (3.11) for each singular region α_r of the real amplitude, we can write:

$$\begin{aligned} R &= \sum_{\alpha_r} R^{\alpha_r}(\mathbf{\Phi}_{\text{R}}) = \sum_{\alpha_r} \left[R^{\text{sing},\,\alpha_r}(\mathbf{\Phi}_{\text{R}}) + R^{\text{remn},\,\alpha_r}(\mathbf{\Phi}_{\text{R}}) \right] , \quad \text{with} \\ R^{\text{sing},\,\alpha_r} &= h^{\alpha_r} R^{\alpha_r}(\mathbf{\Phi}_{\text{R}}) \quad \text{and} \quad R^{\text{remn},\,\alpha_r}(\mathbf{\Phi}_{\text{R}}) = (1-h^{\alpha_r}) R^{\alpha_r}(\mathbf{\Phi}_{\text{R}}) \,, \end{aligned} \tag{4.6}$$

where $\mathbf{\Phi}_{\text{R}}$ is the real phase space and h^{α_r} is the already introduced damping factor. Compared to Eq. (3.11), the full flavour index has been understood. Only the singular contribution $\sum_{\alpha_r} R^{\text{sing},\,\alpha_r}(\mathbf{\Phi}_{\text{R}})$ is then exponentiated in the POWHEG Sudakov Δ_{pwg} and used to generate the first emission according to the POWHEG method in Eq. (3.14), while the finite remnant contribution $\sum_\alpha R^{\text{remn},\,\alpha_r}(\mathbf{\Phi}_{\text{R}})$ can be treated separately, by generating it with standard techniques and feeding it directly into the parton shower.

A standard damping factor [177] is used in both $Z\gamma$ and $Z\gamma$+jet generators, where $h^{\alpha_r} = 0$ when the real squared amplitude in a singular region is greater than five times its soft/collinear approximation, and $h^{\alpha_r} = 1$ otherwise. Additionally, to improve the numerical convergence, the $Z\gamma$+jet generator requires a special setting of the damping factor, which ensures that QED singularities appearing in the real squared matrix element are moved into the remnant contribution. Indeed, not all of the QED singular regions appearing in the real matrix elements have a singularity in their underlying Born amplitude. Accordingly, the associated real singularity is not mitigated by an overall Born suppression factor (as described in more detail below). To deal with this issue, we define, in each singular region α_r, the invariant mass m_{α_r} of the emitter-emitted pair of that singular region, which is the quantity that becomes small close to QCD singularities, and we use as a damping factor

$$\tilde{h}^{\alpha_r} = \frac{(m^2_{\alpha_r})^{-1}}{(m^2_{\alpha_r})^{-1} + c \sum_{i\in(q,\bar{q})} d_{i\gamma}^{-1}} \, h^{\alpha_r} \,, \tag{4.7}$$

where $c \in [0, 1]$ is a free parameter that we choose below, and the sum runs over all (initial- and final-state) quarks with

$$\begin{aligned} d_{i\gamma} &= p^2_{\mathrm{T},\gamma} \qquad \text{when } i \text{ is a quark in the initial state}\,, \\ d_{i\gamma} &= p_i \cdot p_\gamma \quad \text{when } i \text{ is a quark in the final state}\,. \end{aligned} \tag{4.8}$$

The splitting induced by the suppression factor in Eq. (4.7) is such that, when a QED singularity dominates, the event is included in $R^{\mathrm{remn},\,\alpha_r}(\mathbf{\Phi}_{\mathrm{R}})$ ($\tilde{h}^{\alpha_r} \to 0$), and, when the QCD singularity is dominant, the event is moved into $R^{\mathrm{sing},\,\alpha_r}(\mathbf{\Phi}_{\mathrm{R}})$ ($\tilde{h}^{\alpha_r} \to 1$). The numerical constant c controls the transition region between QED singularities in $d_{i\gamma}$ and QCD singularities in $m^2_{\alpha_r}$. Since the QED over QCD coupling ratio $\alpha_{\mathrm{QED}}/\alpha_{\mathrm{s}} \sim 0.1$, we use the value $c = 0.1$, which we have checked to be suitable for an efficient generation of events.

Finally, we exploit another feature of the POWHEG-BOX framework that allows us to improve the numerical convergence in the relevant phase-space regions, namely the suppression factors (that we briefly presented in Section 3.1.2). The usage of suppression factors is in a sense mandatory for processes having singularities at Born level, such as the QED singularities in $Z\gamma$ and $Z\gamma$+jet production, to avoid sampling large statistics in phase-space regions which are eventually removed by the fiducial cuts at analysis level. In order to obtain suitable integration grids that give more weight to the phase-space regions selected by the fiducial cuts, we have introduced a Born suppression factor that vanishes in singular regions related both to QCD and QED emissions. Its precise form is given in Appendix A.1. Since the real phase space is generated directly from the Born one in the POWHEG-BOX, the same factor is also applied to $R^{\mathrm{sing},\,\alpha_r}(\mathbf{\Phi}_{\mathrm{R}})$. For the remnant contribution $R^{\mathrm{remn},\,\alpha_r}(\mathbf{\Phi}_{\mathrm{R}})$, on the other hand, which is QCD regular, but is in our case QED singular, an analogous suppression factor has been chosen, that is given in Appendix A.1, as well.

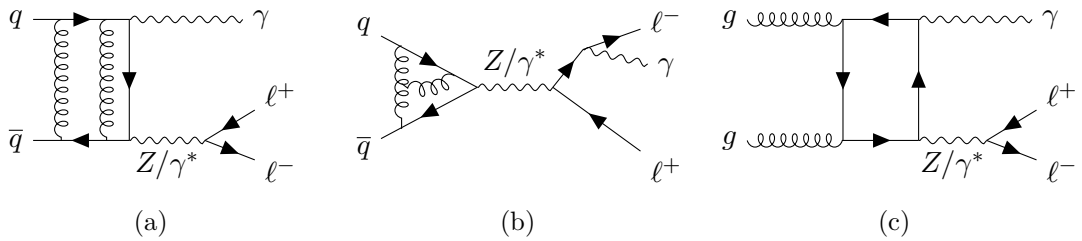

Figure 4.3: Sample Feynman diagrams entering the $\ell^+\ell^-\gamma$ process at NNLO: (a) q-type and (b) ℓ-type two-loop diagrams; (c) loop-induced gluon-fusion contribution.

4.1.2 Ingredients for NNLO+PS accuracy

The discussion so far allows to properly build a `Zγ` and a `ZγJ` generators to match a NLO calculation to PS within the POWHEG-BOX-RES framework. Moreover, the second generator can be used as a starting point to reach NNLO accuracy for observables inclusive over the primary emission by means of MINNLO$_{\text{PS}}$ [213], and in particular the MINNLO$_{\text{PS}}$ extension to genuine $2 \to 2$ hard-scattering of Ref. [217], presented in Section 3.3.1. Once the required resummation ingredients are provided to the desired perturbative accuracy, the whole formalism presented for MINNLO$_{\text{PS}}$ applies here by simply keeping in mind that F=Zγ (or more precisely F=$e^+e^-\gamma$) in this context.

In the introduction to this section we have discussed the contributions relevant to evaluate NLO corrections to $Z\gamma$ and $Z\gamma$+jet production. Those involve tree-level and one-loop amplitudes for the processes $pp \to Z\gamma$ and $pp \to Z\gamma$+jet as well as the tree-level amplitude for $pp \to Z\gamma$+2-jet. The same amplitudes enter the NNLO calculation for $Z\gamma$ production, i.e. at the Born level and as real, virtual one-loop, real-virtual and double-real corrections. The only missing ingredients for the NNLO calculation are the two-loop corrections in the $q\bar{q}$ channel, with sample diagrams shown in Figure 4.3 (a-b), and the loop-induced contribution in the gluon-fusion channel, with a sample diagram shown in Figure 4.3 (c). The latter is effectively only LO accurate and it can be separated from the others in a gauge-invariant way. Its size is rather small, being less than 10% of the NNLO corrections and below 1% of the full $Z\gamma$ cross section at NNLO [81]. We thus refrain from including the loop-induced gluon-fusion contribution in the calculation of this process: we will perform this excercise for a different process in Section 4.4.

For the two-loop corrections we use the $q\bar{q} \to \ell\ell\gamma$ helicity amplitudes calculated in Ref. [261], which have been fully implemented into the MATRIX framework [81, 82], which we have already recalled in the first part of Section 2.1.3. In order to exploit this implementation and evaluate the two-loop helicity amplitudes within our MINNLO$_{\text{PS}}$ calculation we have compiled MATRIX as a `C++` library and linked it to our POWHEG-BOX-RES $Z\gamma$+jet generator.

The one-loop and two-loop corrections to $Z\gamma$ directly enter the perturbative coefficients of the hard function of Eq. (3.69). The first and second order hard-virtual functions as provided by MATRIX are given in a subtraction scheme where IR singularities have

been subtracted according to the *q_T-scheme* (more precisely, choosing the *hard-scheme* of Ref. [262] as a resummation scheme). Those coefficients, that we denote as $H_{\ell_F}^{q_T(1)}$ and $H_{\ell_F}^{q_T(2)}$, are unambiguously defined in eqs. (12) and (62) of Ref. [262] and can be extracted from the one-loop and two-loop virtual amplitudes using the expressions of that paper. To use them within MiNNLO$_{\rm PS}$, we just need to perform a simple scheme conversion, which follows directly from the fact that the hard-scheme is defined such that the collinear coefficient functions do not contain $\delta(1-z)$ terms, which are absorbed into $H_{\ell_F}^{q_T}$ instead. That said, to match the MiNNLO$_{\rm PS}$ resummation scheme we just have to write

$$\begin{aligned} H^{(1)}_{\ell_F=cc'} &= 2H^{q_T(1)}_{\ell_F=cc'} - 4C^{(1),\delta}_{cc}\,, \\ H^{(2)}_{\ell_F=cc'} &= 4H^{q_T(2)}_{\ell_F=cc'} - 4(C^{(1),\delta}_{cc})^2 + 8C^{(2),\delta}_{cc} - 8C^{(1),\delta}_{cc}(H^{q_T(1)}_{\ell_F=cc'} - 2C^{(1),\delta}_{cc})\,, \end{aligned} \tag{4.9}$$

where $C^{(1),\delta}_{cc}$ and $C^{(2),\delta}_{cc}$ are the terms proportional to $\delta(1-z)$ of the first and second order coefficients of the collinear coefficient functions, which in the case of a quark-induced process ($cc' = q\bar{q}$) are given by

$$\tilde{C}^{(1),\delta}_{qq} = -C_F\frac{\pi^2}{24}\,, \tag{4.10}$$

$$\tilde{C}^{(2),\delta}_{qq} = \frac{9C_F^2\pi^4 + 2C_AC_F(4856 - 603\pi^2 + 18\pi^4 - 2772\zeta_3) + 4n_fC_F(-328 + 45\pi^2 + 252\zeta_3)}{10368}\,,$$

where $C_F = 4/3$ and $C_A = 3$, and n_f is the number of light quark flavours. Note that $\tilde{C}^{(1/2)}_{qq}$ in the MiNNLO$_{\rm PS}$ convention can be obtained from the ones of Ref. [262] by simply adding $\tilde{C}^{(1/2),\delta}_{qq} \times \delta(1-z)$.

Once the $D_{\ell_F}(p_T)$, which contains the NNLO corrections for $Z\gamma$, are available, a proper spreading factor $F^{\rm corr}(\mathbf{\Phi}_{Z\gamma J})$ is required to distribute the Born-like NNLO corrections in the $\mathbf{\Phi}_{Z\gamma J}$ kinematics. As we already mentioned, a key ingredient of $F^{\rm corr}(\mathbf{\Phi}_{Z\gamma J})$ is the definition of a phase-space projection $\mathbf{\Phi}_{Z\gamma J} \to \mathbf{\Phi}_{Z\gamma}$. For a process involving a final state photon, such mapping should be carefully defined. One should notice that the projection usually employed within MiNNLO$_{\rm PS}$ (see for instance appendix A of Ref. [213]) does not preserve the full Born kinematics: while it keeps all invariant masses and the rapidity of the $Z\gamma$ system unchanged, it does alter for instance the transverse momentum of the photon. As a result, the photon transverse momentum after the $\mathbf{\Phi}_{Z\gamma J} \to \mathbf{\Phi}_{Z\gamma}$ projection is neither bounded from below by the technical generation cuts nor controlled by the phase-space suppression factor introduced for the $\mathbf{\Phi}_{Z\gamma J}$ kinematics in Section 4.1.1. This induces a singular behaviour through the Born and virtual amplitudes in both the Sudakov form factor and the luminosity factor. We have therefore added the requirement $p_{T,\gamma} \geq 10\,\text{GeV}$ in the projected $\mathbf{\Phi}_{Z\gamma}$ kinematics. This technical cut is below the $p_{T,\gamma}$ threshold used at analysis level and it can be controlled through the input card. Its effect is strictly beyond accuracy, affecting only regular contributions at large $p_{T,\ell\ell\gamma}$. In fact, as discussed in Appendix A.2, our projection can lead to configurations with $p_{T,\gamma} \to 0$ only for events where the jet is back to back to the $Z\gamma$ system, and the Z and the photon are aligned with each other. It is then clear that $p_{T,j} > p_{T,\gamma}$ for such events and that there

are no large logarithms associated to $p_{\mathrm{T},j}$. Indeed, we have varied the cutoff down by a factor of ten, finding changes at the level of the numerical precision of less than a 1%.

For the treatment of higher-order terms and the details on the scale settings, we made use of some standard MiNNLO$_{\mathrm{PS}}$ features that we discussed in Section 3.3.2. The effect of the large logarithmic terms in the high transverse-momentum regions has been handled by making use of the modified logarithms of Eq. (3.81), with $p = 6$. Landau singularities have been regulated according to Eq. (3.85), with a choice of the non-perturbative scale $Q_0 = 0.5\,\mathrm{GeV}$. Moreover, we made use of the possibility to set the renormalization and factorization scales at large $Z\gamma$ transverse momentum according to the prescription of Eq. (3.86).

Finally, we conclude this section by reporting two further non-standard settings related to the showering of the $Z\gamma$ MiNNLO$_{\mathrm{PS}}$ events. First, we turn on by default the Powheg-Box `doublefsr` option, which was introduced and discussed in detail in Ref. [263]. When this option is turned on, both $q \to qg$ splittings and $g \to q\bar{q}$ splittings are treated symmetrically for the definition of the starting scale of the shower. This considerably reduces the appearance of spikes in distributions due to events with large weights that pass fiducial cuts after showering, and it ensures a proper treatment of observables sensitive to radiation off such configurations. Furthermore, for the Pythia8 shower [124], we set the flag `SpaceShower:dipoleRecoil 1`, in order to make use of a less global recoil scheme, which mildly affects the kinematics of the $Z\gamma$ system.

4.1.3 Phenomenological results: input and settings

In this section, we present NNLO+PS accurate predictions for $Z\gamma$ production for 13 TeV collisions at the LHC. Our results have been obtained by using the G_μ-scheme, where the electroweak coupling is defined as

$$\alpha_{G_\mu} = \frac{\sqrt{2} G_\mu m_{\mathrm{W}}^2 \sin^2\theta_W}{\pi}\,, \tag{4.11}$$

where $\cos^2\theta_W = m_{\mathrm{W}}^2/m_{\mathrm{Z}}^2$ and the input parameters are set to

$$\begin{aligned} &m_{\mathrm{W}} = 80.385\ \ \mathrm{GeV}\,, \quad m_{\mathrm{Z}} = 91.1876\ \mathrm{GeV}\,, \quad G_\mu = 1.16637 \times 10^{-5}\ \mathrm{GeV}^{-2}\,, \\ &\Gamma_W = 2.085\ \mathrm{GeV}\,, \qquad \Gamma_Z = 2.4952\ \mathrm{GeV}\,. \end{aligned} \tag{4.12}$$

We use $n_{\mathrm{f}} = 5$ massless quark flavours, and we choose the corresponding NNLO PDF set of NNPDF3.0 [264] with a strong coupling constant of $\alpha_s(m_{\mathrm{Z}}) = 0.118$. For the fixed-order predictions we use the PDF set at the respective order in QCD perturbation theory. To be precise, in the case of MiNLO′ and MiNNLO$_{\mathrm{PS}}$ the PDF grids are handled according to what described in Section 3.3.2: the grids are read using the LHAPDF interface [224] and copied into HOPPET grids [225]. Then, for scales below the internal PDF infrared cutoff the HOPPET package consistently performs the DGLAP evolution of the PDFs keeping the number of active flavours fixed to the one at the internal PDF infrared cutoff. The calculation of $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ in Eq. (3.84) requires the evaluation of different PDF convolutions and the computation of polylogarithms. For the latter we made use of the HPLOG package [265].

Our setting of the central renormalization and factorization scales ($\mu_{\mathrm{R},0}$, $\mu_{\mathrm{F},0}$) is in line with the MiNNLO$_{\mathrm{PS}}$ (MiNLO$'$) prescriptions, as we also summarized in the previous section. In order to validate the FO accuracy reached by MiNNLO$_{\mathrm{PS}}$ for inclusive observables, we also generate fixed-order NNLO results by exploiting the Matrix framework, where the loop-induced gluon-fusion contribution has been dropped. For all fixed-order results we adopt the following setting of the central renormalization and factorization scales:

$$\mu_{\mathrm{R},0} = \mu_{\mathrm{F},0} = m_{\ell\ell\gamma}\,, \tag{4.13}$$

where $m_{\ell\ell\gamma}$ is the invariant mass of the $Z\gamma$ system. The different scale choice between a NLO/NNLO fixed-order and a MiNLO$'$/MiNNLO$_{\mathrm{PS}}$ calculation induces effects beyond the nominal accuracy, in addition to the different treatment of higher-order terms. As a result, minor differences between the fixed-order and matched predictions are expected even for more inclusive observables. Nevertheless, the results should largely agree within scale uncertainties, at least in cases where scale uncertainties are expected to be a reliable estimate of missing higher-order corrections. For that, we employed the standard 7-point scale variations. We reiterate that, when performing scale variations for MiNLO$'$ and MiNNLO$_{\mathrm{PS}}$, additional sources of uncertainties, which are absent in a fixed-order calculation, are included through a scale dependence in the Sudakov form factor.

We have already discussed many of the technical details for the implementation of a $Z\gamma$+jet generator in previous sections, where we stressed that the usage of generation cuts and individual suppression factors at Born level for the singular real contributions and the remnant contributions is essential to get a good convergence of the Monte Carlo integration and an efficient event generation. In addition, folding of the radiation variables (ξ, y, ϕ) [5] has been used, with a choice of $(1, 5, 1)$ for the folding parameters, to evaluate the double-real correction ($Z\gamma$+2-jet) more often, which further improves the numerical convergence. Despite all those efforts to achieve a better numerical performance, we had to produce $\sim$ 100 million $Z\gamma$ events with our MiNNLO$_{\mathrm{PS}}$ generator to obtain acceptable statistical uncertainties and predict integrated cross sections in the fiducial setups considered here at the level of a few permille. Still, our comparison of differential distributions to NNLO predictions suffers from some fluctuations. That reflects the fact that $Z\gamma$ production is probably one of the most involved diboson process, featuring various complications, in particular considering its substantial complexity with respect to the QED singularity structure.

We employ the Pythia8 parton shower [124] with one of the A14 tunes [266] (specifically `py8tune 21`) to dress the hard event with further soft/collinear QCD radiation and use the default Powheg setting for the parton-shower starting scale. Hadronization effects are studied in Section 4.1.7 when comparing against data. Otherwise, the showered results

[5] These are the standard variables used within Powheg to parametrize the radiation phase space:

$$\xi = 2E_{\mathrm{rad}}/\sqrt{s} \quad y = \cos\theta \quad \phi \tag{4.14}$$

where E_{rad} is the radiation energy, $s = (k_\oplus + k_\ominus)^2$ the initial-state partonic energy, and finally θ and ϕ the emission and azimuthal angles with respect to a reference direction, respectively.

	`ATLAS-setup-1` [232]	`ATLAS-setup-2` [238]
Lepton cuts	$p_{\mathrm{T},\ell} > 25\,\mathrm{GeV} \quad \lvert\eta_\ell\rvert < 2.47$ $m_{\ell\ell} > 40\,\mathrm{GeV}$ –	$p_{\mathrm{T},\ell_1} > 30\,\mathrm{GeV} \quad p_{\mathrm{T},\ell_2} > 25\,\mathrm{GeV}$ $\lvert\eta_\ell\rvert < 2.47 \quad m_{\ell\ell} > 40\,\mathrm{GeV}$ $m_{\ell\ell} + m_{\ell\ell\gamma} > 182\,\mathrm{GeV}$
Photon cut	$p_{\mathrm{T},\gamma} > 15\,\mathrm{GeV} \quad \lvert\eta_\gamma\rvert < 2.37$	$p_{\mathrm{T},\gamma} > 30\,\mathrm{GeV} \quad \lvert\eta_\gamma\rvert < 2.37$
Separation cuts	$\Delta R_{\ell\gamma} > 0.7$ $\Delta R_{\ell,\,j} > 0.3 \quad \Delta R_{\gamma j} > 0.3$	$\Delta R_{\ell\gamma} > 0.4$ –
Jet definition	anti-k_{T} algorithm with $R = 0.4$ $p_{\mathrm{T},j} > 30\,\mathrm{GeV} \quad \lvert\eta_j\rvert < 4.4$	– –
Photon Isolation	Frixione isolation with $n = 1 \quad \epsilon_\gamma = 0.5 \quad \delta_0 = 0.4$	Frixione isolation with $n = 2 \quad \epsilon_\gamma = 0.1 \quad \delta_0 = 0.1$ $+ \quad E_{\mathrm{T}}^{\mathrm{cone0.2}}/p_{\mathrm{T},\gamma} < 0.07$

Table 4.1: Fiducial cuts in two different ATLAS setups denoted as `ATLAS-setup-1` and `ATLAS-setup-2`. See text for details.

do not include any effects from hadronization or underlying event models. Moreover, the photon is required to be generated only at the hard scattering level: contributions from a QED shower or the decay of unstable particles is not included. Finally, we keep photons stable by preventing any photon conversion effect, i.e. no $\gamma \to \ell^+\ell^-$ or $\gamma \to \bar{q}q$ splittings.

We present results for two sets of fiducial cuts, which are summarized in Table 4.1. The first one is identical to that used in Refs. [81, 267] and motivated by an earlier ATLAS analysis [232]. We refer to it as `ATLAS-setup-1` in the following. The second one was instead used in the most recent ATLAS 13 TeV measurement of Ref. [238] using the full Run-2 data and named `ATLAS-setup-2` in the following. We make use of `ATLAS-setup-1` only for validation purposes and to show the importance of NNLO+PS matching, while `ATLAS-setup-2` is also used to compare MiNNLO$_{\mathrm{PS}}$ predictions with the most updated experimental measurement available for $Z\gamma$ production. Both setups in Table 4.1 involve standard transverse-momentum and pseudorapidity thresholds to identify leptons and photons, as well as a lower invariant-mass cut on the lepton pair. `ATLAS-setup-2` places an additional requirement on the sum of the invariant masses of the $Z\gamma$ system and of the lepton pair. This cut suppresses the contribution from ℓ-type diagrams, where the photon is emitted from the final state leptons (cf. Figure 4.1(b)), enhancing t-channel production through q-type diagrams (cf. Figure 4.1(a)). Moreover, separation cuts between two particles $(i,\, j)$ are applied in $\Delta R_{ij} = \sqrt{\Delta\eta_{ij}^2 + \Delta\phi_{ij}^2}$, where $\Delta\eta_{ij}$ and $\Delta\phi_{ij}$ are their differences in the pseudorapidity and the azimuthal angle, respectively. In both setups leptons are separated from the isolated photon, while only `ATLAS-setup-1` imposes an additional separation of jets from leptons and from the isolated photon,

	ATLAS-setup-1 [232]		ATLAS-setup-2 [238]	
	$\sigma_{\text{inclusive}}$ [pb]	$\sigma/\sigma_{\text{NNLO}}$	$\sigma_{\text{inclusive}}$ [fb]	$\sigma/\sigma_{\text{NNLO}}$
LO	$1.5032(1)^{+11.2\%}_{-11.9\%}$	0.656	$271.83(2)^{+6.8\%}_{-7.8\%}$	0.508
NLO	$2.1170(5)^{+2.8\%}_{-4.3\%}$	0.924	$456.6(1)^{+3.6\%}_{-3.0\%}$	0.853
NNLO	$2.290(3)^{+0.9\%}_{-1.0\%}$	1.000	$535.3(6)^{+2.7\%}_{-2.5\%}$	1.000
MiNLO′	$2.222(8)^{+8.8\%}_{-11.0\%}$	0.970	$516(4)^{+8.8\%}_{-6.5\%}$	0.964
MiNNLO$_{\text{PS}}$	$2.299(5)^{+1.6\%}_{-1.4\%}$	1.004	$529(2)^{+4.0\%}_{-3.2\%}$	0.988
ATLAS	–		$533.7 \pm 2.1_{(\text{stat})} \pm 12.4_{(\text{syst})} \pm 9.1_{(\text{lumi})}$	

Table 4.2: Predictions for fiducial cross sections of $Z\gamma$ production at LO, NLO, and NNLO, as well as using the MiNLO′ and MiNNLO$_{\text{PS}}$ calculations, in the two ATLAS setups. For comparison, a column with the ratio to the NNLO cross section is shown. In the last row the ATLAS measurement of Ref. [238] is reported.

which in turn requires a jet definition [6]. As a consequence, we employ `ATLAS-setup-1` to study jet observables and show NLO/LO accuracy of MiNNLO$_{\text{PS}}$ predictions for $Z\gamma$+jet/$Z\gamma$+2-jet configurations. Finally, isolation criteria for the photon are needed, as detailed in Section 4.1.1, which is done by means of Frixione isolation in both setups. In `ATLAS-setup-2`, Frixione isolation is applied within a smaller cone and a second isolation criterium is added by requiring the scalar sum of the transverse energy of all stable particles (except neutrinos and muons) within a cone around the photon of size $R = 0.2$ ($E_{\text{T}}^{\text{cone0.2}}$) to be less than 7% of the photon transverse momentum (see Ref. [238] for more details). Note that we apply the latter isolation criterium only when analyzing events after parton showering, but not at Les-Houches-Event (LHE) or fixed-order level.

4.1.4 Fiducial cross sections

In Table 4.2 we report predictions for the $Z\gamma$ cross section in the two fiducial setups at LO, NLO and NNLO, and for MiNLO′ and MiNNLO$_{\text{PS}}$ matched to Pythia8. The fixed-order results have been obtained with Matrix [81, 82]. Although MiNNLO$_{\text{PS}}$ and NNLO calculations entail a different treatment of terms beyond accuracy, in both setups the agreement of their predicted cross sections is remarkably good. One should bear in mind, however, that in `ATLAS-setup-2` there is a slight difference in the treatment of the isolated photon at fixed order, which does not include the $E_{\text{T}}^{\text{cone0.2}}/p_{\text{T},\gamma} < 0.07$ cut, as discussed in the previous section. We further notice from Table 4.2 that the scale

[6] For these and the other results presented in this work, jets are constructed using the standard anti-k_{T} algorithm [268] with a jet radius of $R = 0.4$ as implemented in the FastJet package [269] (see Table 4.1). Indeed, FastJet is nowadays the standard program when it comes to performing jet manipulations. We want to take this opportunity to also stress that `ATLAS-setup-1` is inclusive over jets and that jet-separation cuts are only applied when a jet is present.

uncertainties of the NNLO and MiNNLO$_{\text{PS}}$ predictions are larger in `ATLAS-setup-2`. This is caused by the additional $m_{\ell\ell} + m_{\ell\ell\gamma}$ cut and the stronger cut on the photon transverse momentum in that fiducial setup, rendering the predictions more sensitive to additional QCD radiation, that is effectively described at a lower perturbative accuracy.

Comparing MiNNLO$_{\text{PS}}$ and MiNLO$'$ predictions, the inclusion of NNLO corrections through MiNNLO$_{\text{PS}}$ has a relatively moderate effect for the fiducial cross section of $+3.5\%$ in `ATLAS-setup-1` and $+2.5\%$ in `ATLAS-setup-2`. In fact, in both cases (and particularly evident for the latter setup) MiNLO$'$ predictions are actually closer to the NNLO results than to the NLO ones. After all, the Sudakov form factor in Eq. (3.68) is essentially the same for MiNNLO$_{\text{PS}}$ and MiNLO$'$, and MiNLO$'$ predictions already contain various contributions beyond NLO accuracy, including all real corrections at NNLO through the merging of NLO corrections to $Z\gamma$+jet production. Still, by reaching NNLO accuracy through the MiNNLO$_{\text{PS}}$ procedure the predictions get even closer to the NNLO results and the uncertainty bands substantially decrease, by almost a factor of ten in `ATLAS-setup-1` and by more than a factor of two in `ATLAS-setup-2`. Indeed, the MiNNLO$_{\text{PS}}$ scale uncertainties are comparable with the NNLO ones. The fact that they are slightly larger is expected since the MiNNLO$_{\text{PS}}$ procedure probes lower scales both in the PDFs and in α_s, and it includes scale variations also for the Sudakov form factor.

Finally, we find an excellent agreement of our NNLO+PS accurate MiNNLO$_{\text{PS}}$ predictions with the cross section measured by ATLAS in Ref. [238], which are perfectly compatible within the quoted experimental errors.

4.1.5 Comparison against MiNLO$'$ and NNLO

We now turn to discussing differential distributions in the fiducial phase space. In this section we compare our MiNNLO$_{\text{PS}}$ predictions with MiNLO$'$ and NNLO results. This serves two purposes. On the one hand, MiNNLO$_{\text{PS}}$ distributions are validated for one-jet and two-jet observables against the ones obtained with MiNLO$'$ and for Born-level observables (inclusive over QCD radiation) against NNLO predictions. On the other hand, this allows us to show the importance of NNLO+PS matching with respect to less accurate results. To these ends, we discuss selected distributions which are particularly significant to show the performance of MiNNLO$_{\text{PS}}$ predictions. The figures of this subsection, together with those of Subsections 4.1.6 and 4.1.7 are organized as follows: the main frame shows the predictions from MiNNLO$_{\text{PS}}$ matched to Pythia8 (blue, solid line), together with all other results relevant for the given comparison. In an inset we display the bin-by-bin ratio of all the histograms that appear in the main frame to the MiNNLO$_{\text{PS}}$ one. The bands indicate the theoretical uncertainties that are computed from scale variations.

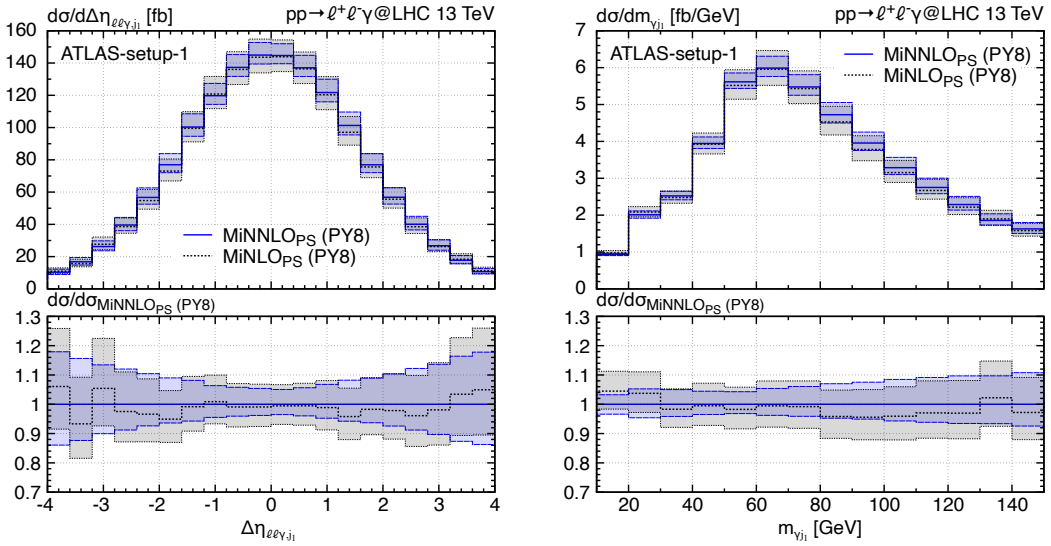

Figure 4.4: Distribution in the pseudorapidity difference of the colour-singlet and the hardest jet (left plot) and in the invariant mass of the photon and the hardest jet (right plot) for MiNNLO$_{\text{PS}}$ (blue, solid line) and MiNLO$'$ (black, dotted line).

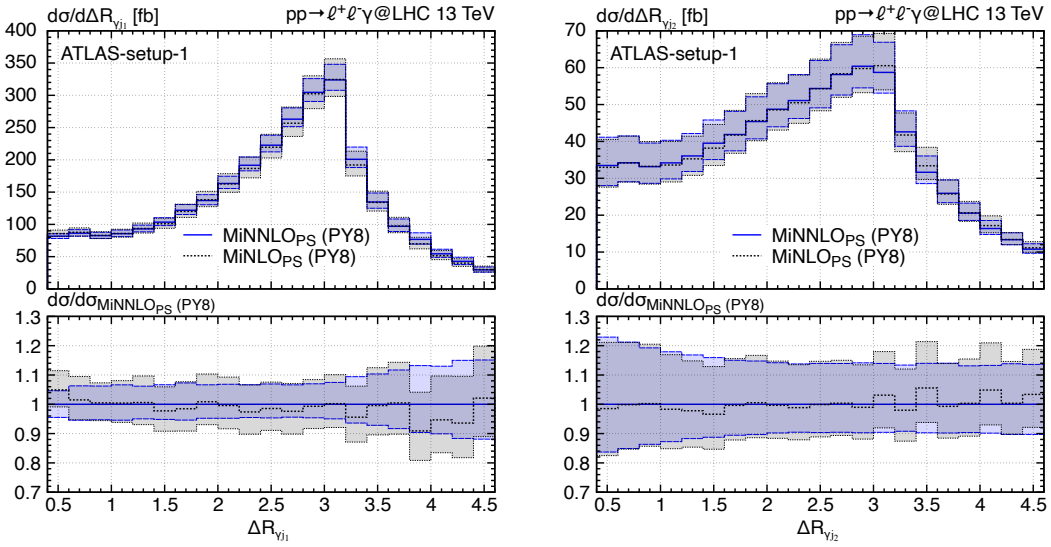

Figure 4.5: Distribution in the ΔR separation between the photon and hardest jet (left plot), and between the photon and the second-hardest jet (right plot) for MiNNLO$_{\text{PS}}$ (blue, solid line) and MiNLO$'$ (black, dotted line).

We start by discussing quantities that involve jets in the final state in Figure 4.4 and Figure 4.5, where MiNNLO$_{\text{PS}}$ (blue, solid line) and MiNLO$'$ (black, dotted line) matched to Pythia8 are compared in `ATLAS-setup-1`. Since for these observables MiNLO$'$ and MiNNLO$_{\text{PS}}$ have the same formal accuracy, the two predictions are expected not to differ from each other significantly (i.e. not beyond uncertainties), both in terms of shapes and size of scale uncertainty bands. In particular, such agreement serves as a validation that NNLO corrections are properly spread by the factor in Eq. (3.78) in the jet-resolved phase space of $Z\gamma$+jet production without altering its NLO accuracy. As a matter of fact, the left plot of Figure 4.4 shows that MiNNLO$_{\text{PS}}$ and MiNLO$'$ predictions agree well within uncertainties for the pseudorapidity difference between the $Z\gamma$ system and the hardest jet ($\Delta\eta_{\ell\ell\gamma,\,j_1}$). Furthermore, the size of the uncertainty bands are comparable over the whole pseudorapidity range. In a similar manner, the ratio between MiNNLO$_{\text{PS}}$ and MiNLO$'$ is nearly flat for the invariant mass of the photon and the hardest jet ($m_{\gamma j_1}$) in the right plot of Figure 4.4. Here, we further observe the effect of the Frixione isolation, which dampens the distribution in the photon-jet collinear limit. Also, in Figure 4.5 MiNNLO$_{\text{PS}}$ and MiNLO$'$ predictions agree well for the distance between the photon and the leading and subleading jet in the η-ϕ plane ($\Delta R_{\gamma j_1}$ and $\Delta R_{\gamma j_2}$). As $\Delta R_{\gamma j_2}$ involves the second-hardest jet, both MiNNLO$_{\text{PS}}$ and MiNLO$'$ are only LO accurate, which is also evident from the broadening of the uncertainty bands. We have examined a large number of other quantities involving jets (not shown here) observing a similar behaviour in all cases.

Next, in Figure 4.6, Figure 4.7 and Figure 4.8, we compare MiNNLO$_{\text{PS}}$ (blue, solid line) against MiNLO$'$ (black, dotted line) and NNLO predictions from Matrix [81, 82] (red, dashed line) for Born-level observables (inclusive over QCD radiation). By and large, we observe a very good agreement of MiNNLO$_{\text{PS}}$ and NNLO predictions, especially considering the fact that they differ from each other in the choice of the renormalization and factorization scales, and in the treatment of higher-order contributions. What can be appreciated is the clear reduction of the scale uncertainties of MiNNLO$_{\text{PS}}$ predictions with respect to the MiNLO$'$ ones up to a size which is comparable to the NNLO ones. In particular, Figure 4.6 displays the pseudorapidity distribution of the $Z\gamma$ system ($\eta_{\ell\ell\gamma}$) in each of the two fiducial setups. The ratio of NNLO over MiNNLO$_{\text{PS}}$ is close to one in both cases, with uncertainty bands of one to two percent in `ATLAS-setup-1`. In `ATLAS-setup-2`, on the other hand, the bands are roughly twice as large, as already observed for the integrated cross section, which is due to the higher sensitivity to phase-space regions related to real QCD radiation.

Then, in Figure 4.7 we show the distributions in the invariant mass ($m_{\ell\ell}$) and transverse momentum ($p_{\text{T},\ell\ell}$) of the lepton pair in `ATLAS-setup-1`. The qualitative behaviour of MiNNLO$_{\text{PS}}$ with respect to MiNLO$'$ and NNLO predictions in the ratio inset is relatively similar to the one of the $Z\gamma$ rapidity distribution. We can appreciate the Z-boson resonance in the $m_{\ell\ell}$ distribution as well as a broader, but smaller, enhancement around

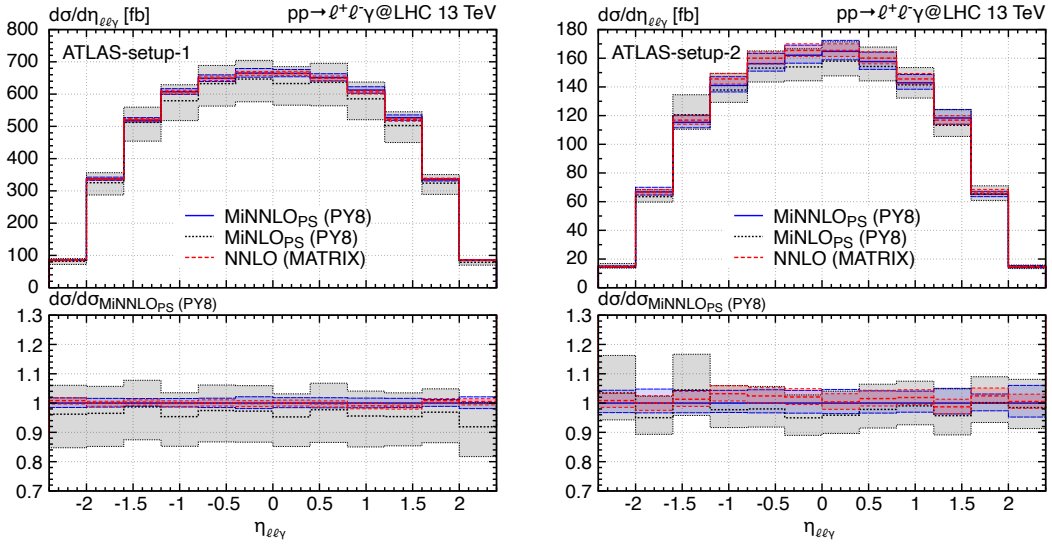

Figure 4.6: Distribution in the pseudorapidity of the $Z\gamma$ system in `ATLAS-setup-1` (left plot) and in `ATLAS-setup-2` (right plot) for MiNNLO$_{\rm PS}$ (blue, solid line), MiNLO$'$ (black, dotted line) and NNLO (red, dashed line).

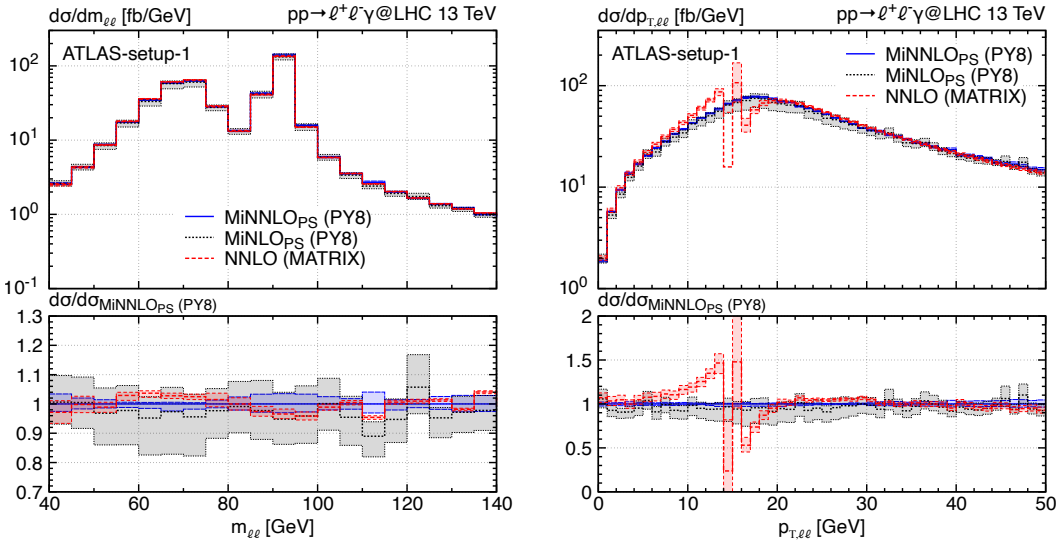

Figure 4.7: Distribution in the invariant mass (left plot) and in the transverse momentum (right plot) of the lepton pair for MiNNLO$_{\rm PS}$ (blue, solid line), MiNLO$'$ (black, dotted line) and NNLO (red, dashed line).

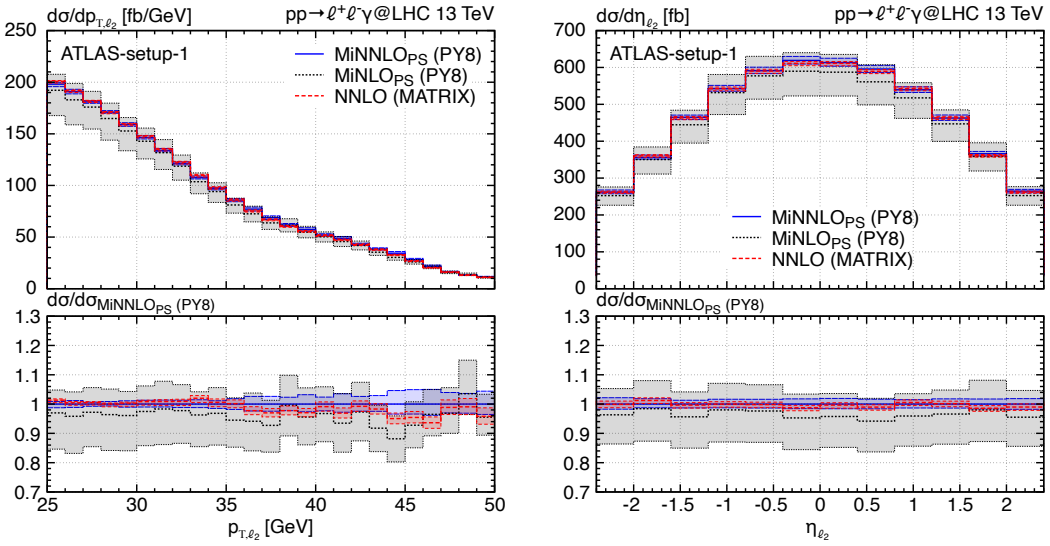

Figure 4.8: Distribution in the transverse momentum (left plot) and in the pseudorapidity (right plot) of the second-hardest lepton for MiNNLO$_{\text{PS}}$ (blue, solid line), MiNLO$'$ (black, dotted line) and NNLO (red, dashed line).

$m_{\ell\ell} \sim 70\,\text{GeV}$, caused by the Z-boson resonance in $m_{\ell\ell\gamma}$ [7]. Close to the latter resonance, we observe that the MiNNLO$_{\text{PS}}$ prediction is $\sim 4\%$ below the NNLO prediction, without overlapping uncertainty bands. The same effect has been observed in correspondence of the enhanced region of the distribution in the $m_{\ell\ell\gamma}$ invariant mass (not reported here). We have checked that in this region of phase space the $Z\gamma$ system tends to be softer. Therefore, all-order terms related to $p_{\text{T},\ell\ell\gamma}$ have a larger impact. Moreover, in the $p_{\text{T},\ell\ell}$ distribution we observe another interesting behaviour, this time related to the NNLO prediction. The NNLO result develops a perturbative instability (Sudakov shoulder) [271] around $p_{\text{T},\ell\ell} \sim 15\,\text{GeV}$ caused by an incomplete cancellation of virtual and real contributions from soft-gluon emissions, which is logarithmically divergent, but integrable. The reason is that the fiducial cut $p_{\text{T},\gamma} > 15\,\text{GeV}$ (see Table 4.1) for LO kinematics implies $p_{\text{T},\ell\ell} = p_{\text{T},\gamma} > 15\,\text{GeV}$, so that the $p_{\text{T},\ell\ell}$ distribution is not filled below 15 GeV at LO. Thus, the fixed-order result is NNLO accurate only for $p_{\text{T},\ell\ell} > 15\,\text{GeV}$, while for $p_{\text{T},\ell\ell} < 15\,\text{GeV}$ at least one QCD emission is necessary, which is described only at NLO accuracy. At the same time, the prediction becomes sensitive to soft-gluon effects at threshold, resulting in an instability at fixed order. Indeed, the parton shower cures this behaviour and yields a physical prediction at threshold for both MiNLO$'$ and MiNNLO$_{\text{PS}}$. This is one example where a NNLO calculation is insufficient and

[7] The peak at $m_{\ell\ell} \sim 70\,\text{GeV}$ is caused by the resonance in the $\ell^+\ell^-\gamma$ invariant mass, where the photon emission from the final state leptons is responsible for the downward shift of the peak (see Figure 4.1(b)). As commented for instance in Ref. [270], this peak is highly dependent on the fiducial cut on the photon transverse momentum: by gradually decreasing this cut, the smaller peak in the $m_{\ell\ell}$ distribution moves towards the larger peak to the point where the two eventually merge.

NNLO+PS matching is required.

Finally, in Figure 4.8 we consider distributions in the second-hardest lepton, showing its transverse momentum (p_{T,ℓ_2}) on the left and its rapidity (η_{ℓ_2}) in the right plot. Similar conclusions as made before for $m_{\ell\ell}$ and $\eta_{\ell\ell\gamma}$ apply also for these observables, so no further comments are needed. We reiterate however that, while the central predictions of MiNLO′ and MiNNLO$_{\mathrm{PS}}$ are generally close to each other, since MiNLO′ already includes many terms beyond NLO accuracy for $Z\gamma$ production, scale uncertainties are substantially reduced in case of MiNNLO$_{\mathrm{PS}}$, down to the level of the NNLO ones.

4.1.6 $Z\gamma$ transverse-momentum spectrum against NNLO+N^{3}LL

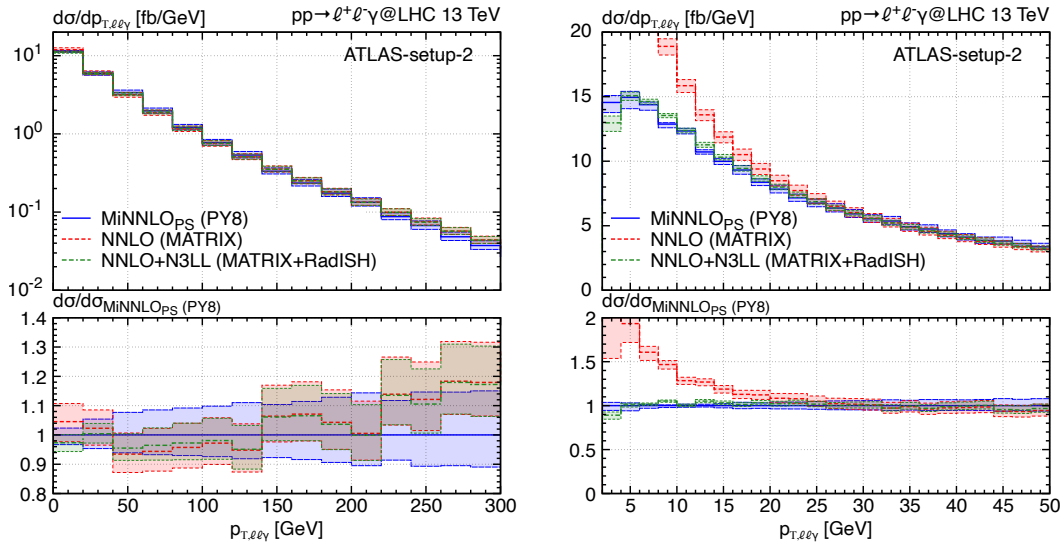

Figure 4.9: Distribution in the transverse momentum of the $Z\gamma$ system in a wider range (left plot) and at small $p_{\mathrm{T},\ell\ell\gamma}$ (right plot) for MiNNLO$_{\mathrm{PS}}$ (blue, solid line), NNLO (red, dashed line) and NNLO+N^{3}LL (green, double-dash-dotted line).

We continue our discussion of differential distributions with the transverse-momentum spectrum of the $Z\gamma$ system ($p_{\mathrm{T},\ell\ell\gamma}$). In Figure 4.9 we compare the $p_{\mathrm{T},\ell\ell\gamma}$ distribution in `ATLAS-setup-2` obtained with MiNNLO$_{\mathrm{PS}}$ against a more accurate prediction at NNLO+N^{3}LL (green, double-dash-dotted line), using the analytic resummation of large logarithmic contributions within Matrix+RadISH [272, 273]. The latter framework allows to match NNLO results with high-accuracy resummation through the RadISH formalism [105, 274, 275] for different transverse observables. For comparison we also show the NNLO result, which is effectively only NLO accurate for this distribution. One should bear in mind that those two predictions include the full heavy-quark mass dependence, which has an impact at large values of $p_{\mathrm{T},\ell\ell\gamma}$ (see discussion in Appendix A.3).

The NNLO+N^{3}LL prediction uses [273]

$$\mu_{\mathrm{R},0} = \mu_{\mathrm{F},0} = \sqrt{m_{\ell\ell}{}^2 + p_{\mathrm{T},\gamma}{}^2} \quad \text{and} \quad Q_{\mathrm{res},0} = \frac{1}{2} m_{\ell\ell\gamma} \tag{4.15}$$

as central scales, where $Q_{\mathrm{res},0}$ is the central resummation scale. $Q_{\mathrm{res},0}$ is varied by a factor of two up and down, while taking the envelope together with the 7-point μ_{R} and μ_{F} variation for the total scale uncertainty.

The $p_{\mathrm{T},\ell\ell\gamma}$ spectrum is shown in two different ranges in Figure 4.9. From the wider range in the left plot we notice that despite the different scale settings in the three calculations their predictions are in reasonable agreement at large $p_{\mathrm{T},\ell\ell\gamma}$ values. The fact that for $p_{\mathrm{T},\ell\ell\gamma} \gtrsim 150\,\mathrm{GeV}$ the NNLO and NNLO+N^{3}LL predictions become successively harder than the MiNNLO$_{\mathrm{PS}}$ one is related to the heavy-quark mass effects (see Appendix A.3). At large $p_{\mathrm{T},\ell\ell\gamma}$ all predictions are effectively NLO accurate, which is indicated by the enlarged scale-uncertainty bands. At small $p_{\mathrm{T},\ell\ell\gamma}$ the fixed-order result becomes unphysical, as the distribution is logarithmically divergent in the $p_{\mathrm{T}} \to 0$ limit, which is visible already in the left plot of Figure 4.9, but can be better appreciated in the zoomed version on the right. In this region, only the calculations that properly account for the resummation of soft QCD radiation by means of an analytic procedure (NNLO+N^{3}LL) or through parton-shower simulations (MiNNLO$_{\mathrm{PS}}$) provide a meaningful description. Even though at small $p_{\mathrm{T},\ell\ell\gamma}$ the Matrix+RadISH computation is N^{3}LL accurate, while the parton shower has a lower logarithmic accuracy, MiNNLO$_{\mathrm{PS}}$ and NNLO+N^{3}LL predictions are in excellent agreement down to transverse-momentum values (almost) in the non-perturbative regime.

4.1.7 Comparison of differential distributions against ATLAS data

Finally, we employ our MiNNLO$_{\mathrm{PS}}$ generator to compare NNLO+PS accurate predictions directly to ATLAS results from the most recent 13 TeV measurement of Ref. [238], which relies on the full 139 fb^{-1} Run-2 data. The comparison, carried out in `ATLAS-setup-2`, is presented in Figure 4.10. The experimental data are given as green points with error bars that refer to the experimental uncertainty. Six observables are shown: the transverse momentum ($p_{\mathrm{T},\gamma}$) and the pseudorapidity (η_γ) of the photon, the transverse momentum ($p_{\mathrm{T},\ell\ell\gamma}$) and the invariant mass ($m_{\ell\ell\gamma}$) of the $Z\gamma$ system, together with their ratio $p_{\mathrm{T},\ell\ell\gamma}/m_{\ell\ell\gamma}$ and the difference in the azimuthal angle between the lepton pair and the photon ($\Delta\phi_{\ell\ell,\gamma}$).

To assess effects from hadronization, in addition to our partonic MiNNLO$_{\mathrm{PS}}$ result (blue, solid) we also show a curve where the hadronization of the partonic events is modelled through Pythia8 (magenta, dash-dotted). The hadronic final states are kept stable and multi-parton interactions are turned off to avoid secondary photons and maintain a sufficiently simple analysis. By and large, we find minor contributions from hadronization for the observables considered here. They are at the level of the statistical uncertainties and well within scale variations.

Overall, we observe a remarkably good agreement with data both in the predicted shapes of the distributions and in the normalization, especially given the fact that the

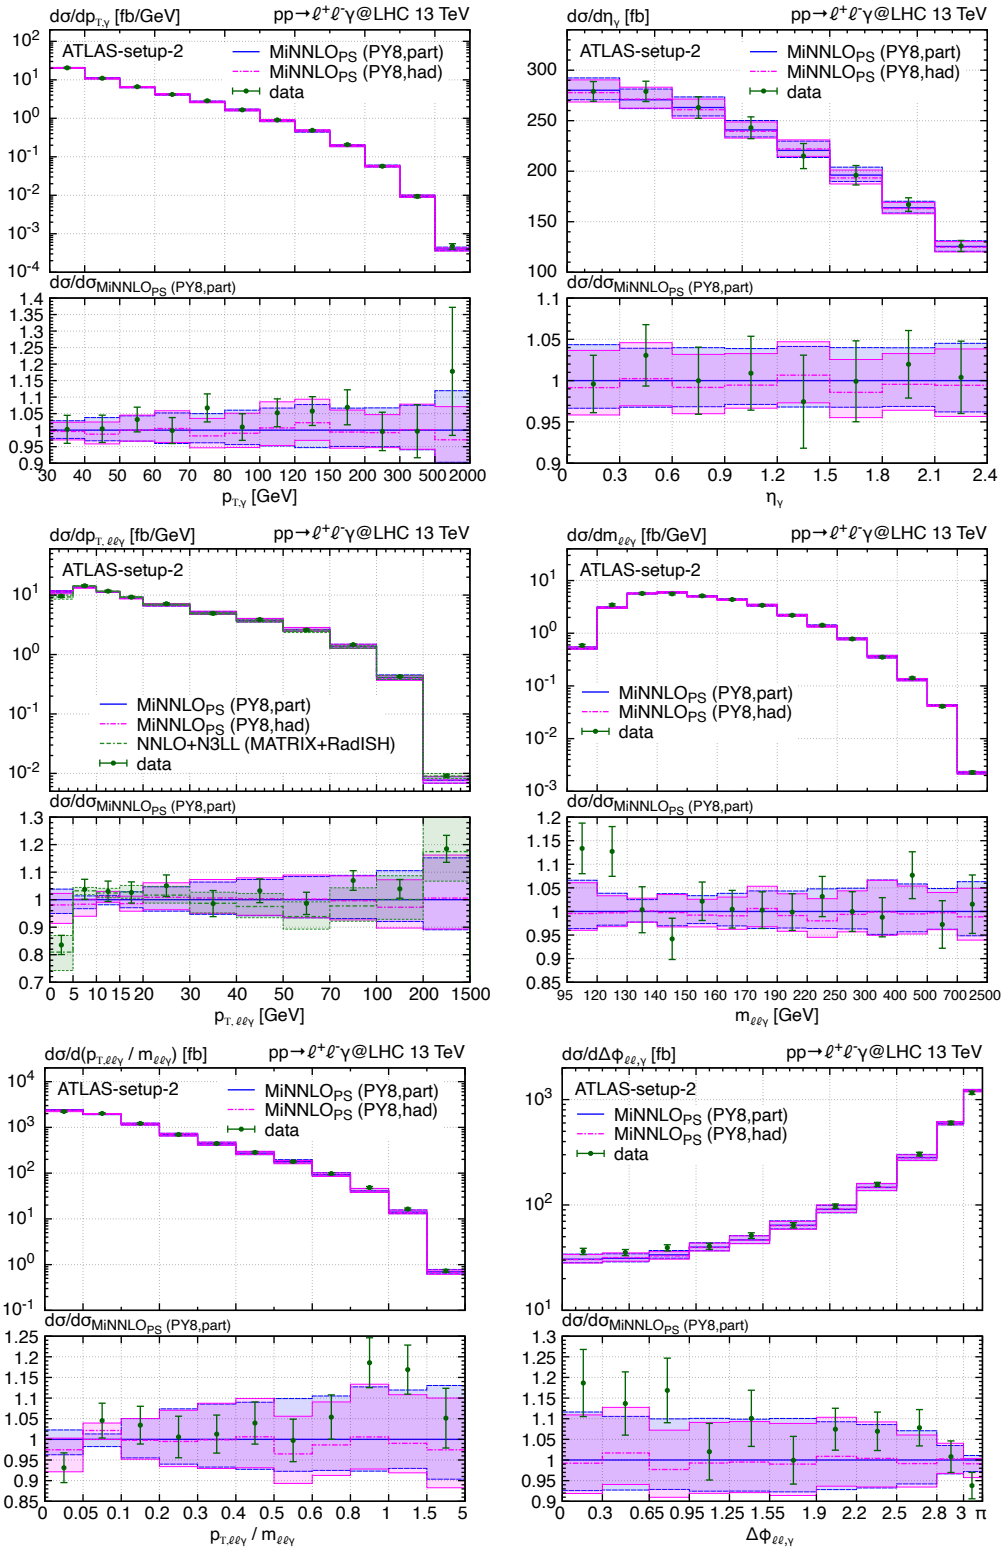

Figure 4.10: MiNNLOPS predictions at parton level (blue, solid line) and including hadronization (magenta, dash-dotted) compared to ATLAS 13 TeV data (green points with error bars). For $p_{\mathrm{T},\ell\ell\gamma}$ also NNLO+N^{3}LL (green, double-dash-dotted line) is shown.

theoretical and the experimental uncertainties are at the few-percent level only. All data points agree with our predictions within the experimental error bars, with the exception of only very few bins, where the agreement is reached within twice the experimental error. This is a clear improvement over the NLO-accurate event simulations employed in the data-theory comparison in figure 6 of Ref. [238], both in terms of accuracy (i.e. to describe the data) and in terms of precision (i.e. regarding theoretical uncertainties). Moreover, looking at the comparison of NNLO predictions to data in figure 7 of Ref. [238], it is clear that some (more inclusive) observables are equally well described at fixed order, while for observables sensitive to QCD radiation, such as $p_{\mathrm{T},\ell\ell\gamma}$ and $\Delta\phi_{\ell\ell,\gamma}$, NNLO predictions are not sufficient, and the matching to a parton shower is essential. In conclusion, our MiNNLO$_{\mathrm{PS}}$ calculation combines the two most important aspects (NNLO and parton-shower effects) to provide the most accurate and most precise $Z\gamma$ predictions to date, which will be essential to find potential deviations from the SM for this process in future.

It is worth discuss in more detail the $p_{\mathrm{T},\ell\ell\gamma}$ distribution in Figure 4.10. In this plot we have also added the more accurate NNLO+N^{3}LL prediction, as introduced in Section 4.1.6 with the scale setting of Eq. (4.15) and including heavy-quark mass effects. We recall that the latter, studied in Appendix A.3, induce differences at large $p_{\mathrm{T},\ell\ell\gamma}$ which are visible in the last bin. Despite the good agreement of MiNNLO$_{\mathrm{PS}}$ with data, the analytically resummed result is performing even better, especially in the first few bins, where the higher accuracy in the resummation of large logarithmic contributions is important. Although MiNNLO$_{\mathrm{PS}}$ and NNLO+N^{3}LL agree quite well (cf. the discussion in Section 4.1.6), this shows that for an observable like $p_{\mathrm{T},\ell\ell\gamma}$ it can be very useful to resort to tools that predict a single distribution more accurately, if available. Nevertheless, it is reassuring that our accurate multi-purpose MiNNLO$_{\mathrm{PS}}$ simulation, with all its flexibility to predict essentially any IR-safe observable, provides a very good description of such distributions as well.

We further notice that the deviation at small $m_{\ell\ell\gamma}$ is due to missing QED corrections as shown in Ref. [238]. Indeed, it is a well-known effect that the emission of additional QED radiation by final state leptons subtracts some energy, which causes the peak of invariant mass distributions involving leptons to shift to lower values. Our MiNNLO$_{\mathrm{PS}}$ computation renders the inclusion of such effects into NNLO-accurate predictions feasible by using a QED shower as implemented for instance within Pythia8, which could be very useful in an experimental analysis. Despite its feasibility, we decided not to focus on QED effects at all in this work. Further considerations on missing QED and EW effects on top of NNLO QCD corrections are presented in Section 4.4.

4.2 Some BSM studies in the $Z(\to \nu\bar{\nu})\gamma$ channel

As we have already discussed in the previous section quite at length, the production of a Z boson in association with a photon is an important vector-boson pair production process in various respects. In the current section we want to put special emphasis on its importance in search for new physics by considering the same process when the Z boson decay into a pair of neutrinos.

Different reasons make this process interesting in the context of searches for new

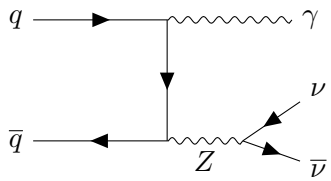

(a) $q\bar{q}$-initiated t-channel diagram

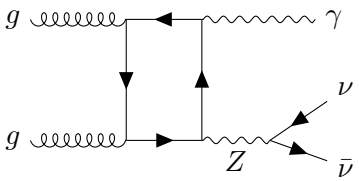

(b) gg-initiated loop-induced diagram

Figure 4.11: Sample Feynman diagrams for $\nu\bar{\nu}\gamma$ production entering at (a) LO and (b) NNLO.

physics. First of all, the measurement of non-zero $ZZ\gamma$ or $Z\gamma\gamma$ couplings, which are absent in the SM, would be direct evidence of BSM physics. Moreover, $Z\gamma$ final states are relevant in direct searches for BSM particles. Additionally, in the $Z \to \nu\bar{\nu}$ decay channel $Z\gamma$ production constitutes an irreducible background to dark-matter searches in the photon plus missing energy final state. The low accuracy of available $\nu\bar{\nu}\gamma$ event simulations is actually one of the limiting factors in current dark-matter analyses [276], despite the fact that a substantial effort has been made to improve the perturbative accuracy of $Z\gamma$ production.

Due to the relevance that improved theoretical predictions can have in supporting experimental analyses, we present here results for $Z\gamma$ production in the $Z \to \nu\bar{\nu}$ decay channel, where NNLO+PS accuracy is achieved with MiNNLO$_{\rm PS}$. To this end, the $Z\gamma$ MiNNLO$_{\rm PS}$ generator for $\ell^+\ell^-\gamma$ production, that we presented in the previous section, was extended in Ref. [218] to deal with the $\nu\bar{\nu}\gamma$ final state and to include in the event generation the effects of anomalous triple gauge couplings (aTGCs), specifically the $ZZ\gamma$ and $Z\gamma\gamma$ verteces.

We start our discussion by considering the process

$$pp \to \nu\bar{\nu}\gamma + X\,, \tag{4.16}$$

with $\nu \in \{\nu_e, \nu_\mu, \nu_\tau\}$ and X any additional QCD radiation. In our calculation, we account for all relevant topologies leading to this final state, and we include interferences, off-shell effects and spin correlations. At leading order the process is quark–anti-quark ($q\bar{q}$) induced in the SM and proceeds only via a t-channel quark exchange with both the isolated photon and the Z boson coupling to the quark line. This is different from the $\ell^+\ell^-\gamma$ final state, where Drell-Yan-like s-channel topologies are allowed as well, since the charged leptons can emit an isolated photon, while the neutrinos can not. A representative LO Feynman diagram is shown in Figure 4.11(a). At NNLO in QCD perturbation theory the loop-induced gluon-fusion contribution enters the cross section, see Figure 4.11(b). However, as already discussed in the previous section, this contribution is very small – at the (sub-)percent level – and, therefore, we neglect its effect in our results.

All ingredients necessary to reach NNLO+PS predictions for the process in Eq. (4.16) within MiNNLO$_{\rm PS}$ have been obtained as described in details in the previous section, relying on MCFM [242] or through OpenLoops [245–247] for the tree-level and one-loop

(a) $Z\gamma Z^*$ vertex (b) $Z\gamma\gamma^*$ vertex

Figure 4.12: Anomalous couplings between three gauge bosons that are relevant for $Z\gamma$ production.

amplitudes, while the two-loop contributions have been extracted from the MATRIX framework [81, 82] that is based on the calculation of Ref. [261]. This calculation allows us to retain NNLO QCD accuracy in the event generation for $\nu\bar{\nu}\gamma$ production interfaced to a parton shower, which is necessary for a complete and realistic event simulation. In particular, that offers the possibility to include multiple photon emissions through a QED shower, as well as non-perturbative QCD effects using hadronization and underlying event models. It is well known that these corrections can have a substantial impact on the lepton momenta, jet-binned cross sections and other more exclusive observables measured at the LHC.

In addition to the SM simulation, we have implemented the leading contributions from aTGCs in the $Z\gamma$ MINNLO$_{\text{PS}}$ generator for both the $\ell^+\ell^-\gamma$ and the $\nu\bar{\nu}\gamma$ final states, but we will focus on the latter when presenting phenomenological results below. Since couplings between three charge-neutral weak bosons are forbidden by the SM gauge symmetry, their contributions can arise only from BSM theories. This is one of the reasons why they provide a powerful way of searching for new physics. Extensions of the standard model gauge structure through aTGC can be described by means of two equivalent approaches [249, 277–279]: the vertex-function and the Lagrangian approach. Both descriptions can be embedded in a self-consistent effective-field-theory (EFT) framework, as presented in Refs. [280, 281]. Here, we follow the vertex-function approach, as this is usually employed by the experimental analyses (see for instance Ref. [282]).

For $Z\gamma$ production, two different neutral aTGCs enter the cross section, namely $Z\gamma V$ with $V = Z^*, \gamma^*$, which are shown in Fig 4.12. The form of these effective interactions can be constraint by imposing Lorentz and electromagnetic gauge invariance, as well as Bose statistics. The latter for instance forbids vertices such as $Z\gamma Z^*$ or $Z\gamma\gamma^*$ with all gauge bosons being on-shell, since in either case two identical particles are involved in the interaction. Note that, since the virtual photon and Z boson essentially couple to massless fermions, one has $\partial_\mu V^\mu = 0$ with V^μ being either the photon or the Z boson vector field. Moreover, we include only terms of dimension less than or equal to eight for practical reasons, as discussed in Ref. [249]. This choice, which avoids the proliferation of couplings that are in principle allowed by symmetry, is justified from an EFT perspective, where only a limited amount of higher-dimensional operators is expected to contribute

to the physical process at a given energy scale. With these minimal requirements, the effective interaction can be parametrized as [278] [8]:

$$\Gamma^{\alpha\beta\mu}_{Z\gamma V}(q_1, q_2, p) = \frac{i(p^2 - m_V^2)}{\Lambda^2}\Bigg(h_1^V (q_2^\mu g^{\alpha\beta} - q_2^\alpha g^{\mu\beta}) + \frac{h_2^V}{\Lambda^2} p^\alpha \Big(p \cdot q_2 \; g^{\mu\beta} - q_2^\mu p^\beta \Big) - h_3^V \varepsilon^{\mu\alpha\beta\nu} q_{2\,\nu} - \frac{h_4^V}{\Lambda^2} \varepsilon^{\mu\beta\nu\sigma} p^\alpha p_\nu q_{2\,\sigma} \Bigg), \quad (4.17)$$

where q_1 and q_2 are the momenta of the on-shell Z and γ gauge bosons, respectively, and p is the momentum of the off-shell boson V. One should bear in mind that in principle additional terms arise when all gauge bosons are considered to be off shell [284]. However, since the $Z\gamma$ analyses select isolated photons and measure predominantly contributions from Z bosons close to their mass shell, Eq. (4.17) provides the dominant effects also when including the leptonic decay of the final-state Z boson. The two anomalous vertices $Z\gamma Z^*$ and $Z\gamma\gamma^*$ are obtained by choosing the $V = Z^*, \gamma^*$ and setting m_V^2 accordingly. We remark here that the overall factor $(p^2 - m_V^2)$ causes the interaction to vanish for $p^2 = m_V^2$, which is required by Bose symmetry for $V = Z$ and by electromagnetic gauge invariance for $V = \gamma$. The effective couplings parametrizing the interaction are given by h_i^V with $i \in \{1 \dots 4\}$ in Eq. (4.17) and Λ is a mass scale conventionally chosen to be the Z boson mass m_Z. Note that a different scale choice for Λ just amounts to a rescaling of all h_i^V couplings [277].

The h_1^V and h_2^V anomalous couplings are CP-violating, and would only appear in a UV completion of the SM allowing for new particles with CP-violating interactions with the SM ones. Being CP odd, these terms can not interfere with the SM sector and they can just contribute to the cross section at quadratic level. On the contrary, the CP-preserving couplings h_3^V and h_4^V enter the cross section also with linear terms through interference with the SM amplitudes. In principle one may think that this fact renders the CP-violating couplings more difficult to constrain [279]; however, the linear term involves an interference between the t-channel and s-channel diagrams (of the SM and BSM contribution, respectively), which is strongly suppressed. Note also that the experimental sensitivity is affected by the dimensionality of the coupling itself. In particular, $h_{2/4}^V$ induce dimension eight terms, which grow with two extra powers of the energy scale with respect to the dimension six couplings $h_{1/3}^V$. Thus better limits can be obtained for $h_{2/4}^V$ [279]. There are many explicit new-physics models that introduce such aTGCs (see for instance Ref. [285] for a more detailed discussion). Indeed, any new fermionic particle can generate h_3^V at one-loop through a triangle diagram, while h_4^V arises only at a higher loop level or from non-perturbative effects as in certain technicolour models.

Our implementation of aTGCs within the $Z\gamma$ MiNNLO$_{\text{PS}}$ generator follows closely the one in Ref. [251]. All relevant diagrams involve $q\bar{q}$-initiated topologies where an (off-shell) Z boson or a photon is produced in the s-channel and splits into a Z boson and a photon through the anomalous vertices in Figure 4.12, with a subsequent decay of

[8] Very recently Ref. [283] pointed out that the CP-even part of the effective interaction in Eq. (4.17) is not consistent with an EFT formulation incorporating the full spontaneously-broken EW gauge invariance, and that a new anomalous coupling, dubbed h_5^V, should be included.

the Z boson (into charged leptons or neutrinos). We stress again that using Eq. (4.17) and considering the (off-shell) decay of the Z boson assumes that the experiments mostly measure Z bosons close to their mass shell, which is indeed a reasonable assumption. The relevant tree-level and one-loop amplitudes have been taken from MCFM [242], while we extended the calculation of the $q\bar{q} \to \ell^+\ell^-\gamma$ and $q\bar{q} \to \nu\bar{\nu}\gamma$ two-loop helicity amplitudes of Ref. [261] with the relevant anomalous contributions directly within MATRIX [245–247], using the $q\bar{q}V^*$ form factor [286–288] for the loop corrections to the tree-level amplitudes with aTGCs.

4.2.1 Results within the SM and beyond

We now turn to presenting phenomenological results for $pp \to \nu\bar{\nu}\gamma$ production at the LHC with $\sqrt{s} = 13\,\text{TeV}$ centre-of-mass energy for $\nu \in \{\nu_e, \nu_\mu, \nu_\tau\}$. All results have been obtained with $n_f = 5$ massless quark flavours and the corresponding NNLO set of the NNPDF3.0 [264] parton distribution functions with a strong coupling $\alpha_s(m_\text{Z}) = 0.118$. The electroweak parameters are evaluated in the G_μ scheme with the electroweak coupling $\alpha_{G_\mu} = \sqrt{2}G_\mu m_\text{W}^2 \sin^2\theta_W/\pi$ and the mixing angle $\cos^2\theta_W = m_\text{W}^2/m_\text{Z}^2$. The input parameters are set to $G_\mu = 1.16637 \times 10^{-5}\,\text{GeV}^{-2}$, $m_\text{Z} = 91.1876\,\text{GeV}$, $m_\text{W} = 80.385\,\text{GeV}$, $\Gamma_Z = 2.4952\,\text{GeV}$, and $\Gamma_W = 2.085\,\text{GeV}$. We note that predictions for MINLO′ are obtained by setting to zero the $\mathcal{D}_{\ell_{Z_\gamma}}$ terms in Eq. (3.76). The scale setting for MINNLO$_\text{PS}$ (MINLO′) is fixed by the method itself and it is the same that we presented in Section 4.1.3. In the fixed-order results we set the central renormalization and factorization scales to the transverse mass of the $Z\gamma$ system. In all cases we use 7-point scale variations to estimate the uncertainties related to missing higher-order contributions.

For the results discussed in this section, we consider three sets of fiducial cuts, which are summarized in Table 4.3. The first one (`fiducial-setup-1`) corresponds to an earlier ATLAS analysis [232] and is used for validation purposes. To study the effects of aTGCs we use `fiducial-setup-2`, which is also employed to compare MINNLO$_\text{PS}$ predictions to a recent $\nu\bar{\nu}\gamma$ measurement by ATLAS [282]. The last setup (`DM-setup`) has instead been chosen to study the importance of NNLO+PS predictions for reducing the uncertainties of the $\nu\bar{\nu}\gamma$ background in dark-matter searches in the photon plus missing energy channel and it is inspired by a recent dark-matter search [276]. All three setups include standard cuts on the identified photon and the missing transverse energy, a jet definition and Frixione smooth-cone isolation [255] for the photon (for the notation, see Section 4.1.1). In `fiducial-setup-2` one category inclusive over QCD radiation and one with a jet veto is considered. The `DM-setup`, on the other hand, considers a quite special choice for the smooth-cone parameters as it combines a fixed (lower) threshold with a fraction of the photon transverse momentum. Note also that in the `DM-setup` various inclusive and exclusive categories in $p_{\text{T,miss}}$ are considered, as discussed below, where the $p_{\text{T,miss}}$ cut given in Table 4.3 is just the loosest one.

Validation of the event generator

In Figure 4.21 we start by comparing MINNLO$_\text{PS}$ (blue, solid line), MINLO′ (black, dotted line), and fixed-order NNLO predictions (red, dashed line). For reference we also show matched NLO (brown, dash-double-dotted line) and fixed-order NLO (green,

	`fiducial-setup-1`	`fiducial-setup-2`	`DM-setup`
Photon cuts	$p_{\mathrm{T},\gamma} > 100\,\mathrm{GeV}$ $\lvert\eta_\gamma\rvert < 2.37$	$p_{\mathrm{T},\gamma} > 150\,\mathrm{GeV}$ $\lvert\eta_\gamma\rvert < 2.37$	$p_{\mathrm{T},\gamma} > 150\,\mathrm{GeV}$ $\lvert\eta_\gamma\rvert < 1.37$ or $1.52 < \lvert\eta_\gamma\rvert < 2.37$ $\Delta\phi_{\gamma,\vec{p}_{\mathrm{T,miss}}} > 0.4$
Neutrino cuts	$p_{\mathrm{T,miss}} > 90\,\mathrm{GeV}$	$p_{\mathrm{T,miss}} > 150\,\mathrm{GeV}$	$p_{\mathrm{T,miss}} > 200\,\mathrm{GeV}$
Jet cuts	—	Inclusive: $N_{\mathrm{jet}} \geq 0$ Exclusive: $N_{\mathrm{jet}} = 0$	$N_{\mathrm{jet}} \leq 1$
Jet definition	anti-k_{T} with $R = 0.4$ $p_{\mathrm{T},j} > 30\,\mathrm{GeV}$ $\lvert\eta_j\rvert < 4.4$ $\Delta R_{\gamma j} > 0.3$	anti-k_{T} with $R = 0.4$ $p_{\mathrm{T},j} > 50\,\mathrm{GeV}$ $\lvert\eta_j\rvert < 4.5$ $\Delta R_{\gamma j} > 0.3$	anti-k_{T} with $R = 0.4$ $p_{\mathrm{T},j} > 30\,\mathrm{GeV}$ $\lvert\eta_j\rvert < 4.5$ $\Delta\phi_{j,\vec{p}_{\mathrm{T,miss}}} > 0.4$
Photon Isolation	Frixione isolation $n = 1$ $\epsilon_\gamma = 0.5$ $\delta_0 = 0.4$	Frixione isolation $n = 1$ $\epsilon_\gamma = 0.1$ $\delta_0 = 0.1$	Frixione isolation $n = 1$ $E_{\mathrm{T}}^{\mathrm{ref}} = 2.45\,\mathrm{GeV} + 0.022\,p_{\mathrm{T},\gamma}$ $\delta_0 = 0.4$

Table 4.3: Definition of fiducial cuts of two ATLAS measurements, `fiducial-setup-1` [232] and `fiducial-setup-2` [282], and of the `DM-setup` that is inspired by the dark-matter search of Ref. [276].

double-dash-dotted line) predictions in the ratio panel. This comparison is done at the LHE level and it serves the purpose of numerically validating the NNLO accuracy of the MiNNLO$_{\mathrm{PS}}$ predictions as well as indicating the importance of NNLO corrections and matching to the parton shower. The left plot shows the rapidity distribution of the photon. The agreement between MiNNLO$_{\mathrm{PS}}$ and NNLO predictions is excellent, with fully overlapping uncertainty bands. In line with the discussion of previous section, we remind that small differences (within uncertainties) between MiNNLO$_{\mathrm{PS}}$ and fixed-order results are expected due to the different scale settings and treatment of terms beyond accuracy. Moreover, when performing scale variations for MiNLO$'$ and MiNNLO$_{\mathrm{PS}}$, an additional scale dependence is kept in the Sudakov form factor [213] for a more conservative uncertainty estimate, which is absent in fixed-order calculations. This is reflected in the slightly larger MiNNLO$_{\mathrm{PS}}$ uncertainty band. Compared to MiNLO$'$, however, we observe a clear reduction of scale uncertainties from about 6% to 3% for MiNNLO$_{\mathrm{PS}}$ and roughly an effect of +3% in normalization from the inclusion of NNLO corrections through MiNNLO$_{\mathrm{PS}}$.

Compared to the NLO LHE and fixed-order results in the ratio panel, the effect of NNLO corrections is even at the 10% level, indicating that these lower-order calculations are

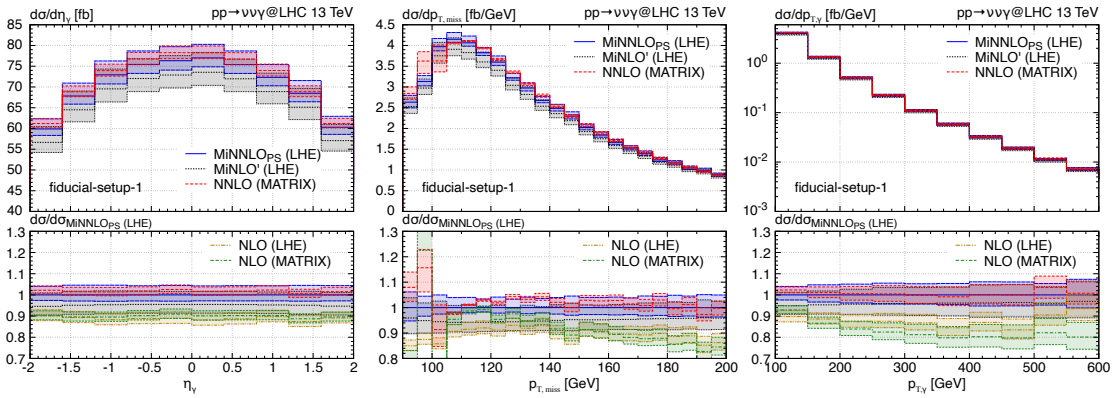

Figure 4.13: Selected plots at LHE level for validation of MiNNLO$_{\text{PS}}$ against NNLO. See text for details.

insufficient for an accurate description of this process. Moreover, the NLO uncertainty bands appear to be unnaturally small, not even overlapping with the MiNNLO$_{\text{PS}}$ prediction and much smaller than for the MiNLO$'$ result, which, in fact, correctly includes some terms beyond NLO accuracy.

Also, at high values of the missing energy ($p_{\text{T,miss}}$) and of the transverse momentum of the photon ($p_{\text{T},\gamma}$) Figure 4.21 shows that MiNNLO$_{\text{PS}}$ and fixed-order NNLO predictions are in excellent agreement. At small values of the missing transverse energy, on the other hand, we observe that the NLO and the NNLO curves develop an interesting feature, similarly to what we encountered in Figure 4.7 for the transverse momentum of the charged-lepton system. This is a consequence of the fiducial cut of $p_{\text{T},\gamma} > 100\,\text{GeV}$, which induces a perturbative instability [271] in $p_{\text{T,miss}}$ at the threshold, as the region $p_{\text{T,miss}} \leq 100\,\text{GeV}$ becomes sensitive to soft-gluon effects and is effectively filled only starting from NLO. This behaviour of fixed-order results is unphysical and cured in the matched predictions (already at LHE level), as can be seen for instance for MiNNLO$_{\text{PS}}$. We have further checked that the impact of including additional emissions through the parton shower is very moderate for the rather inclusive observables under consideration, with (positive) effects at the few-percent level.

Anomalous coupling effects

We continue by studying the effects of aTGCs on differential distributions. The search for aTGCs in $Z\gamma$ production has received great attention in the past, at LEP [289], at the Tevatron [191, 290], and at the LHC at 7 TeV and 8 TeV [227, 229–232, 234–236]. We stress that, the $Z \to \bar{\nu}\nu$ decay channel has a higher sensitivity to aTGCs due to a higher branching ratio compared to the $Z \to \ell^+\ell^-$ decay channel. Indeed, the most recent 13 TeV ATLAS analysis [282] uses the $\nu\bar{\nu}\gamma$ final state to set the most stringent limits on the aTGCs under consideration thus far, which are of the order of $\pm 10^{-4}$ for h_3^V and $\pm 10^{-7}$ for h_4^V (see Table 8 of Ref. [282] for the exact bounds). This ATLAS

analysis did not make use of any form-factor suppression, which is sometimes applied to prevent unitarity violation at high energy, caused by the introduction of aTGCs at the amplitude level [291]. Indeed, in an EFT perspective the terms entering the vertex function in Eq. (4.17) would arise from a set of gauge-invariant operators of dimension eight (or higher) [281], whose validity range is limited by the given new-physics scale. Here, we also refrain from using any artificial form-factor suppression.

Different combinations of the aTGCs have been obtained through reweighting at event-generation level (i.e stage 4 in POWHEG), while ensuring sufficient statistics in the relevant phase-space regions by accounting for the resonance structure of both the t-channel SM and s-channel BSM topologies and by applying suitable suppression factors to increase the sampling in the high-energy tails [9]. Even though all eight anomalous couplings (h_i^V with $i \in \{1...4\}$ and $V \in \{Z, \gamma\}$) are consistently implemented in our code, we limit our study to the CP-conserving ones, which do not interfere with the CP-violating ones. Moreover, since also the $V = Z$ and $V = \gamma$ couplings have been shown to only mildly interfere with each other [277] and to have qualitatively a very similar impact, we focus on the pair (h_3^Z, h_4^Z) here. The most relevant phase-space regions to constrain aTGCs in $\nu\bar{\nu}\gamma$ production are the high-energy tails of the $p_{\mathrm{T,miss}}$ and $p_{\mathrm{T},\gamma}$ spectra. In Figure 4.14 we show MINNLO$_{\mathrm{PS}}$ predictions for each of these two observables in `fiducial-setup-2` with a jet veto ($N_{\mathrm{jet}} = 0$), which is experimentally applied to reduce the SM background in the tails of the distributions. MINNLO$_{\mathrm{PS}}$ results include parton shower and hadronization effects as provided by PYTHIA8 [124], with the A14 tune [266]. We present individual variations of h_3^Z (two upper plots), individual variations of h_4^Z (two central plots), and combinations of (h_3^Z, h_4^Z) (two bottom plots), all within the currently allowed limits [282]. In all plots, the SM results ($h_3^Z = 0, h_4^Z = 0$) are shown with a blue, solid line. As one sees from the individual variations of h_3^Z and h_4^Z, both negative and positive values of the aTGCs lead to a similar positive effect on the spectra, which is a consequence of the very small interference of the s-channel BSM amplitudes with the t-channel SM amplitude, so that the quadratic term in the anomalous couplings dominates. This is also the reason why the experimental limits on the aTGCs in $Z\gamma$ production are almost symmetric. For values of h_3^Z at the edges of experimentally allowed ranges, we start observing deviations of 5-10% from the SM for transverse-momentum values of 400-500 GeV, with a steep increase afterwards, reaching already 100% around 700-800 GeV. For h_4^Z, whose constraints are at least three orders of magnitude more stringent, 5-10% effects in the tails of transverse-momentum distributions manifest themselves starting from 600-700 GeV, with rapidly increasing effects at larger p_{T} as well. Looking at the simultaneous variations of (h_3^Z, h_4^Z), it is clear that different sign combinations constructively interfere, with a mild difference between the two possible sign combinations. On the other hand, same sign combinations of (h_3^Z, h_4^Z) interfere destructively.

[9] To this end, we have added the aTGC coefficients as inputs to the POWHEG reweighting information, introduced the flag `anommode 1` that enables the s-channel resonance histories associated with the aTGCs to be included through the `build_resonance_histories` routines of POWHEG-BOX-RES, and implemented in the code suitable suppression factors for the $\bar{B}$ function and the POWHEG remnant contribution that can be activated respectively via `bornsuppfact 1` and `remnsuppfact 1`.

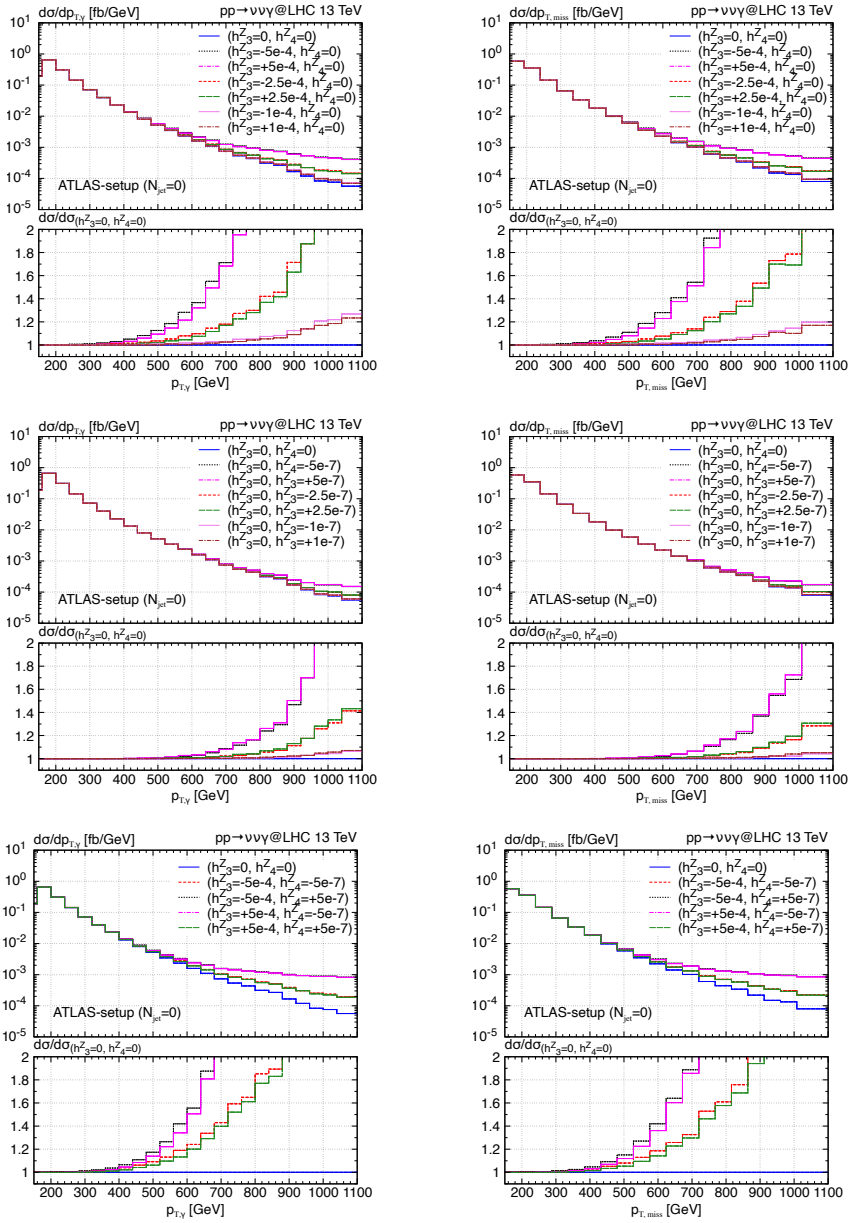

Figure 4.14: MiNNLO$_{\mathrm{PS}}$ predictions for the photon transverse momentum (left column) and the missing transverse momentum (right column) for different values of h_3^Z and h_4^Z. The SM results are reported with a blue, solid line. Individual variations of h_3^Z (two upper plots), individual variations of h_4^Z (two central plots) and their combined variations (two bottom plots) are considered for different values within the experimentally allowed ranges defined in Table 8 of Ref. [282]. All MiNNLO$_{\mathrm{PS}}$ results include parton shower and hadronization effects, as provided by Pythia8.

category	SRI1	SRI2	SRI3	SRI4
$p_{\mathrm{T,miss}}$ [GeV]	> 200	> 250	> 300	> 375
MiNLO′ [fb]	$27.35(18)^{+6.0\%}_{-3.5\%}$	$12.95(11)^{+6.5\%}_{-4.2\%}$	$6.65(8)^{+6.8\%}_{-4.7\%}$	$2.77(6)^{+8.0\%}_{-5.8\%}$
MiNNLO$_{\mathrm{PS}}$ [fb]	$29.09(18)^{+2.9\%}_{-1.9\%}$	$13.77(12)^{+3.2\%}_{-2.2\%}$	$7.07(8)^{+3.2\%}_{-2.4\%}$	$2.95(6)^{+4.2\%}_{-3.2\%}$

Table 4.4: Fiducial cross section in various inclusive $p_{\mathrm{T,miss}}$ categories in the `DM-setup`.

category	SRE1	SRE2	SRE3
$p_{\mathrm{T,miss}}$ [GeV]	200–250	250–300	300–375
MiNLO′ [fb]	$14.4(1.4)^{+5.4\%}_{-3.2\%}$	$6.30(8)^{+6.2\%}_{-3.7\%}$	$3.88(5)^{+5.8\%}_{-3.9\%}$
MiNNLO$_{\mathrm{PS}}$ [fb]	$15.32(15)^{+2.7\%}_{-1.6\%}$	$6.69(7)^{+3.2\%}_{-2.0\%}$	$4.12(5)^{+2.5\%}_{-1.8\%}$

Table 4.5: Fiducial cross section in various exclusive $p_{\mathrm{T,miss}}$ categories in the `DM-setup`.

Improving background in dark matter searches

Now we turn to discussing the importance of NNLO+PS predictions for $\nu\bar{\nu}\gamma$ productions in the context of dark-matter searches in the photon plus missing energy ($\gamma + E_{\mathrm{T,miss}}$) channel, which is one of (if not the) most important signature to detect dark matter at the LHC, see Refs. [276, 292, 293]. Other $X + E_{\mathrm{T,miss}}$ signatures (where X is a visible particle) have been extensively studied at the LHC in the past years: for a jet ([294, 295]), for a heavy quark ([296, 297]), for a vector boson ([295, 298, 299]), and for a Higgs boson ([300, 301]). In Ref. [276], which is the most recent dark-matter study in the $\gamma + E_{\mathrm{T,miss}}$ channel, the results are interpreted both in terms of simplified dark matter models [302–304] and of effective field theories of axion-like particles (ALPs) [305]. As one can see from Table 4 and 5 of Ref. [276], for instance, the dominant SM background is the $\nu\bar{\nu}\gamma$ process, which also dominates the uncertainties of the expected SM events. Depending on the category in $p_{\mathrm{T,miss}}$ considered in Ref. [276], which were used to improve the sensitivity of the analysis, the uncertainties on the expected $\nu\bar{\nu}\gamma$ events range from 4% to almost 15%. Ref. [276] based its $\nu\bar{\nu}\gamma$ predictions on a merged calculation of 0-jet and 1-jet events at NLO+PS within the `Sherpa 2.2` MC event generator [129, 306].

Here, we consider MiNLO′ predictions, which have the same formal accuracy as the merged `Sherpa 2.2` results quoted in Ref. [276], and study the reduction of scale uncertainties when including MiNNLO$_{\mathrm{PS}}$ corrections in each $p_{\mathrm{T,miss}}$ category. Among the categories in $p_{\mathrm{T,miss}}$ considered in Ref. [276], four of them are of inclusive type (SRI1-SRI4) and three of them of exclusive type (SRE1-SRE3). The different $p_{\mathrm{T,miss}}$ ranges defining each of the seven categories are summarized in Table 4.4 and 4.5. These tables also report MiNLO′ and MiNNLO$_{\mathrm{PS}}$ predictions for the cross sections in those categories with

the respective scale uncertainties. Also in this case, results have been showered using Pythia8, with the inclusion of hadronization effects. One should bear in mind that the experimental analysis is performed at the level of events measured in the detector, so that an immediate comparison to Table 4 and 5 of Ref. [276] is not possible. However, both the relative MiNNLO$_{\text{PS}}$ correction and the reduction of scale uncertainties from MiNLO$'$ to MiNNLO$_{\text{PS}}$ give a good indication of the expected improvements. Moreover, by and large, relative uncertainties at detector-event level and at the fiducial level can be assumed to be similar, so that even a direct comparison between our fiducial MiNNLO$_{\text{PS}}$ results and Ref. [276] is not meaningless. We find roughly a 6% correction in the central value for all categories by including NNLO corrections through MiNNLO$_{\text{PS}}$. Moreover, even though uncertainties are larger in categories with more stringent cuts, we still observe an overall reduction in the uncertainty bands by roughly a factor of two. Comparing our uncertainties to the ones of the $\nu\bar{\nu}\gamma$ backgrounds reported in Table 4 and 5 of Ref. [276], we find smaller uncertainties of MiNNLO$_{\text{PS}}$ already in the SRI1 category ($+2.9\%$ and -1.9% compared to $\pm 4.4\%$), while with increasing $p_{\text{T,miss}}$ cut the quoted uncertainties on the $\nu\bar{\nu}\gamma$ events in Ref. [276] increase significantly, up to $\pm 13\%$ in SRI4, while our MiNNLO$_{\text{PS}}$ uncertainties amount to less than 5%, bearing in mind that the translation of the uncertainties of the predictions at detector-event level to those at fiducial level is not immediate, as explained above. Similarly, in the exclusive $p_{\text{T,miss}}$ categories in Table 4.5 the MiNNLO$_{\text{PS}}$ uncertainties stay within about 3%, while the quoted uncertainties in Ref. [276] of the $\nu\bar{\nu}\gamma$ predictions range from 6% to 9%. In conclusion, MiNNLO$_{\text{PS}}$ predictions for $\nu\bar{\nu}\gamma$ production will allow the experiments to substantially improve the dominant background uncertainty in dark-matter searches in the photon plus missing energy channel.

Comparing against ATLAS data

Finally, we compare in Figure 4.15 our $\nu\bar{\nu}\gamma$ MiNNLO$_{\text{PS}}$ predictions, including the effects from hadronization, against 13 TeV ATLAS data with an integrated luminosity of 36.1 fb^{-1} [282]. The first plot shows the distribution in the number of jets. The agreement between the MiNNLO$_{\text{PS}}$ results and the data points for the different jet cross sections is reasonable, being within at most two standard deviations. Note that, starting from NNLO accuracy for the 0-jet cross section, the accuracy of the MiNNLO$_{\text{PS}}$ calculation decreases by one order for each jet multiplicity, with the >2-jet multiplicity described only by the shower. This is the reason why the prediction undershoots the data in this bin. The other plots in Figure 4.15 show the $p_{\text{T},\gamma}$ and $p_{\text{T,miss}}$ spectra with and without a jet veto ($N_{\text{jet}} = 0$). Also here the agreement between MiNNLO$_{\text{PS}}$ predictions and data is very good, with deviations of typically one or at most two standard deviations. There is however one exception: the measured result in the last bin in the inclusive ($N_{\text{jet}} \geq 0$) $p_{\text{T},\gamma}$ spectrum is many standard deviations away from the prediction. The data point seems to be way too low, when following the trend of the distribution. Indeed, it has a very large error and is actually compatible with zero. Moreover, looking at the $N_{\text{jet}} = 0$ result, this bin has actually a higher measured cross section than in the inclusive case, which appears inconsistent considering the fact that the $N_{\text{jet}} = 0$ cross section should be part

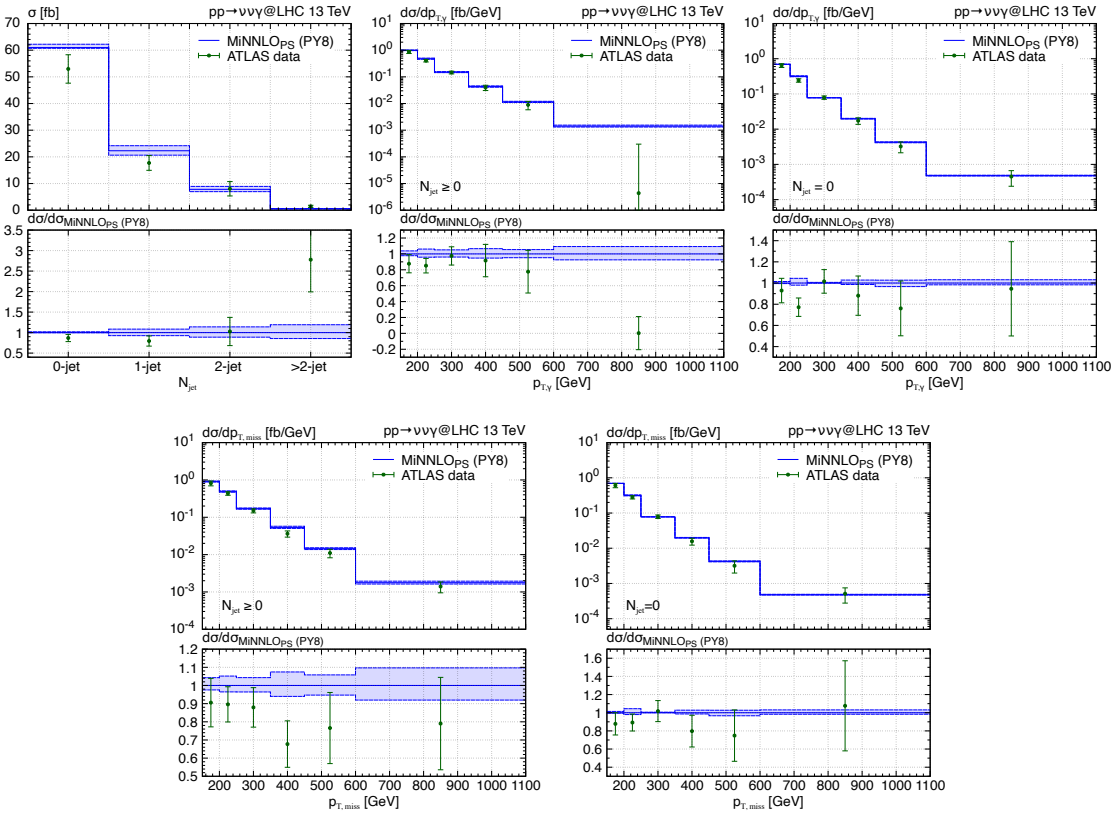

Figure 4.15: Distribution in the number of jets (N_{jet}), in the transverse momentum of the photon ($p_{T,\gamma}$) and in the missing transverse momentum ($p_{T,\text{miss}}$) compared to 13 TeV ATLAS data [282]. The latter two are shown both inclusively ($N_{\text{jet}} \geq 0$) and with a jet veto ($N_{\text{jet}} = 0$).

of the $N_{\text{jet}} \geq 0$ one. A possible explanation could be that, because of the additional jet activity, some events are discarded, for instance due to the photon isolation requirements. Indeed, looking at the $p_{\text{T,miss}}$ spectrum no such behaviour is observed.

4.3 A NNLO+PS W^+W^- generator

Having the largest cross section among massive diboson processes, W^+W^- production is one of the most important signature at the LHC. Moreover, unlike $Z\gamma$ production, it provides direct access to triple-gauge couplings, which appear already in the leading perturbative contribution to the cross section. Therefore, the measurement of this process at the LHC is a direct probe of the gauge symmetry structure of electroweak interactions and of the mechanism of EW symmetry breaking in the SM. Moreover, W^+W^- final states are an irreducible background to Higgs measurements in the $H \to W^+W^-$ decay channel and to direct searches for BSM particles decaying into two leptons, missing energy, and/or jets. On the other hand, an accurate theoretical modelling of this signal is even more important if one considers that its final state can not be fully reconstructed by experiments, due to the presence of neutrinos. Additionally, the theoretical description of the W^+W^- cross section is challenged by the correct modelling of the jet veto, which is applied by the experimental analyses to suppress backgrounds involving top-quarks ($t\bar{t}$ and tW) and essentially consists in requiring the transverse momentum of jets to be below a given threshold. A strict veto against jets in the final state increases the sensitivity to higher-order QCD effects due to potentially large logarithms of the ratio of the small jet-veto scale over the large invariant mass of the system. Such terms challenge the reliability of fixed-order predictions and induce large uncertainties in theory predictions that are typically not covered by scale-variation procedures, especially when extrapolating cross sections measured in the fiducial region to the total phase space. This makes clear that the inclusion of higher perturbative orders in the expansion of the W^+W^- cross section is not sufficient, and such calculations should be supplemented by all-order resummation.

With our NNLO+PS W^+W^- generator, introduced in Ref. [219], the most accurate FO predictions for this process can be interfaced to a transverse-momentum ordered PS. In a previous version of the W^+W^- generator [211], NNLO+PS accuracy was obtained by an a-posteriori reweighting of MiNLO′ results by the NNLO calculation, according to Eq. (3.63); in Ref. [219] this target accuracy is achieved via MiNNLO$_{\text{PS}}$, where NNLO corrections are included in a more efficient way, directly at event generation level without any approximation.

As a starting point, we consider the process

$$pp \to \ell^+ \nu_\ell \, \ell'^- \bar{\nu}_{\ell'} + X \,, \tag{4.18}$$

for any combination of massless leptons $\ell, \ell' \in \{e, \mu, \tau\}$ with different flavours $\ell \neq \ell'$. Again, X refers here to any additional QCD activity. For simplicity and without loss of generality we consider only the process $pp \to e^+\nu_e\mu^-\bar{\nu}_\mu + X$ here, which we will refer to as W^+W^- production in the following. By including all resonant and non-resonant topologies leading to this process, off-shell effects, interferences and spin correlations are taken into account. Sample LO diagrams are shown in Figure 4.16, including

(a) double-resonant t-channel W^+W^- production;

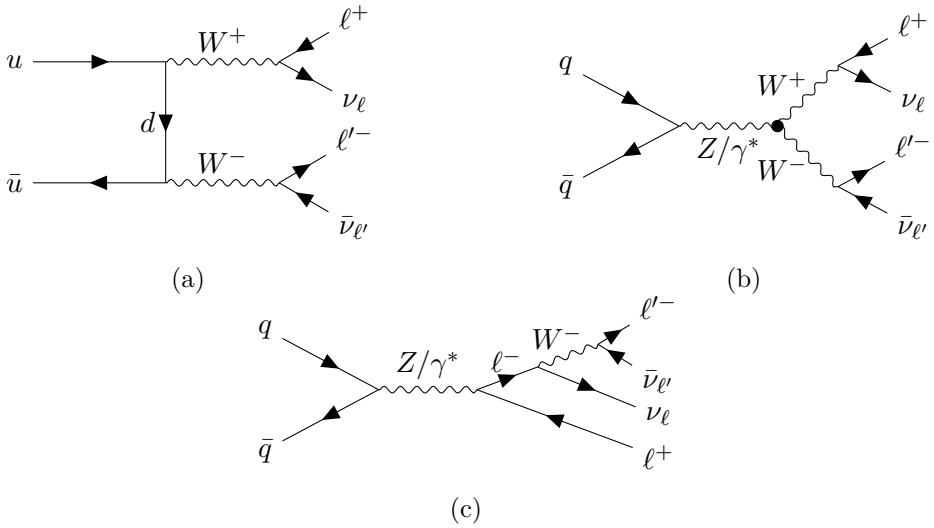

Figure 4.16: Sample LO diagrams in the different-flavour channel $\ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$ for (a) t-channel W^+W^- production, (b) s-channel $Z/\gamma^\star \to W^+W^-$ production, and (c) DY-type production.

(b) double-resonant s-channel $Z/\gamma^\star \to W^+W^-$ topologies via a triple-gauge coupling, with either the W^+W^- pair, or the Z boson and one W boson being resonant;

(c) double-resonant DY-type production, where both the Z boson and the W boson can become simultaneously resonant.

As usual, the appropriate identification of these topologies is also employed within the resonance-aware Powheg-Box-Res framework, to correctly identify [10] the relevant resonance histories according to the richer-resonance-history criterium (see footnote 1 of Section 3.1.1 and reference therein).

The corresponding production of opposite-charge same-flavour leptons $pp \to \ell^+\nu_\ell\ell^-\bar{\nu}_\ell + X$ involves the same type of W^+W^- diagrams as shown in Figure 4.16, but also additional ZZ diagrams as shown in Figure 4.28. By focusing on the different-flavour case ($\ell \neq \ell'$) we avoid the complications originating from the mixing of the W^+W^- and ZZ topologies. In fact, as shown in Refs. [307–309], W^+W^- and ZZ interference effects can be largely neglected and, to a very good approximation, predictions for the two processes can be added incoherently.

An important aspect of W^+W^- production is that its cross section is subject to a severe contamination from top-quark contributions. As anticipated at the beginning of this section, not only does this affect W^+W^- measurements at the LHC, which usually employ a jet veto, a b-jet veto, or both to suppress top-quark backgrounds, but it also renders the

[10] To this end, we had to adapt the Powheg-Box-Res code (specifically the `build_resonance_histories` routine) to automatically find all relevant resonance histories for W^+W^-+jet production. This was required, because the automatic generation of resonance histories is not fully functional for processes with a jet in the final state.

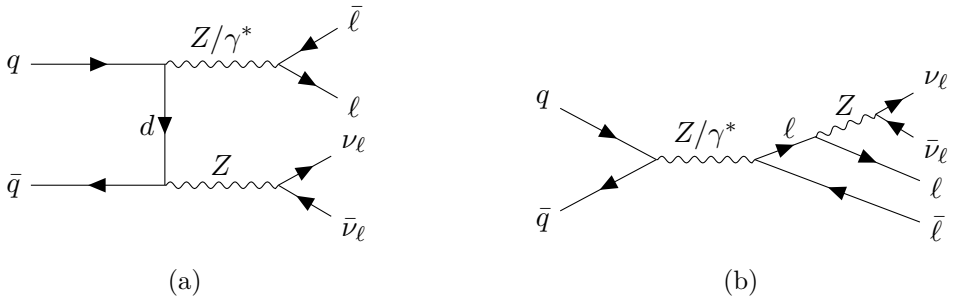

Figure 4.17: Sample additional LO diagrams appearing in the same-flavour channel $\ell^+\nu_\ell\ell^-\bar{\nu}_\ell$ for (a) t-channel ZZ production, and (b) DY-type production.

theoretical definition of the W^+W^- cross section cumbersome. Indeed, resonant top-quark contributions enter radiative corrections to W^+W^- production through interference with real-emission diagrams involving two bottom quarks in the final state. Those interference terms are numerically so large that they easily provide the dominant contribution to the cross section. Specifically, in the inclusive phase space genuine W^+W^- contributions are more than one order of magnitude smaller. Therefore, the consistent removal of the top-quark contamination is mandatory to define a top-free W^+W^- cross section. To this end, we exploit the four-flavour scheme (4FS), where bottom quarks are treated as being massive, do not enter in the initial state and diagrams with real bottom-quark radiation are separately finite. This allows us to drop all contributions with final-state bottom quarks, thereby cancelling the top-quark contamination and obtaining top-free W^+W^- results. We note that there exists an alternative approach to define a top-free W^+W^- cross section that can be used in the five-flavour scheme (5FS). However, this approach is much less practical as it requires the repeated evaluation of the cross section (and distributions) with increasingly small values of the top-quark width Γ_t to extract the top-free W^+W^- cross section as the contribution that is not enhanced by $1/\Gamma_t$. Indeed, it was shown in Ref. [310] at the inclusive level and in Ref. [311] for the fully-differential case that the 4FS and the 5FS definitions of the W^+W^- cross section agree at the level of ∼1-2%. For the sake of simplicity, we make use of the easier 4FS approach throughout our calculation.

Many of the ingredients entering our W^+W^-+jet generator have been recycled from the previous implementation of a W^+W^-+jet generator, mentioned at the beginning of this discussion, developed in Ref. [195] for POWHEG-BOX-V2 [177]. Using that one as a starting point, we have integrated it into the POWHEG-BOX-RES framework, in order to take advantage of its efficient phase-space sampling. Then, we have upgraded the generator to achieve NNLO accuracy for W^+W^- production by means of the MINNLO$_{\text{PS}}$ method, that has been adapted in such a way to consistently work in a 4FS.

As far as the physical amplitudes are concerned, all tree-level real and double-real matrix elements (i.e. for $\ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$+1,2-jet production) are evaluated through the POWHEG-

Box interface to Madgraph 4 [11] [312] developed in Ref. [313]. The $\ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$+jet one-loop amplitude is obtained from GoSam 2.0 [12] [314], neglecting one-loop fermion box diagrams, which have been shown to give a negligibly contribution, but also to slow down the code substantially [13](cf. Ref. [195]). The Born-level and one-loop amplitudes for $\ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$ production have been extracted from MCFM [242]. The (one-loop and) two-loop $q\bar{q} \to \ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$ helicity amplitudes that were derived in Ref. [315] are obtained through their implementation in Matrix by suitably adapting the interface created in Ref. [217]. Those amplitudes are known only in the massless approximation, but the effect of including massive quark loops is expected to be negligible because of the smallness of closed fermion-loop contributions. For a fast evaluation of the two-loop amplitudes, we have generated interpolation grids, as discussed in detail in the next section.

The calculation of $\mathcal{D}_{\ell_F}(p_T)$ in Eq. (3.84) involves the evaluation of several convolutions with the parton distribution functions, which are performed through hoppet [225], as standard within the MiNNLO$_{PS}$ code. Moreover, the collinear coefficient functions require the computation of polylogarithms, for which we employ the hplog package [265].

Finally, we report some of the most relevant (non-standard) settings we have used to produce W^+W^- events. Many of them are identically to the ones discussed in Section 4.1.2, and we just recall some of them here for completeness. In particular, to avoid spurious contributions from higher-order logarithmic terms at large p_T we consistently introduce modified logarithms with the choice of $p = 6$, as defined in eq. (10) of Ref. [214]. At small p_T, we use the standard MiNNLO$_{PS}$ scale setting in eq. (14) of Ref. [214], while we activate the option `largeptscales 1` to set the scales entering the NLO W^+W^-+jet cross section at large p_T as in eq. (19) of Ref. [214]. We use those scale settings with the parameter $Q_0 = 0\,\text{GeV}$, and instead regularize the Landau singularity by freezing the strong coupling and the PDFs for scales below 0.8 GeV. We turn on the Powheg-Box option `doublefsr 1`, which was introduced and discussed in detail in Ref. [263]. As far as the parton-shower settings are concerned, we have used the standard ones and we also employed the more conservative global recoil scheme for these studies.

4.3.1 Fast evaluation of the two-loop amplitude

As discussed before, the two-loop helicity amplitudes for the production of a pair of off-shell vector bosons were computed in Ref. [315] and the relevant coefficients functions to construct the amplitudes can be obtained from the publicly available code `VVAMP` [14] [316]. Using those results all $q\bar{q} \to 4$ leptons amplitudes have been implemented in the Matrix

[11] The Madgraph program efficiently generates all tree-level amplitudes for a given process. The growing complexity of the problem with the increasing number of external states is handled by working at the helicity-amplitude level, instead of the squared-amplitude level (as done by more traditional *trace techniques*).

[12] The GoSam program can generate and compute one-loop amplitudes in a fully automated way including QCD and EW corrections within the SM and even for some BSM theories.

[13] Note that there is an option in the Makefile of our code to include the one-loop fermion box diagrams.

[14] The `VVAMP` package provides the one-loop and two-loop helicity amplitudes for EW vector boson pair production, including the decay into 4 leptons, for the quark-antiquark and the gluon-gluon fusion channels.

framework [81, 82]. To exploit this implementation for our calculation, we have compiled MATRIX as a C++ library and linked it to our MiNNLO$_{\rm PS}$ generator using the interface created in Ref. [217].

The evaluation of these two-loop amplitudes turns out to be the bottleneck of the calculation. In fact, it takes on average $\bar{t}_{\texttt{VVAMP}} \approx 1.9\,\text{s}$ to evaluate a single phase-space point, while the evaluation of the tree- and one-loop amplitudes is orders of magnitude faster. Therefore, even though we provide the option to run the code using the exact two-loop amplitudes (by setting to zero the input card parameter `use_interpolator`[15]), all of the results of this paper have been obtained using a four-dimensional cubic spline interpolation procedure for the set of independent two-loop coefficient functions that are required for the evaluation of the two-loop helicity amplitudes. In the following, we present this procedure in detail.

Coefficient functions of the $q\bar{q} \to \ell^+\nu_\ell\ell'^-\bar{\nu}_{\ell'}$ helicity amplitudes

We start by recalling some relevant formulae of Ref. [315] for the helicity amplitudes. Specifically, the physical process is denoted by:

$$q(p_1) + \bar{q}(p_2) \to W^+(p_3) + W^-(p_4) \to \ell^+(p_5) + \nu_\ell(p_6) + \ell'^-(p_7) + \bar{\nu}_{\ell'}(p_8)\,, \qquad (4.19)$$

where p_i are the momenta of the corresponding particles and each of the two off-shell W bosons decays into a neutrino–lepton pair, such that $p_3 = p_5 + p_6$ and $p_4 = p_7 + p_8$. We denote by $\mathrm{M}_{\lambda\lambda_1\lambda_2}$ the bare helicity amplitudes of a general vector-boson pair production process, where λ represents the handedness of the partonic current, while λ_1 and λ_2 stand for the helicities of the two leptonic currents. There are in general just two independent helicity amplitudes M_{LLL} and M_{RLL}, since all the other helicity configurations can be recovered by permutations of external legs [315]. The bare helicity amplitudes are the building blocks of the dressed helicity amplitudes $\mathcal{M}_{\lambda LL}$, which are process specific and for W^+W^- production read

$$\mathcal{M}^{W^+W^-}_{\lambda LL}(p_1,p_2;p_5,p_6,p_7,p_8) = \frac{(4\pi\alpha_{\rm EW})^2}{2\,\sin^2\theta_W}\,\frac{\mathrm{M}_{\lambda LL}(p_1,p_2;p_5,p_6,p_7,p_8)}{(p_3^2 - m_W^2 + i\,\Gamma_W m_{\rm W})(p_4^2 - m_W^2 + i\,\Gamma_W m_{\rm W})}\,, \qquad (4.20)$$

where $\lambda = L, R$. In the previous expression, $\alpha_{\rm EW}$ refers to the EW coupling constant, θ_W to the mixing angle, and $m_{\rm W}$ and Γ_W to the W-boson mass and decay width, respectively. Since a W boson can just couple to left-handed lepton currents (see discussion in Section 1.1), it is clear that $\mathcal{M}_{\lambda RL} = \mathcal{M}_{\lambda LR} = \mathcal{M}_{\lambda RR} = 0$. As shown in Ref. [315], for four-dimensional external states the expression of the bare helicity amplitudes can be written in a compact form using the spinor-helicity formalism (whose basics can be found

[15] More information on the practical usage of the interpolator as well as of the whole W^+W^- code is provided in Appendix B.

in many references, see for instance Ref. [317]):

$$\begin{aligned}\mathrm{M}_{\lambda LL}(p_1,p_2;p_5,p_6,p_7,p_8) = ([i5]\langle 5j\rangle + [i6]\langle 6j\rangle)\Big\{ & E_1^\lambda \langle 15\rangle\langle 17\rangle[16][18] \\ & + E_2^\lambda \langle 15\rangle\langle 27\rangle[16][28] + E_3^\lambda \langle 25\rangle\langle 17\rangle[26][18] \\ & + E_4^\lambda \langle 25\rangle\langle 27\rangle[26][28] + E_5^\lambda \langle 57\rangle[68]\Big\} \\ & + E_6^\lambda \langle 15\rangle\langle j7\rangle[16][i8] + E_7^\lambda \langle 25\rangle\langle j7\rangle[26][i8] \\ & + E_8^\lambda \langle j5\rangle\langle 17\rangle[i6][18] + E_9^\lambda \langle j5\rangle\langle 27\rangle[i6][28]\,, \end{aligned} \tag{4.21}$$

where the two indices i and j are determined by the handedness of the partonic current: $(i,j) = (1,2)$ for $\lambda = L$ and $(i,j) = (2,1)$ for $\lambda = R$. Eq. (4.21) depends on nine complex scalar coefficients E_j^λ, which are functions of the invariant masses p_3^2 and p_4^2 of the two vector bosons and of the two Mandelstam invariants $\tilde{s}$ and $\tilde{t}$, defined as

$$\tilde{s} = (p_1+p_2)^2\,, \qquad \tilde{t} = (p_1-p_3)^2\,. \tag{4.22}$$

Each coefficient E_j^λ receives a contribution from four different classes of diagrams $\mathcal{C}$

$$E_j^\lambda = \delta_{i_1 i_2} \sum_{\mathcal{C}} Q_{q\,q}^{\lambda,W^+W^-,[\mathcal{C}]} E_j^{[\mathcal{C}]}\,, \qquad j = 1,\ldots,9\,, \tag{4.23}$$

where i_1, i_2 represent the colours of the incoming quark and anti-quark, respectively, and $Q_{q\,q}^{\lambda,W^+W^-,[\mathcal{C}]}$ denotes a coupling factor, which is the only process specific ingredient entering Eq. (4.23). Following the labeling introduced in Ref. [315] for the diagram classes, we have for W^+W^- production:

- class A and B, including all diagrams where the two vector bosons are attached to the fermion line, with the W^+ boson adjacent to the incoming quark or antiquark, respectively, whose coupling factors read

$$Q_{q\,q}^{L,W^+W^-,[A]} = Q_{q\,q}^{L,W^+W^-,[B]} = \frac{1}{2\sin^2\theta_W}\,, \tag{4.24}$$

  which are identical to zero for $\lambda = R$;

- class C, containing diagrams where both vector bosons are attached to a fermion loop, where

$$Q_{q\,q}^{\lambda,W^+W^-,[C]} = \frac{n_g}{4\,\sin^2\theta_W}\,, \qquad \text{for } \lambda = L,R\,, \tag{4.25}$$

  with n_g being the number of massless quark generations running in the loop;

- class F_V, collecting form-factor diagrams where the production of the two W bosons is mediated either by a virtual photon ($V = \gamma^*$) or a Z boson ($V = Z$), as shown

in Figure 4.16(b) [16]. In that case we have

$$Q_{q\,q}^{L,W^{\pm}W^{\mp},[F_Z]} = \frac{\mp 1}{\sin^2\theta_W}\frac{(I_q^3 - e_q\sin^2\theta_W)}{\tilde{s} - m_Z^2 - i\Gamma_Z m_Z}\,, \quad Q_{q\,q}^{R,W^{\pm}W^{\mp},[F_Z]} = \frac{\pm e_q}{\tilde{s} - m_Z^2 - i\Gamma_Z m_Z}\,,$$
$$Q_{q\,q}^{\lambda,W^{\pm}W^{\mp},[F_{\gamma^*}]} = \frac{\mp e_q}{\tilde{s}}\,, \qquad \text{for } \lambda = L, R\,, \tag{4.26}$$

where e_q and I_q^3 are the electric charge and isospin number of the incoming quark q, and m_Z and Γ_Z the Z-boson mass and decay width, respectively.

Since the functions E_j^{λ} admit a perturbative expansion as

$$E_j^{\lambda} = E_j^{\lambda,(0)} + \left(\frac{\alpha_s}{2\pi}\right)E_j^{\lambda,(1)} + \left(\frac{\alpha_s}{2\pi}\right)^2 E_j^{\lambda,(2)} + \mathcal{O}(\alpha_s^3)\,, \tag{4.27}$$

the two-loop contribution to the helicity amplitude $\mathrm{M}_{\lambda LL}$ is fully determined once the 45 complex coefficients $E_j^{[\mathcal{C}],(2)}$ are known. In contrast with the helicity amplitude itself, which is a complex-valued function of the full kinematics, the coefficients $E_j^{[\mathcal{C}],(2)}$ just depend on four Lorentz scalars. Therefore, an interpolation procedure that approximates the $E_j^{[\mathcal{C}],(2)}$ coefficients is clearly more feasible. This choice considerably reduces the complexity of the interpolation problem, since it decreases the dimensionality of the space on which the functions are interpolated, at the minor cost of increasing the number of functions to approximate. In essence, this turns our problem into a four-dimensional interpolation of 90 real-valued functions.

However, one should bear in mind that $E_j^{[F_V],(2)}$ does not depend on the type of the vector boson V, so that in our case $E_j^{[F_{\gamma^*}],(2)} = E_j^{[F_Z],(2)} = E_j^{[F],(2)}$. Moreover, any loop correction to the corresponding form-factor diagrams just amounts to a function $\mathcal{F}(\tilde{s})$ which multiplies the tree-level structure, so that at two loops

$$E_j^{[F],(2)} = \mathcal{F}^{(2)}(\tilde{s})E_j^{[F],(0)}\,. \tag{4.28}$$

The tree-level coefficients evaluate to constants:

$$E_j^{[F],(0)} = 0\,, \quad j = 1,...,4\,,$$
$$E_6^{[F],(0)} = E_7^{[F],(0)} = +4\,, \qquad E_5^{[F],(0)} = E_8^{[F],(0)} = E_9^{[F],(0)} = -4\,. \tag{4.29}$$

The dependence on $\tilde{s}$ in $\mathcal{F}^{(2)}(\tilde{s})$ just enters through the ratio of $\tilde{s}$ with the squared of the renormalization scale μ_R. By setting $\mu_R = \sqrt{\tilde{s}}$, the non-vanishing $E_j^{[F],(2)}$ coefficients also become constants, which have been computed in the literature long ago [318–321]. Note that the correct renormalization scale dependence will be recovered through the

[16] Note that another class of form-factor diagrams exists, containing two-loop corrections to DY-type production (see Figure 4.16(c)). This class is evaluated by MATRIX using the corresponding form factor returned by `VVAMP`, as well. Since those form factors are constants, as discussed below, their contribution is handled without interpolation.

MiNNLO$_{\rm PS}$ formulae (cf. appendix D of Ref. [213]). Therefore, only the coefficient functions belonging to families $\mathcal{C} = \{A, B, C\}$ need to be interpolated, reducing the number of real-valued functions that require interpolation from 90 to 54.

Finally, as discussed in Ref. [315] (see eqs.(4.8) to (4.12)), the functions $E_j^{[\mathcal{C}],(2)}$ with $\mathcal{C} = \{A, B, C\}$ are not completely independent from each other, but they can be related via crossings of external legs, either permuting the quark and the anti-quark or the two vector bosons. If we stick to the first permutation, which is the one we exploit here, one can show that the bare helicity amplitudes can be related by exchanging the quark and the anti-quark as

$$\mathrm{M}^{[A]}_{LLL} = -\,\mathrm{M}^{[B]}_{RLL}(p_1 \leftrightarrow p_2), \qquad \mathrm{M}^{[C]}_{LLL} = -\,\mathrm{M}^{[C]}_{RLL}(p_1 \leftrightarrow p_2)\,, \tag{4.30}$$

from which one can derive thirteen equations that express all of the $E_j^{[A],(2)}$ functions through $E_j^{[B],(2)}$ and the four $E_j^{[C],(2)}$ coefficients with $j = \{1, 2, 6, 8\}$ through the remaining ones of the same family. This symmetry property allows us to explicitly generate precomputed grids (according to the procedure outlined in next subsection), which are needed for the interpolation, only for 28 real-valued functions out of 54, while evaluating the remaining ones through crossing symmetry. Note that, in principle, relations similar to Eq. (4.30) can be derived by exploiting the crossing relation of the two vector-boson lines. This second set of relations can be used to further reduce to 20 the number of precomputed grids storing the real-valued functions. We have not made use of this option for the results presented in this work, but the latter is available in the public version of the MiNNLO$_{\rm PS}$ W^+W^- generator as a further way of reducing the memory requirement for the code at running time.

Generation of interpolation grids

As a first step, we have generated rectilinear grids (i.e comprised of congruent parallelotopes) for each of the 28 non-constant independent two-loop coefficient functions $E_j^{[\mathcal{C}],(2)}$ defined in eqs. (4.23) and (4.27), whose exact values have been computed through `VVAMP` on a set of given phase-space points $(p_3^2,\, p_4^2,\, \tilde{s},\, \tilde{t})$ and stored. All results have been obtained by fixing the centre-of-mass energy to $\sqrt{s} = 13\,\mathrm{TeV}$.

As it turns out, a suitable parametrization of the four-dimensional phase-space points $(p_3^2,\, p_4^2,\, \tilde{s},\, \tilde{t})$ is crucial to obtain a good interpolation performance. Moreover, a finer binning is required in those phase-space regions that give a large contribution to the overall integral of the multi-differential cross section, such as resonance-enhanced regions in p_3^2 and p_4^2 around the two W-boson masses. To this end, our grids are defined on a four-dimensional unit hypercube $[0, 1]^4$ with fifty equally spaced bins, where each element $(x_1,\, x_2,\, x_3,\, x_4) \in [0, 1]^4$ is uniquely mapped to a physical phase-space point. The first two axes x_1 and x_2 are related to the invariant masses p_3^2 and p_4^2 through the transformations

$$\begin{aligned} p_3^2 &= m_{\rm W}\Gamma_W\, \tan(z_1(x_1)) + m_W^2\,, \quad \text{with} \quad z_{1,\,\rm min} < z_1(x_1) < z_{1,\,\rm max}\,, \\ p_4^2 &= m_{\rm W}\Gamma_W\, \tan(z_2(x_2)) + m_W^2\,, \quad \text{with} \quad z_{2,\,\rm min}(p_3^2) < z_2(x_2) < z_{2,\,\rm max}(p_3^2)\,, \end{aligned} \tag{4.31}$$

where z_1 and z_2 are continuous functions of x_1 and x_2, respectively. The lower and upper bounds on z_1 and z_2 have been chosen in such a way that the physical range of invariant-mass values is covered. Specifically, $z_{1,\,\mathrm{min}}$ and $z_{1,\,\mathrm{max}}$ are fixed by the choice $40\,\mathrm{GeV}^2 < p_3^2 < s/100$. Through energy conservation $z_{2,\,\mathrm{min}}$ and $z_{2,\,\mathrm{max}}$ depend directly on p_3^2, as it has been made explicit in Eq. (4.31). However, their exact expressions, which we omit here, have been tuned such that the physical mass range of p_4^2 is covered efficiently. The two functions $z_1(x_1)$ and $z_2(x_2)$ are defined piecewise on three subranges of the two intervals in Eq. (4.31). In the central subrange, z_1 and z_2 correspond to a linear mapping, which guarantees that p_3^2 and p_4^2 follow a Breit-Wigner distribution. For the other two subranges of z_1 and z_2, polynomial functions are used such that the off-shell regions are covered by a sufficient number of grid points.

The variables x_3 and x_4 also have a physical interpretation, since they are related to the relativistic velocity β_{W^+} and the cosine of the scattering angle θ_{W^+} of one of the vector bosons in the center-of-mass frame. In particular, we define

$$\beta_{W^+} = a_s + (b_s - a_s)\, x_3\,, \qquad \cos\theta_{W^+} = 1 - 2\,(a_t + (b_t - a_t)\, x_4)\,, \tag{4.32}$$

where $a_{s/t}$ and $b_{s/t}$ determine the range of values allowed for the two physical quantities. Instead of setting $a_{s/t} = 0$ and $b_{s/t} = 1$, we use small cutoffs to avoid numerical instabilities at the kinematic edges. β_{W^+} and $\cos\theta_{W^+}$ can be expressed in terms of $\tilde{s}$ and $\tilde{t}$ as

$$\beta_{W^+} = \frac{\kappa\left(\tilde{s}, p_3^2, p_4^2\right)}{\tilde{s} + p_3^2 - p_4^2}\,, \qquad \cos\theta_{W^+} = \frac{2\tilde{t} + \tilde{s} - p_3^2 - p_4^2}{\kappa\left(\tilde{s}, p_3^2, p_4^2\right)}\,, \tag{4.33}$$

with the Källén function

$$\kappa\left(\tilde{s}, p_3^2, p_4^2\right) \equiv \sqrt{\tilde{s}^2 + p_3^4 + p_4^4 - 2(\tilde{s}\, p_3^2 + p_3^2\, p_4^2 + p_4^2\, \tilde{s})}\,. \tag{4.34}$$

By inverting Eq. (4.33) in the physical region of the process, which is defined by

$$\tilde{s} \geq \left(\sqrt{p_3^2} + \sqrt{p_4^2}\right)^2\,, \qquad \frac{1}{2}(p_3^2 + p_4^2 - \tilde{s} - \kappa) \leq \tilde{t} \leq \frac{1}{2}(p_3^2 + p_4^2 - \tilde{s} + \kappa)\,, \tag{4.35}$$

we can express $\tilde{s}$ and $\tilde{t}$ in terms of the hypercube variables x_3 and x_4.

As illustrated in Ref. [315], the behaviour of the coefficients $E_j^{[\tilde{C}],(2)}$ is not always smooth over the two-dimensional phase space (β_{W^+}, $\cos\theta_{W^+}$) and it can even be divergent close to the highly relativistic ($\beta_{W^+} \to 1$) or highly collinear ($|\cos\theta_{W^+}| \to 1$) regions. One possibility to improve the description of this rapidly changing functional behaviour is to combine different grids to cover the whole phase space (p_3^2, p_3^2, $\tilde{s}$, $\tilde{t}$), instead of simply increasing the number of bins for selected axes. For the case at hand, using four precomputed grids for each of the 28 real-valued (non-constant, independent) functions has proven to significantly improve the performance of the interpolator in some phase-space regions. Even though the definition of the grids is unchanged for x_1 and x_2, by adjusting the values of $a_{s/t}$ and $b_{s/t}$ in Eq. (4.32) we defined four slightly overlapping grids in the (β_{W^+}, $\cos\theta_{W^+}$) phase space, to properly cover the above-mentioned singular regions.

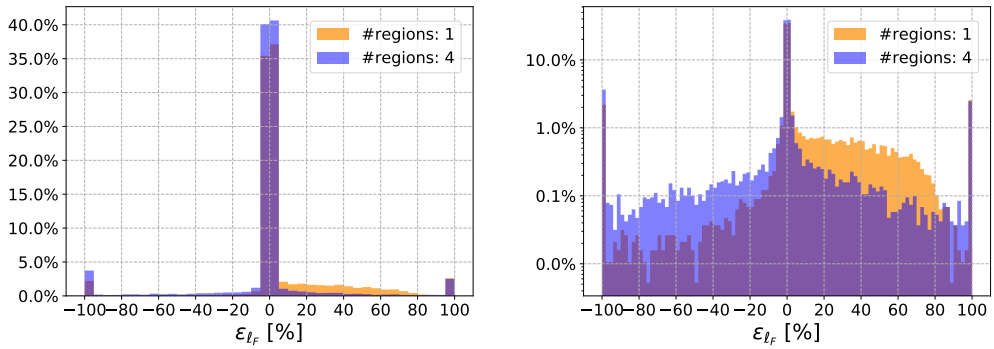

Figure 4.18: Results for the ε_{ℓ_F} accuracy parameter in one sample flavour channel (specifically the $\bar{u}u$ one) using either one (orange histogram) or four (blue histogram) precomputed grids. The same distribution is shown both in linear (left plot) and logarithmic (right plot) scale, where the bins at the edges account for overflow ($|\varepsilon_{\ell_F}| > 100\%$).

Interpolation and validation

At the beginning of each run the grids just need to be read and loaded into memory. Then, for each $E_j^{[\mathcal{C}],(2)}$ any value can be computed by properly interpolating between the values stored in the precomputed grids. To perform this task, we make use of the N-dimensional interpolation library `Btwxt` [322], which just requires the input grids to be rectilinear. The interpolation is achieved through N-dimensional cubic splines [323], which are multivariate piecewise polynomials of degree three. Specifically, `Btwxt` employs cubic Hermite splines, where each polynomial in a given N-dimensional interval is specified by its values and its first derivatives at the corners of the interval itself. The values of the first derivatives are computed according to the Catmull-Rom implementation [324].

In order to quantify the accuracy of our four-dimensional interpolation strategy we use an adimensional parameter ε, which describes the deviation of the interpolated result for the two-loop contribution from its exact expression on a given phase-space point and is defined as

$$\varepsilon_{\ell_F}[\%] = \frac{H^{q_T(2)}_{(\text{int})\ell_F} - H^{q_T(2)}_{(\text{ex})\ell_F}}{H^{q_T(2)}_{(\text{ex})\ell_F}} \cdot 100\,, \tag{4.36}$$

where the four independent Born flavour configurations $\ell_F = \{q\bar{q},\ \bar{q}q;\ \text{for } q = \text{u-type or } q = \text{d-type}\}$ have been considered separately. In Eq. (4.36), $H^{q_T(2)}_{(\text{ex})\ell_F}$ refers to the hard function in the q_T-scheme [262] returned by MATRIX using the evaluation of the exact two-loop coefficients through `VVAMP`, while $H^{q_T(2)}_{(\text{int})\ell_F}$ stands for its value obtained by MATRIX using the interpolation of the two-loop coefficients from the precomputed grid results. For the conversion between the q_T-scheme $H^{q_T(2)}_{(\text{int})\ell_F}$ and the MiNNLO$_{\text{PS}}$ $H^{(2)}_{\ell_F}$ we have used the expressions in Eq. (4.9).

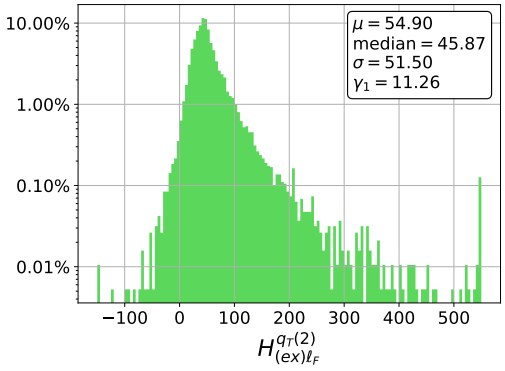

Figure 4.19: Distribution of the values of the two-loop hard function computed using the interpolator without rescue system for the $\bar{u}u$ channel. In the legend the mean (μ), the median, the standard deviation (σ) and the skewness (γ_1) are reported. The edge bins account for overflow.

In Figure 4.18, the distribution of the values of ε_{ℓ_F} is displayed for a selected flavour channel (namely the $\bar{u}u$ one). All the other flavour channels have the same qualitative behaviour. The figure shows the impact on the ε_{ℓ_F} parameter of increasing the number of precomputed grids from one (orange histogram) to four (blue histogram). Besides being essential for a simultaneously accurate description of physical observables over a wide phase-space region, our choice of covering the phase space with four separate grids improves the accuracy of the predictions for the two-loop contribution and yields a more symmetric ε_{ℓ_F} distribution.

We further notice that the bulk of the interpolator predictions (roughly $\gtrsim 80\%$) has an accuracy greater than 5% (i.e $|\varepsilon_{\ell_F}| \leq 5\%$), while almost 95% lie inside the interval $|\varepsilon_{\ell_F}| \leq 100\%$. The remaining fraction of ε_{ℓ_F} values consists of phase-space points where the interpolator poorly reproduces the correct two-loop result. In Figure 4.18, where the ε_{ℓ_F} distribution is reported both in linear and in logarithmic scale, this fraction is clearly visible in the overflow bins at the edges of the histograms. In most cases, these poorly predicted values are associated to phase-space points falling outside the grid boundaries and thus requiring extrapolation of the two-loop coefficient functions $E_j^{[C],(2)}$ outside the grid edges. However, this also means that most of these points lie in kinematic regions where the cross section is strongly suppressed.

To deal with instabilities, a basic rescue-system is introduced. This mechanism takes care of computing the exact $E_j^{[C],(2)}$ coefficient functions with `VVAMP` in all cases where $H^{q_T(2)}_{(\text{int})\ell_F}$ falls outside a process-specific range, where the bulk of the $H^{q_T(2)}_{(\text{ex})\ell_F}$ values lies. Specifically, we have required $-100 < H^{q_T(2)}_{(\text{int})\ell_F} < 500$ as an acceptance interval, where roughly 99.8% of the $H^{q_T(2)}_{(\text{ex})\ell_F}$ distribution is concentrated. In Figure 4.19 this distribution is shown together with the median and the value of the first three moments of the distribution. As it can be inferred from the positive skewness value (or from the fact that

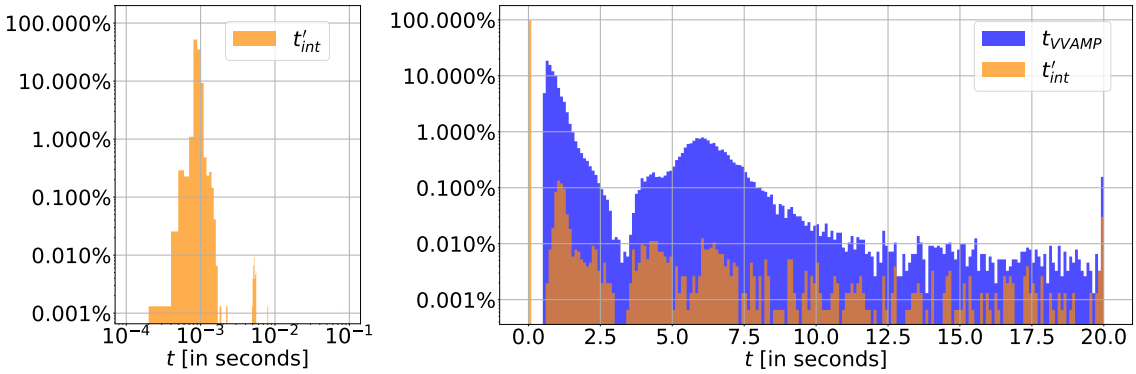

Figure 4.20: Distributions in the evaluation time $t_{\texttt{VVAMP}}$ (blue histogram) and t'_{int} (orange histogram). The last bin accounts for overflow ($t_{\texttt{VVAMP}}, t'_{\text{int}} > 20\,\text{s}$). The left plot resolves the peak of the t'_{int} distribution in double logarithmic scale.

the mean and median do not coincide), the distribution is asymmetric, which is why our acceptance interval for $H^{q_T(2)}_{(\text{ex})\ell_F}$ is not centered around the mean, but rather it extends to higher values to partially account for the long distribution tail on the right of the peak. This simple criterium suffices to catch the small fraction of ε_{ℓ_F} outliers that would compromise the stability of the results. Some phase-space points remain that elude the rescue-system and where the two-loop coefficients are not computed accurately, but we have checked that they have a negligible impact on the physical results, as it will be discussed below (see Figure 4.21).

Clearly, the advantage of using the interpolation approach compared to the full evaluation of the two-loop coefficients is the gain in time performance. Indeed, the average time required by VVAMP ($\bar{t}_{\texttt{VVAMP}}$) and the interpolator ($\bar{t}_{\text{int}}$) to evaluate the two-loop contribution turns out to differ by three orders of magnitude, while the improvement is still roughly a factor forty if one uses the rescue-system ($\bar{t}'_{\text{int}}$):

$$\bar{t}_{\texttt{VVAMP}} \approx 1.9\,\text{s}\,, \qquad \bar{t}_{\text{int}} \approx 0.9 \times 10^{-3}\,\text{s}\,, \qquad \bar{t}'_{\text{int}} \approx 4.5 \times 10^{-2}\,\text{s}\,. \tag{4.37}$$

As complementary information, Figure 4.20 shows the time distributions of $t_{\texttt{VVAMP}}$ (blue histogram) and t'_{int} (orange histogram). The bulk of the VVAMP evaluation times (roughly 80%) is located in the time interval $0.5s < t_{\texttt{VVAMP}} < 2.0\,\text{s}$, with a small, but not negligible fraction of phase-space points requiring a CPU time between $5.0s < t_{\texttt{VVAMP}} < 7.5\,\text{s}$, and about 0.1% exceeding 20 s (visible from the overflow bin). When using the interpolator, more than 99% of the evaluations just require some hundredths of a second. The small number of phase-space points with a CPU time $t'_{\text{int}} > 0.5\,\text{s}$ are associated to the values catched by the rescue-system. Those timings have been obtained on machines with Intel Haswell Xeon E5-2698 processors with 2.3 GHz per core.

Our implementation of a faster evaluation of the two-loop amplitudes through interpolation is tested by looking at its impact on physical predictions, especially on some relevant differential distributions in the inclusive phase space. All results have been

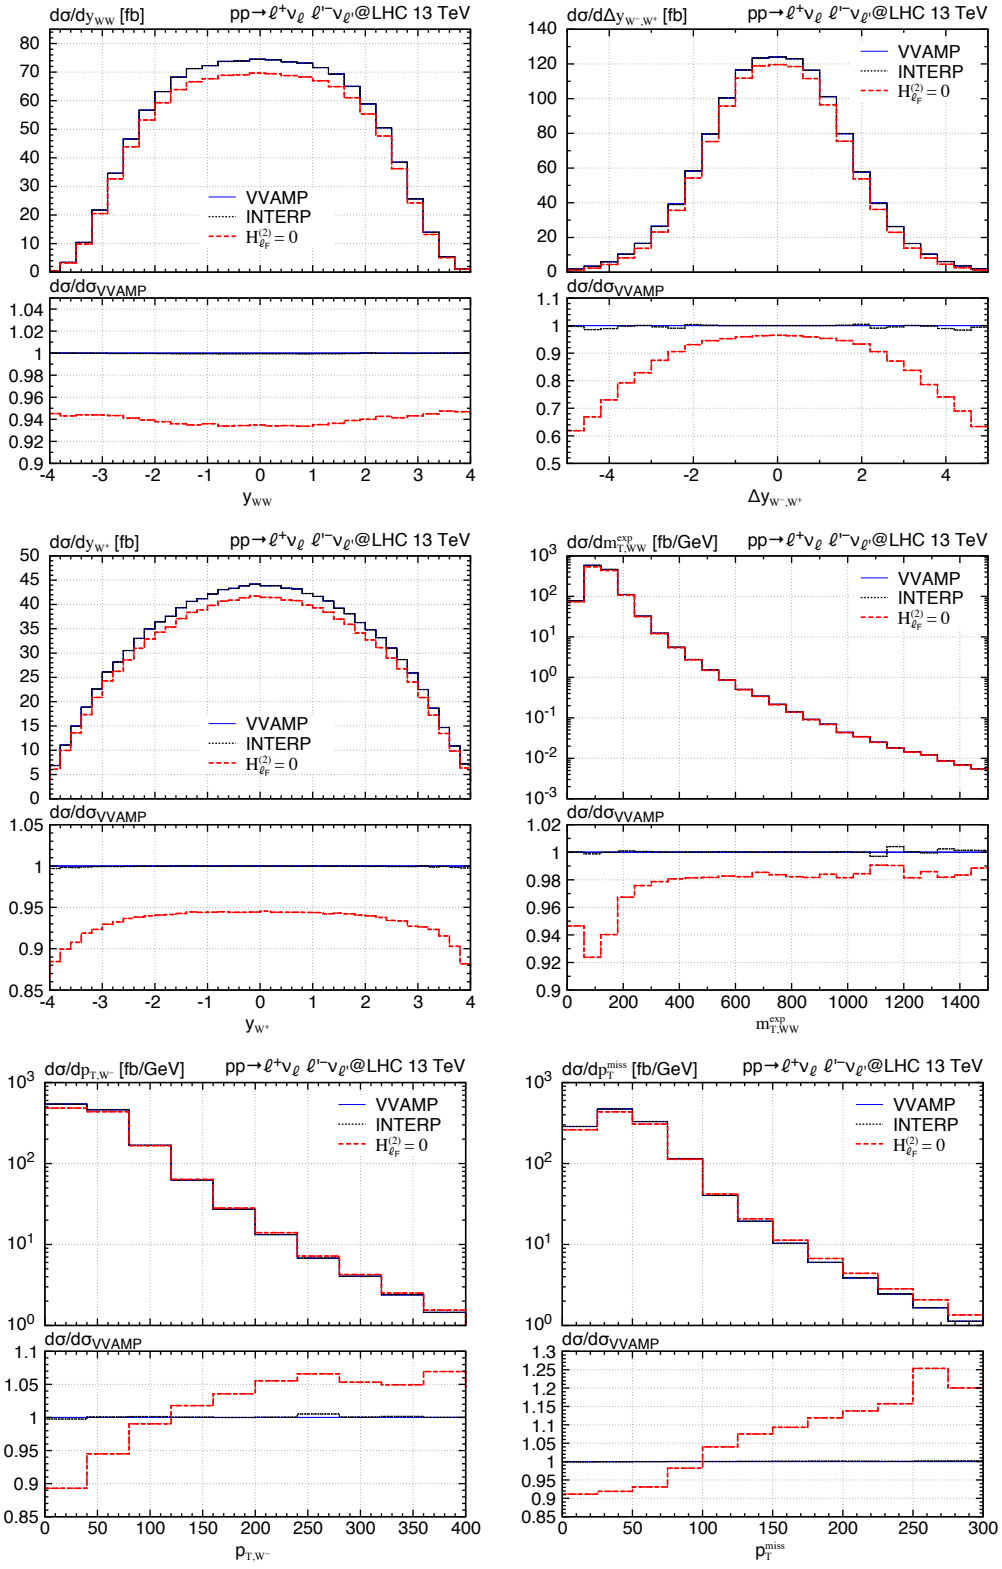

Figure 4.21: POWHEG stage 2 distributions (from left to right and top to bottom) in y_{WW}, $\Delta y_{W^-,W^+}$, y_{W^+}, $m^{\text{exp}}_{\text{T},WW}$, p_{T,W^-} and $p_{\text{T,miss}}$ for the `VVAMP` (blue, solid line), the interpolator (black, dotted line) and the $H^{(2)}_{\ell_\text{F}} = 0$ (red, dotted line) results.

obtained at the level of the Monte Carlo integration of the cross section (i.e POWHEG stage 2), so that no parton shower radiation or hadronization effects, which would not be relevant for the validation, have been included.

First, it is worth mentioning that the code with the interpolation of the two-loop amplitudes reproduces accurately the exact inclusive cross section, with discrepancies of the order of about 0.4 permille, which are well within statistical uncertainties. Then, in Figure 4.21 we show representative plots that compare the exact `VVAMP` predictions (blue, solid line) against the results with interpolation (black, dotted line) for the rapidity of the W^+W^- pair (y_{WW}), the rapidity difference between the two W-bosons ($\Delta y_{W^-,W^+}$), the rapidity of positively-charged W boson (y_{W^+}), the experimental definition of the transverse mass of the W^+W^- pair

$$m_{\mathrm{T},WW}^{\mathrm{exp}} = \sqrt{\left(E_{\mathrm{T},\ell_1} + E_{\mathrm{T},\ell_2} + p_{\mathrm{T,miss}}\right)^2 - \left(\mathbf{p}_{\mathrm{T},\ell_1} + \mathbf{p}_{\mathrm{T},\ell_2} + \mathbf{p}_{\mathrm{T,miss}}\right)^2}\,, \tag{4.38}$$

the transverse momentum of the negatively-charged W boson (p_{T,W^-}) and the missing transverse momentum ($p_{\mathrm{T,miss}}$). We stress that many more distributions than those included in this section have been carefully examined and verified to show a very good agreement between the analytic and interpolated results. Moreover, in order to highlight the phase-space regions where the two-loop corrections provide a large contribution to the cross section, a third curve (red, dotted line) has been included in all plots, obtained by setting the MiNNLO$_{\mathrm{PS}}$ $H_{\ell_{\mathrm{F}}}^{(2)} = 0$. The lower panel of the plots displays the bin-by-bin ratio using the `VVAMP` curve as a reference.

From all plots, it is evident that the interpolator reproduces correctly the differential distributions in all kinematic regimes, with only small fluctuations at very high values of $\Delta y_{W^-,W^+}$ (at most of the order of 2%) or high values of $m_{\mathrm{T},WW}^{\mathrm{exp}}$ or p_{T,W^-} (where differences are well below one percent). These are the regions where the two-loop contribution has the largest impact on the cross section. Indeed, from the y_{WW} distribution it is evident that $H_{\ell_{\mathrm{F}}}^{(2)}$ has a 5-6% effect on the cross section, and contributes uniformly to this observable, while for instance for $|\Delta y_{W^-,W^+}| > 4.5$ and $|y_{W^+}| > 4$ it induces a positive correction that reaches more than 30% and 10%, respectively. For transverse-momentum distributions, such as p_{T,W^-} or $p_{\mathrm{T,miss}}$, the two-loop contribution has a positive impact for relatively low transverse momenta (roughly for $p_{\mathrm{T}} < 100\,\mathrm{GeV}$) of at most 10%, while it yields an increasingly negative correction for very large transverse momenta.

4.3.2 Validation settings

In this section, as well as in Section 4.3.3, 4.3.4 and 4.3.5, we present a detailed comparison of our MiNNLO$_{\mathrm{PS}}$ W^+W^-+jet generator against other theory predictions, which serves the purpose of robustly validating our code. We do not include the loop-induced gluon-fusion contribution in all our NNLO results, since our focus is to study the genuine NNLO corrections to the $q\bar{q}$ initiated process. The leading-order gluon-gluon initiated contribution enters at NNLO: NLO QCD corrections to it are known and can be incoherently added to the predictions presented here through a dedicated calculation. We perform such a calculation for another diboson process in Section 4.4. In this

context, since the loop-induced gluon-fusion channel is known to have a sizable effect for W^+W^- production, our choice not to include it prevents us to directly compare against experimental data, except for a few cases.

In what follows, we consider $\sqrt{s} = 13\,\text{TeV}$ proton–proton collisions at the LHC and present predictions for $pp \to \ell^+\nu_\ell\,\ell'^-\bar{\nu}_{\ell'} + X$ production with $\ell = e$ and $\ell' = \mu$. The EW parameters are determined in the G_μ scheme, therefore computing the EW coupling as $\alpha_{G_\mu} = \sqrt{2}\,G_\mu m_W^2(1 - m_W^2/m_Z^2)/\pi$ and the mixing angle as $\cos^2\theta_W = m_W^2/m_Z^2$. We use the following PDG [325] values as inputs: $G_F = 1.16639\times10^{-5}\,\text{GeV}^{-2}$, $m_W = 80.385\,\text{GeV}$, $\Gamma_W = 2.0854\,\text{GeV}$, $m_Z = 91.1876\,\text{GeV}$, and $\Gamma_Z = 2.4952\,\text{GeV}$. We set the CKM matrix to unity, which, because of unitarity and the fact that we consider only massless external quarks, is a very good approximation, as explained in Ref. [211]. As described at the beginning of Section 4.3, the four-flavour scheme with $n_f = 4$ massless quark flavours and massive bottom and top quarks is used to define a top-free W^+W^- cross section by removing all contributions with final-state bottom quarks. Consistently with this choice, we use the $n_f = 4$ NNLO set of the NNPDF3.0 [264] parton densities. More precisely, we make use of the standard MiNNLO$_{\text{PS}}$ hybrid grid approach, described in Section 3.3.2. For the central renormalization and factorization scales, we also stick to the usual settings for MiNNLO$_{\text{PS}}$ and MiNLO$'$ discussed in Section 3.3.2. Perturbative uncertainties are estimated from customary 7-point variations, i.e. by varying μ_{R} and μ_{F} around the central scale by a factor of two while respecting the constraint $0.5 \leq \mu_{\text{R}}/\mu_{\text{F}} \leq 2$.

We compare our MiNNLO$_{\text{PS}}$ (and MiNLO$'$) predictions to NNLOPS results from a previous implementation of the W^+W^-+jet generator, that we have already mentioned at the beginning of this discussion. Results for the latter generator, presented in Ref. [211], are based on a MiNLO$'$ calculation with $\mu_{\text{R}} = \mu_{\text{F}} = p_{\text{T},WW}$, but use NNLO predictions for the reweighting with

$$\mu_{\text{R}} = \mu_{\text{F}} = \mu_0 \equiv \frac{1}{2}\,(m_{\text{T},W^+} + m_{\text{T},W^-})\,, \quad m_{\text{T},W^\pm} = \sqrt{m^2_{\ell^{(\prime)}\nu_{\ell^{(\prime)}}} + p^2_{\text{T},\ell^{(\prime)}\nu_{\ell^{(\prime)}}}}\,, \tag{4.39}$$

where $m_{\ell\nu_\ell}$ and $p_{\text{T},\ell\nu_\ell}$ ($m_{\ell'\nu_{\ell'}}$ and $p_{\text{T},\ell'\nu_{\ell'}}$) are the invariant masses and the transverse momenta of the reconstructed W bosons. The setting in Eq. (4.39) is therefore the effective scale used in the NNLOPS calculation of Ref. [211], where the perturbative uncertainties are obtained from 7-point scale variations that are assumed correlated in the reweighting. For the $p_{\text{T},WW}$ spectrum and the jet-vetoed cross section we also compare against more accurate analytically resummed predictions obtained with Matrix+RadISH [272, 273, 326], where we have chosen

$$\mu_{\text{R}} = \mu_{\text{F}} = \mu_0 \equiv m_{\text{T},WW}\,, \quad Q_{\text{res}} = m_{WW}\,, \quad m_{\text{T},WW} = \sqrt{m^2_{WW} + p^2_{\text{T},WW}}\,, \tag{4.40}$$

with m_{WW} the invariant mass of the W^+W^- pair. Here, scale uncertainties are obtained not just from customary 7-point variations, but also by varying the resummation scale Q_{res} by a factor of two around its central value, while keeping μ_{R} and μ_{F} fixed to μ_0. For some observables it is instructive to also compare to fixed-order NNLO predictions with both the scale settings in Eq. (4.39) and the ones in Eq. (4.40), which we have obtained

	`fiducial-1-JV` [328]	`fiducial-2-JV` [329]
Lepton cuts	$p_{\mathrm{T},\ell} > 25\,\mathrm{GeV} \quad \lvert\eta_\ell\rvert < 2.4$ $m_{\ell\ell} > 10\,\mathrm{GeV}$	$p_{\mathrm{T},\ell} > 27\,\mathrm{GeV} \quad \lvert\eta_\ell\rvert < 2.5$ $p_{\mathrm{T},\ell\ell} > 30\,\mathrm{GeV} \quad m_{\ell\ell} > 55\,\mathrm{GeV}$
Neutrino cuts	$p_{\mathrm{T,miss}} > 20\,\mathrm{GeV} \quad p_{\mathrm{T,miss,rel}} > 15\,\mathrm{GeV}$	$p_{\mathrm{T,miss}} > 20\,\mathrm{GeV}$
Jet cuts	anti-k_{T} algorithm with $R = 0.4$ $N_{\mathrm{jet}} = 0$ for $p_{\mathrm{T},j} > 25\,\mathrm{GeV}$ $\lvert\eta_{\mathrm{j}}\rvert < 2.5 \quad \Delta R_{ej} > 0.3$ $N_{\mathrm{jet}} = 0$ for $p_{\mathrm{T},j} > 30\,\mathrm{GeV}$ $\lvert\eta_{\mathrm{j}}\rvert < 4.5 \quad \Delta R_{ej} > 0.3$	anti-k_{T} algorithm with $R = 0.4$ $N_{\mathrm{jet}} = 0$ for $p_{\mathrm{T},j} > 35\,\mathrm{GeV}$

Table 4.6: Fiducial cuts used in two different setups, see text for details.

with MATRIX [81, 82]. In this case, perturbative uncertainties are again estimated from 7-point scale variations.

Finally, for the matching to the parton shower we employ PYTHIA8 [124] with the Monash 2013 tune [327] (specifically setting `py8tune 14` in the input card of the process). Since in this study our focus is on the comparison with other theory predictions, we do not include any effect from hadronization, underlying-event models, or a QED shower. Such effects can, however, be directly included and studied by any user of our program through the PYTHIA8 interface of POWHEG-BOX-RES.

Since the W^+W^- cross section is finite at the LO without any cuts, we present results in the fully inclusive W^+W^- phase space, referred to as `setup-inclusive`. Additionally, we consider two sets of fiducial cuts, which are summarized in Table 4.6. The first one corresponds to an earlier ATLAS 13 TeV analysis [328] and it is identical to that used in the NNLOPS calculation of Ref. [211], which allows us to compare directly our MINNLO$_{\mathrm{PS}}$ predictions with the fiducial NNLOPS results of Ref. [211]. We refer to it as `fiducial-1-JV` in the following. We note that `fiducial-1-JV` involves a two-fold jet-veto requirement, vetoing all jets in the rapidity region $|\eta_{\mathrm{j}}| < 2.5$ and separated from the leptons by $\Delta R_{ej} > 0.3$ with $p_{\mathrm{T},j} > 25\,\mathrm{GeV}$ and all jets in the rapidity region $|\eta_{\mathrm{j}}| < 4.5$ and separated from the leptons by $\Delta R_{ej} > 0.3$ with $p_{\mathrm{T},j} > 30\,\mathrm{GeV}$. The second setup instead corresponds to the most recent ATLAS 13 TeV measurement of Ref. [329], and it was used to study high-accuracy resummed predictions for W^+W^- production in Ref. [272]. This setup, referred to as `fiducial-2-JV` in the following, is useful for two reasons. First, at variance with `fiducial-1-JV`, it includes a single jet-veto cut for jets with $p_{\mathrm{T},j} > 35\,\mathrm{GeV}$. This allows us to directly compare against the NNLO+NNLL resummed predictions for the $p_{\mathrm{T},WW}$ spectrum with a jet veto [272]. Note that to facilitate this comparison, we have removed the jet rapidity (η_{j}) requirement from `fiducial-2-JV` [329], which has a numerically tiny effect. Second, `fiducial-2-JV` is used to compare against data, since Ref. [329] measured the fiducial cross section as a function of the jet-veto cut to validate theory predictions for the jet-vetoed W^+W^- cross section. It is worth recalling here once again that jet-veto requirements are crucial for experimental

W^+W^- analyses in order to suppress the large top-quark backgrounds. In addition, we introduce `fiducial-1-noJV` and `fiducial-2-noJV` for the same fiducial setups as given in Table 4.6, but without any restriction on the jet activity. Those are relevant to study the $p_{\mathrm{T},WW}$ distribution inclusive over jet radiation as well as the cross section as a function of the jet-veto cut. Besides jet-veto requirements, the two setups in Table 4.6 involve standard cuts on the transverse momentum ($p_{\mathrm{T},\ell}$) and pseudorapidity (η_ℓ) of the charged leptons as well as a lower cut on the invariant mass of the dilepton pair ($m_{\ell\ell}$) and on the missing transverse momentum ($p_{\mathrm{T,miss}}$). Setup `fiducial-2-JV` includes also a lower cut on the transverse momentum of the dilepton pair ($p_{\mathrm{T},\ell\ell}$), while setup `fiducial-1-JV` cuts on the so-called relative missing transverse momentum ($p_{\mathrm{T,miss,rel}}$). The latter denotes the component of the missing transverse-momentum vector perpendicular to the direction of the closest lepton in the azimuthal plane, and is defined as

$$p_{\mathrm{T,miss,rel}} = \begin{cases} p_{\mathrm{T,miss}} \cdot \sin|\Delta\phi| & \text{for } \Delta\phi < \pi/2\,, \\ p_{\mathrm{T,miss}} & \text{for } \Delta\phi > \pi/2\,, \end{cases} \tag{4.41}$$

where $\Delta\phi$ denotes the azimuthal angle between the missing transverse-momentum vector vector and the nearest lepton.

4.3.3 Integrated cross sections

We start the presentation of phenomenological results by discussing integrated cross sections in Table 4.7. In particular, we report predictions in the fully inclusive and the two fiducial phase spaces introduced in Section 4.3.2 for MiNLO′, MiNNLO$_{\mathrm{PS}}$, NNLOPS [211] as well as two fixed-order NNLO predictions obtained with Matrix [81, 82] using the scale settings of Eq. (4.39) and Eq. (4.40). We summarize our main observations in the following:

- It is clear that NNLO accuracy is crucial for an accurate prediction of the W^+W^- cross section, since the MiNLO′ result is about 14% lower than the MiNNLO$_{\mathrm{PS}}$ one not only for `setup-inclusive`, but also after including the `fiducial-1-JV` and `fiducial-2-JV` cuts. In fact, in all cases the MiNNLO$_{\mathrm{PS}}$ prediction is outside the upper uncertainty boundary of the MiNLO′ one. This is not surprising since for W^+W^- production also at fixed order the NLO uncertainty band does typically not include the central value of the NNLO prediction [211, 311]. Additionally, the precision at NNLO is substantially improved, with scale uncertainties reduced by more than a factor of three.

- The NNLO-accurate predictions compare well against one another. They are compatible within their respective scale uncertainties and the central predictions are all within less than 3%. Indeed, small differences are expected from the fact that those predictions differ by terms beyond NNLO accuracy. Note that the NNLOPS and the NNLO calculations with $\mu_0 = \frac{1}{2}\,(m_{\mathrm{T},W^+} + m_{\mathrm{T},W^-})$ are very close, both in terms of central values and uncertainties, since the former is actually reweighted to the latter prediction in the inclusive phase space. The fact that the inclusive

$\sigma(pp \to \ell^+\nu_\ell\,\ell'^-\nu_{\ell'})$ [fb]	`setup-inclusive`	`fiducial-1-JV`	`fiducial-2-JV`
MiNLO′	$1146.4(3)^{+4.7\%}_{-5.6\%}$	$182.8(1)^{+5.7\%}_{-9.4\%}$	$141.3(1)^{+4.9\%}_{-8.1\%}$
MiNNLO$_{\rm PS}$	$1299.8(8)^{+2.0\%}_{-1.2\%}$	$211.7(3)^{+1.5\%}_{-1.5\%}$	$161.9(3)^{+1.1\%}_{-1.1\%}$
NNLO+PS [211]	$1308.9(3)^{+1.7\%}_{-1.6\%}$	$206.4(1)^{+2.2\%}_{-2.3\%}$	$159.0(1)^{+1.7\%}_{-1.8\%}$
NNLO $\mu_0 = (m^2_{{\rm T},W^+} + m^2_{{\rm T},W^-})^{0.5}/2$	$1306.5(5)^{+1.6\%}_{-1.6\%}$	$206.5(1)^{+1.0\%}_{-0.7\%}$	$158.9(5)^{+0.8\%}_{-0.6\%}$
NNLO $\mu_0 = m_{{\rm T},WW}$	$1284.9(10)^{+1.4\%}_{-1.3\%}$	—	$160.8(3)^{+1.0\%}_{-0.8\%}$
ATLAS−gg [328]	$1481 \pm 59_{\rm (stat)}$ $\pm 154_{\rm (syst)} \pm 108_{\rm (lumi)}$	$236.5 \pm 10_{\rm (stat)}$ $\pm 25_{\rm (syst)} \pm 5.5_{\rm (lumi)}$	—
ATLAS−gg [329]	—	—	$178.5 \pm 2.5_{\rm (stat)}$ $\pm 12.7_{\rm (syst)} \pm 4_{\rm (lumi)}$
CMS−gg [330]	$1289 \pm 68_{\rm (stat)}$ $^{\pm 67_{\rm (exp.\,syst)}}_{\pm 76_{\rm (th.\,syst)}} \pm 42_{\rm (lumi)}$	—	—
CMS−gg [331]	$1316 \pm 65_{\rm (stat)}$ $\pm 23_{\rm (syst)} \pm 38_{\rm (lumi)}$	—	—

Table 4.7: W^+W^- cross sections in the fully inclusive phase space and in the fiducial phase spaces defined in Table 4.6. We compare our MiNLO′ and MiNNLO$_{\rm PS}$ predictions to the NNLOPS results of Ref. [211] and to the NNLO cross section obtained with different settings of $\mu_{\rm R}$ and $\mu_{\rm F}$. All NNLO corrections to $q\bar{q}$-induced W^+W^- production are taken into account, while the loop-induced gg contribution is excluded. In the last rows, the comparison to CMS and ATLAS data is shown. For the measured inclusive cross sections we have assumed a branching fraction of BR$(W^\pm \to \ell^\pm\nu_\ell) = 0.108987$, consistently evaluated with our inputs, and applied one factor for each of the two W bosons. The measured fiducial cross sections have been divided by a factor two so that they correspond to $pp \to \ell^+\nu_\ell\,\ell'^-\nu_{\ell'}$ production with $\ell = e$ and $\ell' = \mu$. In addition, we have subtracted the loop-induced gluon-fusion contribution from the central value of the data. For the inclusive results and the `fiducial-1-JV` result we used the prediction for gg (non-resonant) cross section quoted in table 5 of the ATLAS analysis in Ref. [328]. For the `fiducial-2-JV` result we used the ggLO result in table 2 of Ref. [332]. The ATLAS measurement of Ref. [328] includes resonant Higgs contributions, which have been subtracted from that data as well, using the corresponding prediction quoted in table 5 of that paper.

MiNNLO$_{\rm PS}$ result is about 0.7% below the NNLOPS one is due to the different scale choice and treatment of terms beyond accuracy. Indeed, the second NNLO prediction with a scale choice of $\mu_0 = m_{{\rm T},WW}$ is even lower than the MiNNLO$_{\rm PS}$ one.

- In the fiducial phase space MiNNLO$_{\text{PS}}$ and NNLOPS predictions differ by up to 2.5%, with only a moderate overlap of the uncertainty bands. The main effect here is related to the applied jet veto cuts. Indeed, the actual uncertainties related to jet vetos are generally underestimated, as discussed below.

- Some caution is advised regarding the quoted scale uncertainties. First of all, the quoted uncertainties generally appear to be quite small, and potentially at the edge of providing a realistic estimate of the true uncertainty. Clearly, the MiNLO$'$ uncertainty does not cover the inclusion of NNLO corrections through MiNNLO$_{\text{PS}}$. For MiNNLO$_{\text{PS}}$ and fixed-order the NNLO uncertainties further decrease when the jet-veto requirements are imposed. Such behaviour is not new [333] and especially for the jet-vetoed predictions one may consider more conservative approaches to estimate the perturbative uncertainties (see for instance Refs. [333, 334]).

- Finally, there is a good agreement of MiNNLO$_{\text{PS}}$ results with data from ATLAS and CMS in both inclusive and fiducial phase-space regions. The measured cross sections agree mostly within one and at most within two standard deviations.

4.3.4 Some differential distributions

We now turn to discussing results for some differential distributions, which mainly have validation purposes. Specifically, we compare our MiNNLO$_{\text{PS}}$ to MiNLO$'$ and NNLOPS predictions in the inclusive and the fiducial phase space. This allows us, on the one hand, to study the effect of NNLO corrections through MiNNLO$_{\text{PS}}$ with respect to MiNLO$'$ and, on the other hand, to assess the compatibility of the MiNNLO$_{\text{PS}}$ predictions with the known NNLOPS results.

Unless indicated otherwise, the plots in this section and in Section 4.3.5 are organized as follows: there is a main frame, which shows differential distributions for the MiNNLO$_{\text{PS}}$ (blue, solid), MiNLO$'$ (black, dotted), and NNLOPS (magenta, dash-dotted) predictions. In a lower frame we show bin-by-bin ratios of all curves to the central MiNNLO$_{\text{PS}}$ result. In some cases, where it is instructive to compare to the fixed-order results, we show fixed-order NNLO distributions for $\mu_0 = \frac{1}{2}\,(m_{\text{T},W^+} + m_{\text{T},W^-})$ (green, long-dashed) and/or for $\mu_0 = m_{\text{T},WW}$ (red, dashed). We note that we refrain from showing fixed-order NNLO predictions for most observables as the NNLOPS results correspond to a scale setting of $\mu_0 = \frac{1}{2}\,(m_{\text{T},W^+} + m_{\text{T},W^-})$ and are, in general, numerically very close to the respective fixed-order NNLO cross section.

Inclusive phase space

We start by discussing distributions in the inclusive phase space. We have considered a large number of relevant distributions both of the leptonic final states and of the reconstructed W-bosons. A selection of those which reflect some general features is presented in Figure 4.22. Since experimentally the W bosons can not be directly reconstructed and the fully inclusive phase space can not be covered by the detectors in any case, we

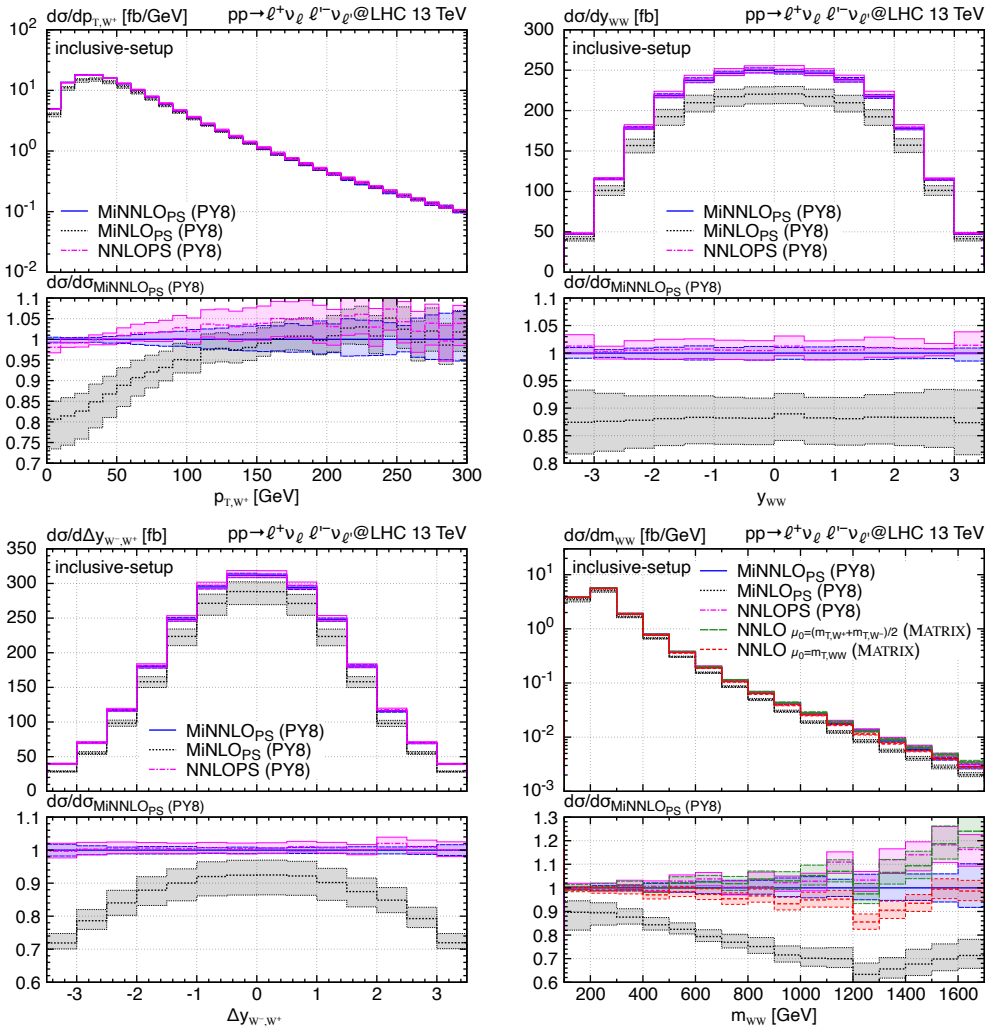

Figure 4.22: Differential distributions in the `setup-inclusive`.

follow here a more theoretical motivation and study observables related to the reconstructed W bosons rather than their decay products. In particular, Figure 4.22 shows the transverse-momentum spectrum of the W^+ boson (p_{T,W^+}), the rapidity distribution of the W-boson pair (y_{WW}), the rapidity difference between the two W bosons ($\Delta y_{W^-,W^+}$), and the invariant-mass distribution of the W-boson pair (m_{WW}).

For the p_{T,W^+} spectrum, the MiNNLO$_{\mathrm{PS}}$ prediction is in full agreement with the NNLOPS result, which is particularly striking in the low-p_{T,W^+} region since scale uncertainties are only at the level of $\pm 1\%$. At larger values of p_{T,W^+}, the uncertainty bands of the NNLOPS accurate predictions widen and reach about $\pm 5\%$. This indicates that this region is predominantly filled by higher-order (real) radiative corrections with at least one jet, and that the formal accuracy is somewhat decreased by one order. Indeed, in the region $p_{\mathrm{T},W^+} \gtrsim 100\,\mathrm{GeV}$ the NNLOPS predictions become fully compatible with

the MiNLO′ result, also in terms of the size of the uncertainty bands. By contrast, for smaller p_{T,W^+} we observe large NNLO corrections with respect to MiNLO′ that reach almost 20% and substantially reduced scale uncertainties.

Also for the y_{WW} and $\Delta y_{W^-,W^+}$ distributions we find fully compatible results with overlapping uncertainty bands when comparing MiNNLO$_{\mathrm{PS}}$ and NNLOPS predictions. While the NNLO corrections compared to MiNLO′ are relatively flat for y_{WW}, we find that the corrections increase substantially at larger values of $\Delta y_{W^-,W^+}$, reaching $\sim +30\%$ for $\Delta y_{W^-,W^+} \gtrsim 3$. This behaviour was observed already in Ref. [211], and our result confirms that this large effect is not a remnant of the scale setting in the NNLOPS calculation, but a genuine NNLO correction.

Similarly sizeable NNLO corrections are observed also at large values of m_{WW}. This is also one of the few regions of phase space that we found where MiNNLO$_{\mathrm{PS}}$ and NNLOPS predictions do not agree at the level of a few percent. While up to $m_{\mathrm{WW}} \lesssim 1\,\mathrm{TeV}$ the MiNNLO$_{\mathrm{PS}}$ and NNLOPS results are fully compatible, they start deviating at larger invariant masses, reaching differences of about 20% at $m_{\mathrm{WW}} = 1.7\,\mathrm{TeV}$. Those differences can be traced back to the different scale settings in the MiNNLO$_{\mathrm{PS}}$ and NNLOPS calculations. Indeed, comparing the additional NNLO results shown for the m_{WW} distribution, we notice a relatively large spread between the green long-dashed curve with central scale $\mu_0 = \frac{1}{2}\,(m_{\mathrm{T},W^+} + m_{\mathrm{T},W^-})$ and the red dashed curve with $\mu_0 = m_{\mathrm{T},WW}$, which is of the same size as (or even slightly larger than) the observed differences between MiNNLO$_{\mathrm{PS}}$ and NNLOPS. As expected, the NNLOPS result is close to the NNLO one with $\mu_0 = \frac{1}{2}\,(m_{\mathrm{T},W^+} + m_{\mathrm{T},W^-})$, while the MiNNLO$_{\mathrm{PS}}$ prediction is somewhat in-between the two NNLO results, but slightly closer to the one with $\mu_0 = m_{\mathrm{T},WW}$. Thus, the origin of the observed differences are terms beyond NNLO accuracy. Although the uncertainty bands increase to about 10% towards large m_{WW}, the two NNLO+PS accurate predictions do not (or only barely) overlap for $m_{\mathrm{WW}} \gtrsim 1\,\mathrm{TeV}$, indicating that plain 7-point scale variations do not represent a realistic estimate of the actual size of uncertainties in that region of phase space. One may ask the question whether one of the two scale choices can be preferred. Although one might assume that m_{WW} would be the natural scale of the m_{WW} distribution, the situation is actually not that clear. This was discussed in some detail in Ref. [211], and it boils down to the fact that for s-channel topologies m_{WW} would be the more natural scale, while for t-channel topologies the transverse masses of the W bosons reflect better the natural scale of the process. Since both topologies appear in W^+W^- production already at the LO and they interfere, it is hard to argue in favour of any of the two scale choices. As a result, and since now there are two NNLO+PS accurate predictions available, the difference between the two should be regarded as an uncertainty induced by terms beyond NNLO accuracy. Moreover, one could introduce a different setting of the hard scales at high transverse momenta in the MiNNLO$_{\mathrm{PS}}$ calculation (i.e. in the W^+W^-+jet part) as a further probe of missing higher-order terms.

In summary, we find that MiNNLO$_{\mathrm{PS}}$ and NNLOPS predictions are in excellent agreement for essentially all observables we considered in the inclusive phase space that are genuinely NNLO accurate. This indicates the robustness of NNLO+PS predictions for such observables. For the few exceptions, like large m_{WW}, we could trace back the origin of the differences to terms beyond accuracy that are induced by the different scale

settings. Moreover, in all cases the NNLO corrections substantially reduce the scale uncertainties with respect to the MiNLO′ prediction. Notice, however, that in the bulk region of the inclusive phase space the MiNNLO$_{\rm PS}$ uncertainty bands are a bit smaller than the NNLOPS ones, as already observed for the fully inclusive cross section.

Fiducial phase space

We continue our comparison by considering differential distributions in the `fiducial-1-JV` phase space. Here, we have selected a set of leptonic observables that are directly measured by the experimental analyses, cf. Refs. [328, 329], and which represent well the general features of all observables we considered. To this end, Figure 4.23 shows the distributions in the transverse momentum of the leading lepton (p_{T,ℓ_1}), in the invariant mass ($m_{\ell\ell}$), transverse momentum ($p_{\mathrm{T},\ell\ell}$), rapidity ($y_{\ell\ell}$) and azimuthal difference ($\Delta\phi_{\ell\ell}$) of the dilepton pair, and in an observable particularly sensitive to new physics effects [335], defined through the separation in η of the two leptons:

$$|\cos(\theta^\star)| = |\tanh\left(\Delta\eta_{\ell\ell}/2\right)| \,. \tag{4.42}$$

Similar to the `setup-inclusive` in the previous section, we find that MiNNLO$_{\rm PS}$ and NNLOPS predictions are compatible with each other. However, as already observed for the fiducial cross section, there is a $\sim 2.5\%$ shift in the central value that is related to the applied jet veto. While this shift is covered by the nominal scale uncertainties with just touching bands, we stress again that a more realistic assessment of the actual uncertainties related to the jet veto would further improve the compatibility between MiNNLO$_{\rm PS}$ and NNLOPS results. Indeed, in terms of shapes of the distributions they are already remarkably close. Also in the case with fiducial cuts, an important observation is that the inclusion of NNLO corrections on top of the MiNLO′ is crucial not only for the correct normalization, but for many observables to capture relevant shape effects. Moreover, the NNLO-accurate predictions are substantially more precise due to their strongly reduced uncertainty bands with respect to MiNLO′. We further notice that the impact of parton-shower emissions on observables with NNLO accuracy is quite moderate. Nevertheless, at phase-space boundaries where the fixed-order accuracy is reduced and the perturbative expansion breaks down due to effects from soft QCD radiation, the parton shower is absolutely crucial for a physical description. For instance, this can observed in the $p_{\mathrm{T},\ell\ell}$ distribution, where we have added the fixed-order NNLO prediction for comparison. Since the $p_{\mathrm{T,miss}} > 20\,\mathrm{GeV}$ requirement in `fiducial-1-JV` setup translates directly into a $p_{\mathrm{T},\ell\ell} > 20\,\mathrm{GeV}$ cut at LO, where the two leptons are back-to-back with the two neutrinos, the region $p_{\mathrm{T},\ell\ell} \leq 20\,\mathrm{GeV}$ is filled only upon inclusion of higher-order corrections and is effectively only NLO accurate. As a result, the boundary region $p_{\mathrm{T},\ell\ell} \sim 20\,\mathrm{GeV}$ becomes sensitive to soft-gluon effects that induce large logarithmic corrections and a perturbative instability [271] at $p_{\mathrm{T},\ell\ell} = 20\,\mathrm{GeV}$ in the fixed-order NNLO prediction. This unphysical behaviour is cured through the matching to the parton shower in the MiNNLO$_{\rm PS}$ and NNLOPS calculations.

It is clear that our new MiNNLO$_{\rm PS}$ predictions compare very well with the previous NNLOPS results, and that the two tools can be used equivalently to produce W^+W^-

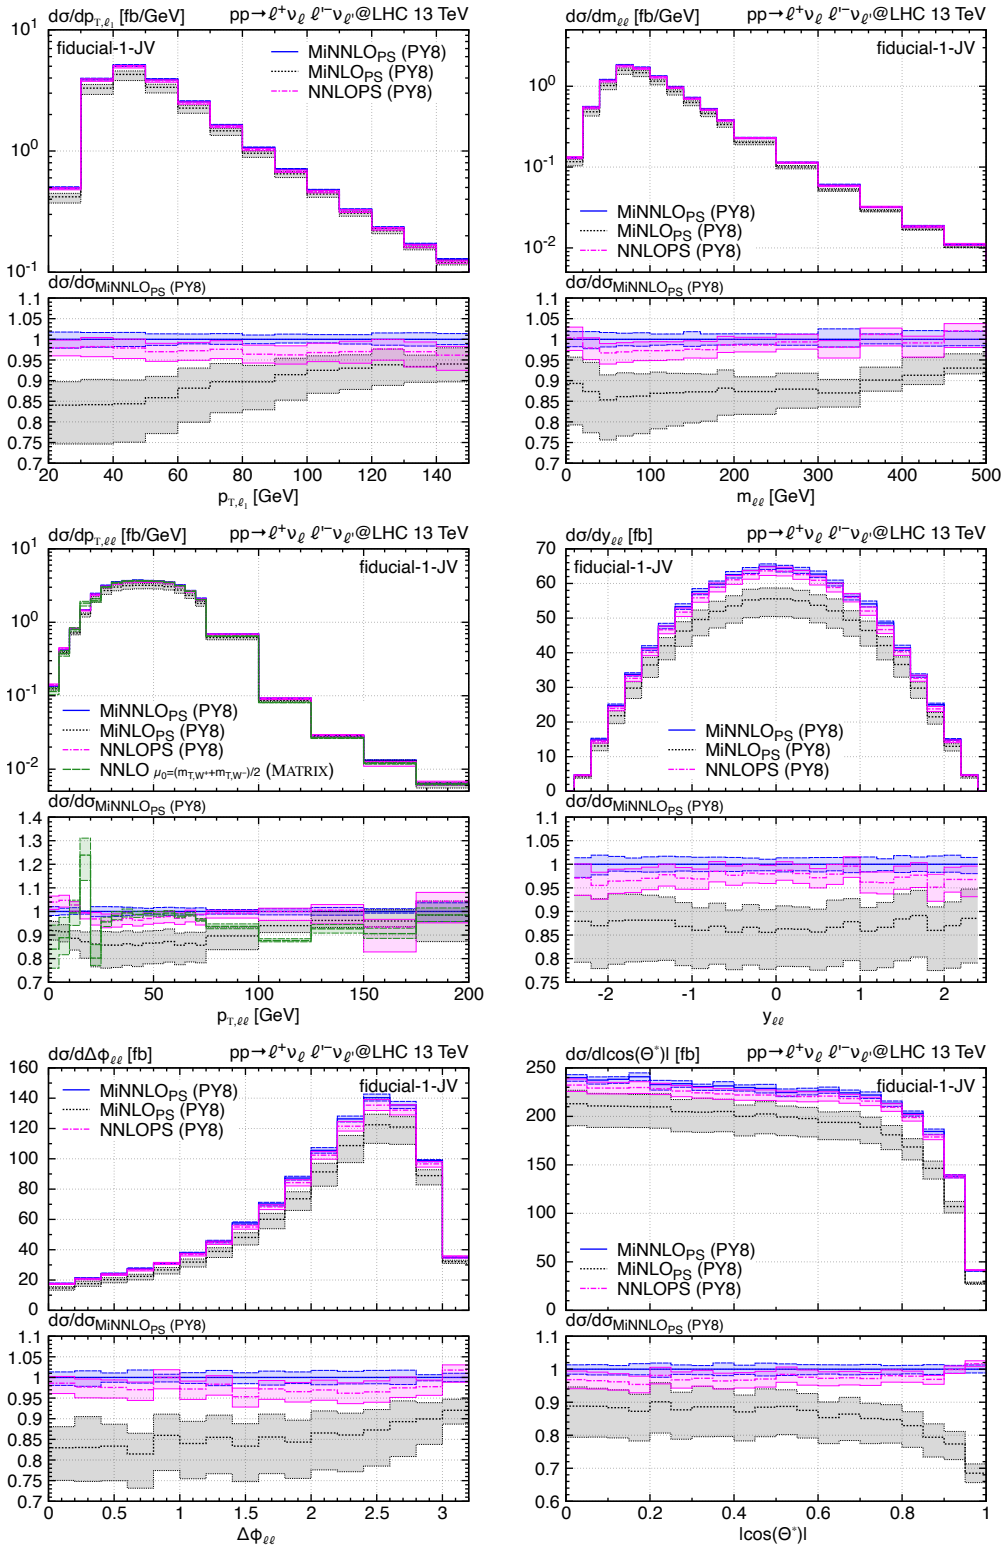

Figure 4.23: Differential distributions in the `fiducial-1-JV` phase space.

cross sections and distributions at NNLO accuracy matched to parton showers. This is also an indication of the robustness of NNLO+PS predictions for observables that are genuinely NNLO accurate. Given the limitation of the NNLOPS calculation regarding the necessity of multi-dimensional reweighting, the advantage of the MiNNLO$_{\rm PS}$ generator is that those results can be obtained directly at the level of the event generation. However, in the few phase-space regions where differences between the two calculations can be observed, those differences indicate relevant corrections beyond NNLO accuracy. Since plain 7-point scale variations do not always cover those discrepancies, they should be regarded as a perturbative uncertainty.

In the next section we will move to observables that are subject to large logarithmic corrections and where differences between the MiNNLO$_{\rm PS}$ and NNLOPS generator are larger. Thus, their assessment as an uncertainty becomes particularly important.

4.3.5 Observables sensitive to soft-gluon effects

In this section, we study distributions sensitive to soft-gluon radiation that require the inclusion of large logarithmic corrections to all orders in QCD perturbation theory either through a parton shower or through analytic resummation. In order to compare parton shower numerical resummation against the analytic one, we additionally include in some plots of this section curves for FO predictions matched to analytic resummation.

We begin the discussion with the results of Figure 4.24, where we study the transverse-momentum spectrum of the W^+W^- pair ($p_{\rm T,WW}$) in the `fiducial-1-noJV` phase space. For these plots, predictions for NNLO+N^{3}LL and NNLO+NNLL are reported with a green, double-dash dotted and a brown, dash-double dotted line, respectively. We refrain from showing the corresponding distribution in `setup-inclusive` and within the `fiducial-2-noJV` phase space, as we found them to be almost identical concerning the relative behaviour of the various predictions. At small values of $p_{\rm T,WW}$, large logarithmic contributions break the validity of the expansion in the strong coupling constant at a given fixed order, which requires their inclusion at all orders in perturbation theory either through a parton shower or through an analytic resummation. The left figure displays the region $0 \leq p_{\rm T,WW} \leq 50\,\text{GeV}$ and, indeed, the NNLO prediction, which is shown in the main frame only, becomes unphysical for small values of $p_{\rm T,WW}$. If we compare MiNNLO$_{\rm PS}$ and NNLOPS results in that region, we observe differences of about -10% to $+5\%$. By and large, those are covered by the respective uncertainty bands. However, it is clear (and expected) that for such an observable, which is sensitive to infrared physics, the differences between the two calculations become larger. In particular, both predictions are only NLO accurate in the tail of the $p_{\rm T,WW}$ distribution and at small transverse momenta the parton shower limits the accuracy of the calculation effectively to leading-logarithmic (LL) or partial (i.e. at leading colour) NLL accuracy. Therefore, differences of the order of those that we observe between MiNNLO$_{\rm PS}$ and NNLOPS are understood. Also, the comparison against the high-accuracy analytic resummation results at NNLO+N^{3}LL and NNLO+NNLL is quite good, which also agree within -10% to $+5\%$ with the MiNNLO$_{\rm PS}$ prediction for $p_{\rm T,WW} < 20\,\text{GeV}$ and are even fully compatible in the intermediate region up to 50 GeV. The resummed predictions do not favour either

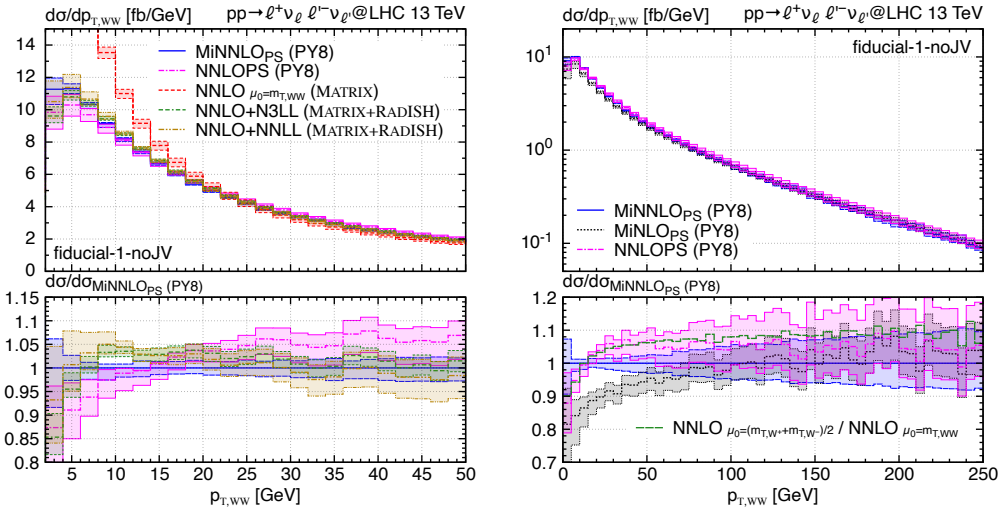

Figure 4.24: Distribution in the transverse momentum of the W^+W^- pair in the `fiducial-1-noJV` phase space, showing a smaller (left plot) and a wider range (right plot).

MiNNLO$_{\text{PS}}$ or NNLOPS results, but rather show similar differences from the two. On the other hand, the agreement is actually quite remarkable considering the fact that the region $p_{\text{T},WW} < 20\,\text{GeV}$ is entirely described by the substantially less accurate parton shower. Given the fact that for some bins the NNLO+N^{3}LL and NNLO+NNLL predictions are outside the uncertainty bands of the NNLO+PS accurate predictions though, the estimated uncertainties from μ_{R} and μ_{F} variations appear insufficient to reflect the actual size of uncertainties and one should consider additional handles to better assess the uncertainties of the parton shower at small $p_{\text{T},WW}$. Indeed, the NNLL prediction has a much larger uncertainty band in this region (induced by the variation of Q_{res}) even though it is more accurate.

In the right plot of Figure 4.24, we show the range $0 \leq p_{\text{T},WW} \leq 250\,\text{GeV}$. In the tail of the distribution, MiNNLO$_{\text{PS}}$ and NNLOPS (as well as MiNLO$'$) predictions are in perfect agreement with fully overlapping uncertainty bands. In the lower frame we show an additional curve given by the ratio of the central fixed-order NNLO prediction with $\mu_0 = \frac{1}{2}\,(m_{\text{T},W^+} + m_{\text{T},W^-})$ to the one with $\mu_0 = m_{\text{T},WW}$. It is very interesting to observe that the ratio corresponds almost exactly to the NNLOPS/MiNNLO$_{\text{PS}}$ ratio at smaller $p_{\text{T},WW}$. We recall that $\mu_0 = \frac{1}{2}\,(m_{\text{T},W^+} + m_{\text{T},W^-})$ is the scale used in the reweighting of the NNLOPS prediction, while $\mu_0 = m_{\text{T},WW}$ is somewhat more similar to the one within the MiNNLO$_{\text{PS}}$ approach. This suggests that the differences originating from terms beyond accuracy at small $p_{\text{T},WW}$ between the MiNNLO$_{\text{PS}}$ and NNLOPS are predominantly induced by the different scale settings in the two calculations. In fact, for any distribution (of the various ones we considered) where the NNLOPS/MiNNLO$_{\text{PS}}$ ratio becomes larger than a couple of percent, we observe that the corresponding ratio of fixed-order NNLO

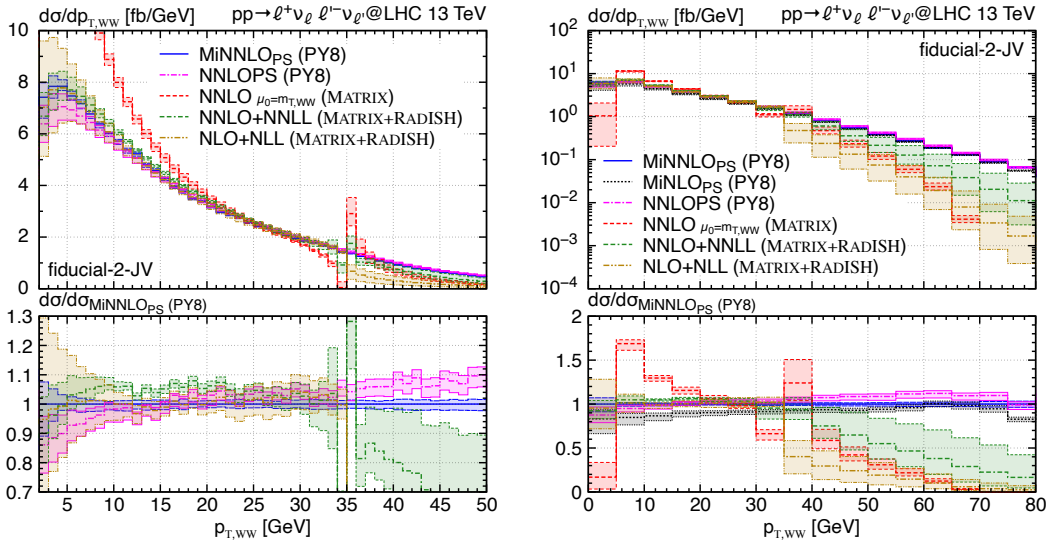

Figure 4.25: Distribution in the transverse momentum of the W^+W^- pair in the presence of a jet veto (`fiducial-2-JV` phase space), showing a smaller (left plot) and a wider range (right plot).

predictions is either very similar or even larger.

In Figure 4.25 we consider the W^+W^- transverse-momentum spectrum in the presence of a jet veto of $p_{\mathrm{T},j_1}^{\mathrm{veto}} = 35\,\mathrm{GeV}$ using the `fiducial-2-JV` setup. Here we have also included curves for NNLO+NNLL (green, double-dash dotted) and NLO+NLL (brown, dash-double dotted) calculation. The relative behaviour between the MiNNLO$_{\mathrm{PS}}$, NNLOPS, NNLO+NNLL and NLO+NLL results at small transverse momenta is relatively similar to the one observed for the $p_{\mathrm{T},WW}$ distribution without jet veto in setup `fiducial-1-noJV`. One main difference is that for this observable, which requires double differential resummation in $p_{\mathrm{T},WW}$ and p_{T,j_1}, the analytically resummed results are less accurate and therefore feature larger uncertainty bands, rendering them more compatible with the showered results. Indeed, the NLL uncertainty band is strongly increased at small $p_{\mathrm{T},WW}$ and much larger than the NNLO+PS ones, which, as argued before, also points to the fact that the scale uncertainties of the latter are somewhat underestimated, given that the parton shower is less accurate than the NLL calculation in that region. Another interesting region for this observable is around $p_{\mathrm{T},WW}$ values of 35 GeV, i.e. of the order of the jet-veto cut. The region $p_{\mathrm{T},WW} \geq p_{\mathrm{T},j_1}^{\mathrm{veto}}$ is filled for the first time at NNLO, which is effectively only LO accurate, since at LO one has $p_{\mathrm{T},WW} = 0$ and at NLO $p_{\mathrm{T},WW} = p_{\mathrm{T},j_1}$. Therefore, large logarithmic contributions challenge the perturbative expansion around $p_{\mathrm{T},WW} = p_{\mathrm{T},j_1}^{\mathrm{veto}}$ and the fixed-order NNLO prediction develops a perturbative instability, as visible in the main frame of the left plot in Figure 4.25. This instability is partially cured by the analytic resummation approach, which resums Sudakov logarithms in the limit where $p_{\mathrm{T},WW}$ and p_{T,j_1} are much smaller than the hard

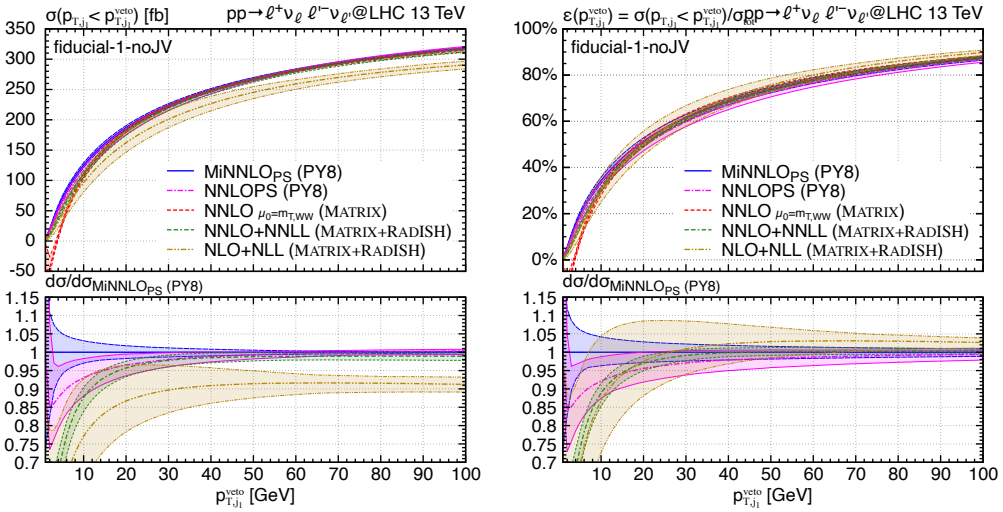

Figure 4.26: Jet-vetoed cross section (left plot) and jet-veto efficiency (right plot) in the `fiducial-1-noJV` phase space.

scale, but not all logarithmic contributions of the form $\log(p_{\mathrm{T},WW} - p_{\mathrm{T},j_1}^{\mathrm{veto}})$, which would require additional resummation when one or more hard jets are present. By contrast, the NNLO+PS calculations cure this instability entirely as they resum all relevant classes of logarithms (although with limited accuracy). Therefore, the MiNNLO$_{\mathrm{PS}}$ and NNLOPS calculations provide a more physical prediction at and above threshold, while below the threshold they are in good agreement with the analytically resummed predictions. If we look at the region above threshold in the right plot of Figure 4.25, we notice that the NNLO result drops substantially for $p_{\mathrm{T},WW}$ values above $p_{\mathrm{T},j_1}^{\mathrm{veto}}$, and also the NNLO+NNLL prediction is only slightly larger. Hence, this region of phase space is almost exclusively filled by the parton shower. Consequently, the transverse-momentum spectrum of a colour singlet in presence of a jet veto could be a good observable to tune the parton shower in experimental analyses.

In Figure 4.26 we study the jet-vetoed cross section as a function of the jet-veto cut $p_{\mathrm{T},j_1}^{\mathrm{veto}}$, which is defined as

$$\sigma(p_{\mathrm{T},j_1} < p_{\mathrm{T},j_1}^{\mathrm{veto}}) = \int_0^{p_{\mathrm{T},j_1}^{\mathrm{veto}}} \mathrm{d}p_{\mathrm{T},j_1} \, \frac{\mathrm{d}\sigma}{\mathrm{d}p_{\mathrm{T},j_1}} \,, \tag{4.43}$$

and the jet-veto efficiency given by

$$\varepsilon(p_{\mathrm{T},j_1}^{\mathrm{veto}}) = \sigma(p_{\mathrm{T},j_1} < p_{\mathrm{T},j_1}^{\mathrm{veto}})/\sigma_{\mathrm{int}} \,, \tag{4.44}$$

where σ_{int} is the integrated cross section in the `fiducial-1-noJV` phase space. Again the results for `setup-inclusive` and `fiducial-2-noJV` are very similar and are not shown. The interesting region is at small jet-veto cuts, where the validity of the perturbative expansion is broken by large logarithmic contributions in $p_{\mathrm{T},j_1}^{\mathrm{veto}}$, while for larger values

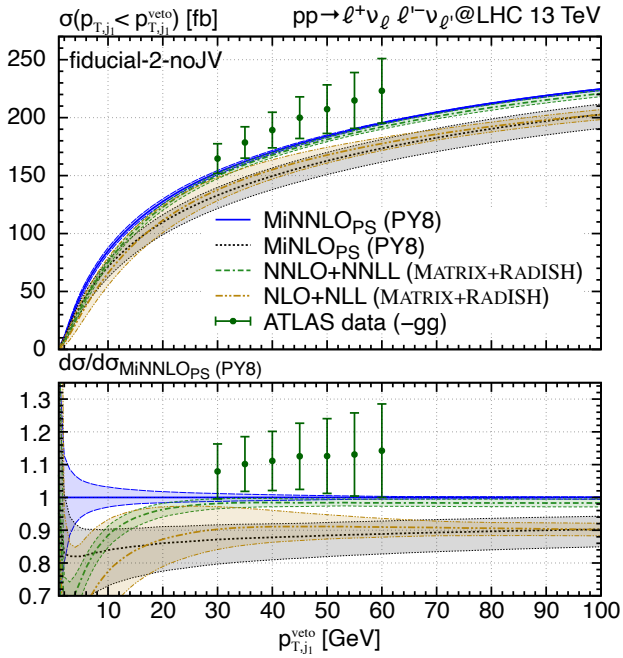

Figure 4.27: Jet-vetoed cross section in the `fiducial-2-noJV` phase space compared to data. As described in the caption of Table 4.7 the data has been adjusted by subtracting the ggLO contribution quoted in table 2 of Ref. [332] and by dividing out a factor of two.

the results tend towards their respective integrated cross sections. As it can be seen from the main frame, in the low $p_{\mathrm{T},j_1}^{\mathrm{veto}}$ region the pure fixed-order result at NNLO becomes indeed unphysical and turns actually negative. When comparing MiNNLO$_{\mathrm{PS}}$ and NNLOPS predictions, we find them to be in reasonable agreement within their respective uncertainties, with the NNLOPS one tending a bit faster towards zero for $p_{\mathrm{T},j_1}^{\mathrm{veto}} \lesssim 20\,\mathrm{GeV}$. In that region, the resummed NNLO+NNLL and NLO+NLL results tend even faster towards zero, with the NNLO+NNLL curve being about 20% below the MiNNLO$_{\mathrm{PS}}$ one at $p_{\mathrm{T},j_1}^{\mathrm{veto}} = 5\,\mathrm{GeV}$. This region is dominated by the parton shower, which resums only the LL (partial NLL) contributions. Clearly, the actual uncertainties in the NNLO+PS calculations are not covered by plain μ_{R} and μ_{F} variations. As argued for the $p_{\mathrm{T},WW}$ spectrum, additional handles would need to be considered to better assess the parton-shower uncertainties for very small $p_{\mathrm{T},j_1}^{\mathrm{veto}}$ cuts. Indeed, the NLL result features much wider uncertainties, despite being more (similarly) accurate in that region of phase space. However, we stress that such small $p_{\mathrm{T},j_1}^{\mathrm{veto}}$ cuts are usually not relevant for experimental W^+W^- analyses. Moreover, as pointed out before, there have been suggestions to include more conservative uncertainty estimates for jet-vetoed predictions [333, 334]. Since those effects are currently not accessed by any W^+W^- measurement, we have not performed further studies in that direction here. For instance, looking at

the fiducial phase-space definitions of Refs. [328, 329] that have been considered in this work, jet-veto cuts of $p_{\mathrm{T},j_1}^{\mathrm{veto}} = 25\,\mathrm{GeV}$, 30 GeV and 35 GeV are used. For those values, MiNNLO$_{\mathrm{PS}}$ predictions are in good agreement with the NNLO+NNLL resummation, and even down to $p_{\mathrm{T},j_1}^{\mathrm{veto}} \sim 20\,\mathrm{GeV}$ they differ by less than 5% with overlapping uncertainties.

When comparing the predicted jet-vetoed cross section as a function of $p_{\mathrm{T},j_1}^{\mathrm{veto}}$ in the `fiducial-2-noJV` setup against data in Figure 4.27, it is clear that the MiNNLO$_{\mathrm{PS}}$ and the NNLO+NNLL predictions are compatible in the relevant region. The agreement with data is good in either case, with the data points either marginally overlapping within one standard deviation or being just outside this range. One should bear in mind however that the normalization of the theory predictions can be increased by $\sim 5\%$ just by using a different PDF set, which yields even better agreement with data, as shown in Ref. [272]. Apart from that, it is clear that the inclusion of NNLO corrections brings the theory predictions closer to data.

4.4 Reaching nNNLO+PS accuracy for ZZ production

Despite the fact that the cross section for ZZ production is smaller than the one for $W^{\pm}Z$ or W^+W^- production, the decay to four charged leptons experimentally provides the cleanest signature among massive diboson processes, since the final state does not involve any particle eluding the detectors. Accordingly, experimental measurements for this signal have already reached a remarkable level of precision. In particular, both ATLAS and CMS collaborations have performed many measurements of ZZ production, the latest of which at 13 TeV discussed in Ref. [336–341]. Those measurements are extremely important, as for the other vector-boson pair production processes, to test the SM and to constrain triple-gauge couplings. Moreover, ZZ is also an irreducible background to on- and off-shell Higgs cross-section measurements, when the Higgs boson decays to four charged leptons. These latter measurements are particularly relevant for the extraction of the Higgs couplings and for constraining the Higgs width (see for instance [342, 343] for some recent analyses). But the interest in these measurements is also motivated by the possibility to disclose new physics scenarios (see for instance Ref. [344]).

In this section we describe how we include NNLO QCD corrections for ZZ production in the quark-initiated channel using the MiNNLO$_{\mathrm{PS}}$ method. Additionally, we also present a NLO+PS Powheg calculation for the loop-induced $gg \to ZZ$ process. Despite the fact that the latter channel starts contributing only at $\mathcal{O}(\alpha_s^2)$, so that its leading order contribution formally enters only at NNLO, it is enhanced by the gluon PDFs. Indeed, we have already commented in the previous section how the missing inclusion of this contribution prevented us to compare the prediction of our W^+W^- generator against data. Also for ZZ production, a fair data-theory comparison requires to account for this channel. When the two contributions are combined, as presented in Ref. [220], these results become the most advanced theoretical predictions for ZZ production at the LHC, since they include the highest perturbative accuracy in QCD available to date.

Our discussion starts from considering the process

$$pp \to \ell^+\ell^-\ell^{(\prime)+}\ell^{(\prime)-} \tag{4.45}$$

for any combination of charged leptons $\ell, \ell' \in \{e, \mu, \tau\}$. While at the matrix-element level our calculation is based on the different-flavour channel $\ell \neq \ell'$, at the event-generation level arbitrary combinations of charged leptons can be considered, both with different flavours $\ell \neq \ell'$ and with same flavours $\ell = \ell'$ (in the latter case interference effects when exchanging the charged leptons, which are typically at the 1-2% level [345], are neglected). Moreover, lepton masses are accomodated via reshuffling of the momenta in the event generation. For simplicity and without loss of generality we consider only the process $pp \to e^+e^-\mu^+\mu^-$ here, which we will refer to as ZZ production in the following. By including all resonant and non-resonant topologies leading to this process, off-shell effects, interferences and spin correlations are taken into account. Sample diagrams are shown in Figure 4.28 and they include:

(a) tree-level double-resonant t-channel ZZ production in the $q\bar{q}$ channel;

(b) tree-level single-resonant s-channel Drell–Yan topologies in the $q\bar{q}$ channel;

(c) loop-induced ZZ production in the gg channel.

The loop-induced gg contribution, including also the single-resonant Higgs mediated diagrams, proceeds through a quark loop and, as we said, enters the cross section at $\mathcal{O}(\alpha_s^2)$, i.e. it is part of the NNLO QCD corrections. Since this contribution is enhanced by the large gluon-gluon luminosity at LHC energies, it yields a relatively large fraction of the NNLO corrections [310, 311]. Moreover, it is known that at NLO QCD [346, 347], i.e. $\mathcal{O}(\alpha_s^3)$, its relative correction is very sizable (about a factor of two). It is likely that these corrections constitute the most significant contribution to ZZ production at $\mathcal{O}(\alpha_s^3)$, since the $\mathcal{O}(\alpha_s^3)$ corrections to the $q\bar{q}$ channel are not expected to be of the same size as those at the previous perturbative order.

For our results we include the most accurate currently available information in QCD perturbation theory for both the $q\bar{q}$-initiated and the loop-induced gg-initiated processes, and match them consistently with a parton shower. Thus, we calculate NNLO+PS predictions in the $q\bar{q}$ channel by means of the MiNNLO$_{\text{PS}}$ method, described in Section 4.4.1, and NLO+PS predictions in the loop-induced gg channel using the POWHEG approach, as

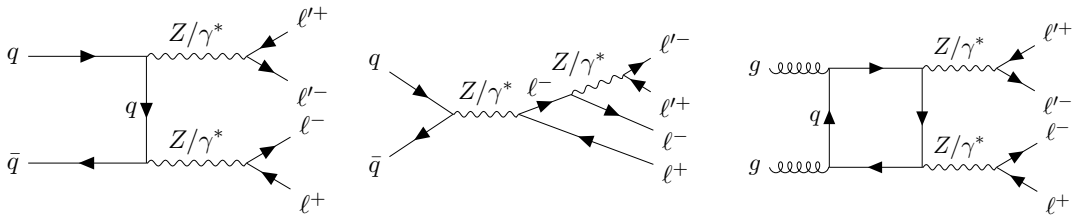

Figure 4.28: Sample Feynman diagrams for ZZ production with four charged leptons in the final state. Panels (a) and (b): tree-level diagrams of the quark annihilation ($q\bar{q}$) channel; Panel (c): loop-induced diagram in the gluon-fusion (gg) channel.

presented in Section 4.4.2. Full top-quark mass effects are included everywhere, except for the two-loop amplitudes. For the $q\bar{q}$ channel they are expected to be small in the relevant phase-space regions, while for the loop-induced gg contribution we have included them approximately via reweighting in the two-loop correction (cf. section 4.4.2). We stress that to avoid any overlap, our MiNNLO$_{\text{PS}}$ implementation of the NNLO+PS calculation in the $q\bar{q}$ channel does not include the loop-induced gg-initiated contribution. In this way, all loop-induced gg contributions are correctly accounted for when combining the former with our NLO+PS predictions for the gg channel. The possibility to incoherently combine the two calculations relies on the fact that NLO corrections to the loop-induced gluon-fusion channel can be formulated as a gauge-invariant set of contributions. We dubbed the level of accuracy reached by our results nNNLO+PS, as the NLO corrections to the gg channel are of $\mathcal{O}(\alpha_s^3)$. These corrections constitute the most significant part of the N^{3}LO calculation, as pointed out before.

4.4.1 MiNNLO$_{\text{PS}}$ for $q\bar{q} \to ZZ$ production

In this section we present the implementation of a NNLO+PS generator within the Powheg-Box-Res framework for ZZ production in the $q\bar{q}$ channel by means of the MiNNLO$_{\text{PS}}$ method. Our NLO+PS generator for the loop-induced gg channel is discussed in the next section. We stress that, while we distinguish these processes as $q\bar{q}$ and gg, in their higher-order corrections of course all the relevant partonic initial states are consistently included, with the exception of the gg loop-induced partonic channel in the NNLO corrections to the $q\bar{q}$ process, since this contribution is already accounted for by our NLO+PS generator for the loop-induced gg channel.

Since no implementation for ZZ+jet production was available in Powheg-Box to date, the first step was to implement this process in the Powheg-Box-Res framework. We have implemented all relevant flavour channels and, in addition, adapted the routine `build_resonance_histories` of Powheg-Box-Res such that it is capable of automatically constructing the correct resonance histories. The tree-level single and double real matrix elements for $e^+e^-\mu^+\mu^-$+1, 2-jet production and the one-loop amplitude for $e^+e^-\mu^+\mu^-$+1-jet production are evaluated through OpenLoops [245–247].

In a second step, we have employed the MiNNLO$_{\text{PS}}$ method to obtain NNLO+PS predictions for ZZ production in the $q\bar{q}$ channel. In particular, we made use of the implementation of the MiNNLO$_{\text{PS}}$ method that was developed for the first time for $Z\gamma$ production, as described in Sections 3.3.1 and 4.1. The respective tree-level and one-loop $q\bar{q} \to e^+e^-\mu^+\mu^-$ amplitudes are also evaluated through OpenLoops, while the two-loop helicity amplitudes have been obtained by extending the interface to Matrix [81] developed in Ref. [217] to ZZ production. The evaluation of the two-loop coefficients in this implementation relies on the code VVamp [316] and is based on the calculation of Ref. [315].

As discussed in Section 4.3.1 for W^+W^- production, the evaluation of the two-loop helicity amplitudes for massive diboson processes is particularly demanding from a computational point of view. In Section 4.3.1 we circumvented the problem by constructing a set of interpolation grids for the two-loop coefficients that achieves their fast on-the-fly

evaluation. For ZZ production we pursue a different strategy: we exploit the possibility of reweighting the events at the generation level (i.e. stage 4 in POWHEG-BOX) to include the two-loop contribution. In combination with a suitable caching system of the two-loop amplitude that we implemented, that allows us to omit the evaluation of the two-loop contribution entirely in the calculation up to stage 4, where it needs to be evaluated only once per event [17]. To be more precise, we have implemented a new flag (`run_mode`), which is accessible from the POWHEG input file, and allows the user to switch between four different ways of running the code (see also description in Appendix B.5). Either the full calculation is performed including the two-loop contributions throughout (`run_mode 1`), or one completely drops the NNLO corrections provided by MiNNLO$_{\rm PS}$, specifically the terms $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$ in Eq. (3.79), thus effectively reproducing MiNLO$'$ predictions (`run_mode 2`). Alternatively, the evaluation of the two-loop amplitude can be omitted only in the grid setup, i.e. stage 1 in POWHEG-BOX, (`run_mode 3`), or completely (`run_mode 4`). For all results presented in this work we run the code with the last option `run_mode 4`, i.e. without evaluating the computationally expensive two-loop amplitude. In this way, also the generation of the events is faster. However, once an event has been accepted, it is reweighted such that the two-loop contribution is included by resetting the `run_mode` keyword in the event reweight information of the POWHEG input file. As a result the two-loop amplitude is evaluated only once for each event, considerably improving the efficiency of the code. Moreover, following the same logic we can also compute MiNLO$'$ weights in parallel to the generation of MiNNLO$_{\rm PS}$ ones using the appropriate setting for `run_mode` in the event reweight information. We have first validated our implementation in an inclusive setup, requiring only a suitable Z-mass window for the opposite-charge same-flavour dilepton pairs. Here we compared the inclusive cross section at the LHE level obtained at stage 4 with the one computed at stage 2 when including the two-loop contribution, finding excellent agreement. Another very robust cross-check of the reweighting procedure is provided by the comparison of the MiNLO$'$ results obtained directly or through reweighting, which also agree perfectly.

For the evaluation of the several convolutions with the PDFs entering the calculation of $\mathcal{D}_{\ell_{\rm F}}(p_{\rm T})$, we employ HOPPET [225]. This is the standard package used in MiNNLO$_{\rm PS}$, as we have already discussed several times, together with the HPLOG package [265], for the computation of the polylogarithms entering the collinear coefficient functions.

Many of the MiNNLO$_{\rm PS}$ settings used to produce NNLO+PS accurate ZZ events are similar to what discussed for W^+W^- in Section 4.3. In this context, to avoid spurious contributions from higher-order logarithmic terms at large $p_{\rm T}$, we replace the logarithm $L = \ln Q/p_{\rm T}$, where we set $Q = m_{4\ell}$, with a modified logarithm $\tilde{L}$ which is identical to L below $p_{\rm T} = Q/2$ and smoothly vanishes at $p_{\rm T}$ equal or larger than Q, as it was introduced in Ref. [223]. As far as the renormalization and factorization scales are concerned, we use the standard MiNNLO$_{\rm PS}$ scale setting at small $p_{\rm T}$, while in the NLO ZZ+jet cross section at large $p_{\rm T}$ the scale setting is changed by activating the option `largeptscales 1`.

[17] Note that in order for the caching to work properly and not having to reevaluate the two-loop amplitude for every scale variation during the event reweighting, we have set the parameter `rwl_group_events 1` in the input file, which ensures that the events are reweighted one-by-one instead of in batches.

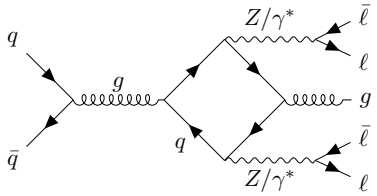

Figure 4.29: Sample Feynman diagram for a $q\bar{q}$-initiated regular contribution to the loop-induced gg channel at $\mathcal{O}(\alpha_s^3)$.

As done in Section 4.3, the Landau singularity is regulated by freezing the strong coupling and the PDFs for scales below 0.8 GeV. Finally, as recommended for processes with jets in the final state already at the LO, we turn on the option `doublefsr 1` of the POWHEG-BOX (see Ref. [263] for details). For the parton shower we have used the standard settings, also for the recoil scheme (namely a global recoil scheme for initial state radiation, with `SpaceShower:dipoleRecoil 0`).

4.4.2 NLO+PS for $gg \to ZZ$ production

As discussed before, the leading-order contribution to the loop-induced gluon-fusion process enters the ZZ cross section at $\mathcal{O}(\alpha_s^2)$. Thus, it constitutes a NNLO correction relative to the LO part of the $q\bar{q}$ channel, but it is significantly enhanced by the large gluon-gluon luminosities. It is therefore mandatory to include also the NLO corrections to the loop-induced gluon-fusion contribution in any precision study of ZZ production that compares theory and data.

We have implemented a NLO+PS generator for loop-induced ZZ production in the gg channel within the POWHEG-BOX-RES framework. Note that in addition to continuum ZZ production as shown in Figure 4.28 (c) we also include the contribution mediated by a Higgs boson (or a single Z boson). The calculation of these loop-induced processes is effectively of similar complexity as a NNLO calculation, as far as the amplitude evaluation is concerned. Despite that, the matching to the parton shower through the POWHEG method, which is essentially automated in POWHEG-BOX-RES, can be applied to loop-induced processes as well, without any further conceptual issue. However, such a NLO calculation requires the evaluation of both one-loop and two-loop helicity amplitudes and the process at hand is numerically substantially more demanding than a tree-level one, since the evaluation time of the one-loop and two-loop amplitudes is much slower and the stability of the one-loop matrix elements with one emitted real parton is challenged in the infrared regions. To cope with these numerical issues, we have implemented and exploited a number of handles within POWHEG-BOX-RES, which will be discussed below.

For the implementation in POWHEG-BOX-RES, we have specified the relevant flavour channels and hard-coded also the resonance channels of the process, as the automatic determination of the latter via the already mentioned routine `build_resonance_histories` is not available yet for loop-induced processes. At NLO, all relevant partonic channels, namely gg, gq, qg, $q\bar{q}$ and the charge-conjugated ones, are included. To unambiguously

define the NLO corrections to the loop-induced gluon-fusion process for each of those initial states, we follow the criterium introduced in Ref. [347], which consists in including all diagrams that involve a closed fermion loop with at least one vector boson attached to it. The one-loop amplitudes with zero and one jet are evaluated through OpenLoops [245–247]. For this purpose, we have adapted the OpenLoops interface in Powheg-Box-Res developed in Ref. [179] to deal with loop-induced processes. As a cross-check, we have also interfaced Recola [18] [348, 349] to Powheg-Box-Res and found full agreement for all one-loop amplitudes. For the two-loop helicity $gg \to \ell^+\ell^-\ell^{(\prime)+}\ell^{(\prime)-}$ amplitudes we exploit their implementation within Matrix [81], which is based on the evaluation of the two-loop coefficients through VVamp [315] from their calculation in Ref. [350]. To this end, we have extended the interface of Powheg-Box-Res to Matrix developed in Ref. [217] to include the $gg \to \ell^+\ell^-\ell^{(\prime)+}\ell^{(\prime)-}$ two-loop amplitudes. When it comes to evaluating these two-loop coefficients through VVamp, the code can become once again very slow, with an average time evaluation per phase-space point lasting from a few seconds to several tens of seconds. Since this leads to a severe bottleneck in the calculation and especially in the event generation, we have implemented a caching system for the two-loop corrections and we include them only through event reweighting. This is very similar in spirit to the way the two-loop hard function is included in the MiNNLO$_{\text{PS}}$ generator in the $q\bar{q}$ channel, as described in the previous section. Our calculation includes the full top-quark mass effects, except for the two-loop $gg \to \ell^+\ell^-\ell^{(\prime)+}\ell^{(\prime)-}$ amplitudes, where they are not known to date [19]. Instead, we follow the same approach as Ref. [347] and include them approximately through a reweighting of the massless two-loop amplitude with the ratio of the one-loop result including massive loops to the one with only massless loops. Since here we are interested in the ZZ signal region, such reweighting is expected to work extremely well. In fact, Ref. [353] recently confirmed that using an asymptotic expansion in the top-quark mass leads to practically identical results as using such reweighting, as long as one sticks to the validity range of the expansion itself.

During the course of the preparation of this work, Ref. [353] presented a completely independent implementation of a NLO+PS generator for loop-induced ZZ production in the gg channel within the Powheg-Box-Res framework. We have compared our calculation to theirs both at the level of individual phase-space points and of the integrated cross sections, and we have found perfect agreement when applying the same approximation for the two-loop virtual corrections. Since, although developed independently, the two calculations are essentially interchangeable (both developed in Powheg-Box-Res using OpenLoops and VVamp), we advocate that it is equivalent to use either code and combine the results subsequently with our MiNNLO$_{\text{PS}}$ generator

[18] Recola is one more computer program for the evaluation of tree- and one-loop amplitudes at NLO QCD and EW using SM Feynman rules, but also supporting different SM extensions, including aTGC, EFT-operator contributions and so on. Different softwares for the automated generation and evaluation of one-loop amplitudes have been encountered so far in addition to Recola, such as OpenLoops (see foonote 2) or GoSam (see footnote 12): what distinguishes the various packages are the actual algorithms and the ideas on which they are based on in order to render the entire workflow as efficient and stable as possible.

[19] For the case of on-shell $gg \to ZZ$ production the full top-quark mass dependence was recently calculated in Refs. [351, 352].

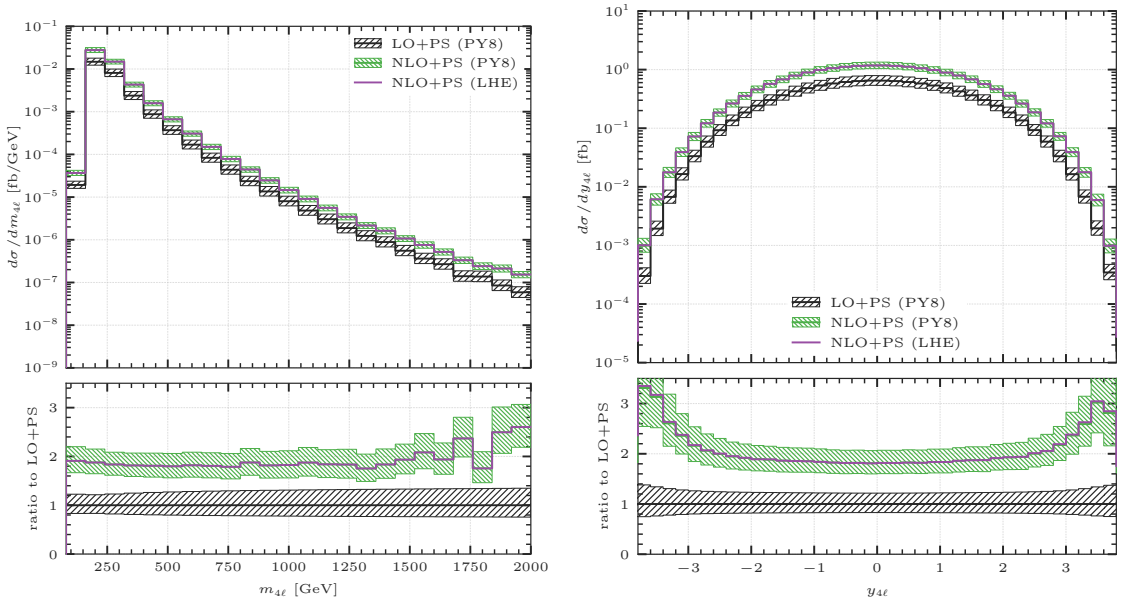

Figure 4.30: Predictions for ZZ production in the loop-induced gg channel at LO+PS and NLO+PS. For reference also the LHE-level central result at NLO is plotted. Shown are the distributions in the invariant mass and rapidity.

in the $q\bar{q}$ channel to obtain nNNLO+PS accurate results.

To better control the numerical stability of the calculation we have implemented settings similar to those reported in Ref. [353]: in particular we apply small (0.5 GeV) generation cuts on the transverse momentum of the four-lepton system and of each Z boson. Moreover, we exploit the stability system of OpenLoops and set the parameter `stability_kill2 0.01` to remove the remaining unstable points. We have further modified the code in such a way that, whenever the real-emission contribution is set to zero by one of the previous stability checks, also the respective counterterms entering the FKS subtraction are set to zero. Finally, we use `withdamp 0` in the input card of the process in order not to split the real cross section into a singular and a remnant contribution (as defined in Eq. (3.11)) as the considerably small value of the latter leads to numerical issues when generating events. The same is true for the regular contribution (introduced in Section 3.1) that contains only the $q\bar{q}$ channel (see Figure 4.29 for a sample diagram): after verifying that it is completely negligible, we have turned it off for all results presented in this work.

Since in the upcoming section we study phenomenological results for the full $pp \to e^+e^-\mu^+\mu^-$ process, we show some plots for the loop-induced gg channel separately in Figure 4.30 and Figure 4.31, both at LO and at NLO. The settings and inputs that we use here correspond to those introduced in Section 4.4.3 in the inclusive setup (`setup-inclusive`) with just a Z-mass window applied between 60 GeV and 120 GeV. The renormalization and factorization scales are set to $\mu_R = \mu_F = \sqrt{m_{4\ell}^2 + p_{T,4\ell}^2}$, where

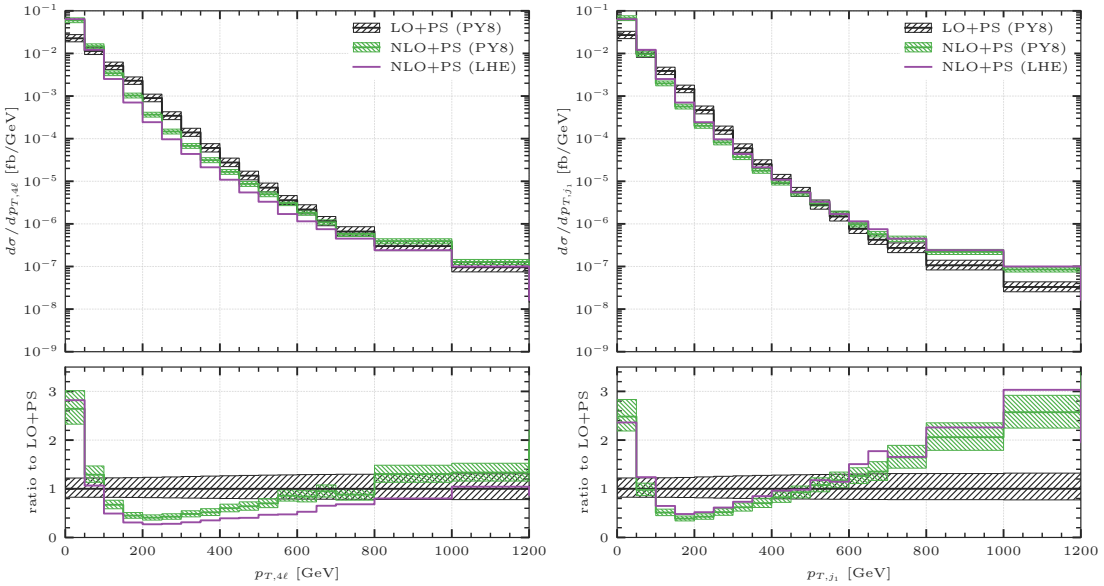

Figure 4.31: Predictions for ZZ production in the loop-induced gg channel at LO+PS and NLO+PS. For reference also the LHE-level central result at NLO is plotted. Shown are the distributions in the transverse momentum of the four-lepton system and in the transverse momentum of the leading jet.

$m_{4\ell}$ and $p_{\mathrm{T},4\ell}$ are the invariant mass and the transverse momentum of the four-lepton system, respectively. Furthermore, the uncertainty bands are obtained through a standard seven-point scale variation, and we employ the PYTHIA8 parton shower [124] with the A14 tune [266] (see Section 4.4.3 for further details) to obtain the showered results presented in Figure 4.30 and Figure 4.31. For the genuine NLO-accurate quantities shown in Figure 4.30, namely $m_{4\ell}$ and the rapidity of the four-lepton system ($y_{4\ell}$), we find results that are completely in line with the findings of previous fixed-order calculations [346, 347]: this is totally expected, since shower effects are negligible for those observables, as one can see from the LHE results. In particular, NLO corrections are sizable and increase the value of the inclusive cross section by almost a factor of two, with scale uncertainties at the level of 10-15%. In certain phase-space regions, like in the tail of the $m_{4\ell}$ distribution, the NLO corrections can even become significantly larger than a factor of two. However, in those regions the relative impact of the loop-induced gg contribution is reduced. When looking at the transverse-momentum spectrum of the four-lepton system ($p_{\mathrm{T},4\ell}$) and of the leading jet (p_{T,j_1}) in Figure 4.31, the importance of matching to the parton shower becomes clear [20]: at LO only the parton shower fills those distributions and at NLO it still provides a substantial correction. In fact, in a fixed-order calculation both observables would diverge, and therefore be unphysical, at small transverse momenta. It is interesting to notice that the LO+PS result is actually

[20] Note that, compared to the LO+PS results quoted for $gg \to ZZ$, predictions at higher accuracy in the presence of an additional radiated jet, possibly including zero- and one-jet merging, have been presented in Refs. [354–356].

higher than the NLO+PS one in the intermediate $p_{\mathrm{T},4\ell}$ region before it falls off steeply. This region is completely filled by the shower, whose starting scale by default is set to $m_{4\ell}$ in the LO calculation. The fact that $m_{4\ell}$ is on average relatively large explains why the shower fills the spectrum even at such high transverse momenta.

4.4.3 Setup and cross section results

We now turn to presenting phenomenological results derived in the framework that we have described so far for ZZ production. We start discussing the setup that we used for the predictions reported in this section for integrated cross sections; the same setup is also employed for the other differential predictions of the following section.

We consider proton–proton collisions at the LHC with a center-of-mass energy of 13 TeV and we specifically present predictions for $pp \to e^+e^-\mu^+\mu^-$ production. We use the complex-mass scheme [254] throughout and set EW inputs to their PDG [2] values: $G_F = 1.16639 \times 10^{-5}$ GeV^{-2}, $m_W = 80.385$ GeV, $\Gamma_W = 2.0854$ GeV, $m_Z = 91.1876$ GeV, $\Gamma_Z = 2.4952$ GeV, $m_H = 125$ GeV and $\Gamma_H = 0.00407$ GeV. We set the on-shell top-quark mass to $m_t = 173.2$ GeV, and $\Gamma_t = 1.347878$ GeV is used. We determine the other EW parameters in the G_μ scheme with the EW coupling $\alpha_{G_\mu} = \sqrt{2}/\pi G_\mu |(m_W^2 - i\Gamma_W m_W)\sin^2\theta_W|$ and the EW mixing angle $\cos^2\theta_W = (m_W^2 - i\Gamma_W m_W)/(m_Z^2 - i\Gamma_Z m_Z)$. We use the NNPDF3.1 [357] NNLO set with $\alpha_s = 0.118$ via the LHAPDF interface [224] for all our predictions. For MiNLO′ and MiNNLO$_{\mathrm{PS}}$, the PDFs are read by LHAPDF and evolved internally by HOPPET [225], as we already described for the other diboson processes presented in previous sections. The central factorization and renormalization scales are set as discussed in Section 3.3.2 for the MiNNLO$_{\mathrm{PS}}$ ZZ generator in the $q\bar{q}$ channel and as given in Section 4.4.2 for the loop-induced gg channel. Scale uncertainties are estimated by varying μ_{F} and μ_{R} around their central values by a factor of two in each direction, while keeping the minimal and maximal values with the constraint $0.5 \leq \mu_{\mathrm{R}}/\mu_{\mathrm{F}} \leq 2$, according to the usual seven-point scale variation.

By combining the MiNNLO$_{\mathrm{PS}}$ $q\bar{q}$ results and loop-induced gg results at (N)LO+PS, we obtain predictions for ZZ production at (n)NNLO accuracy matched to parton showers. For all (n)NNLO+PS predictions presented in what follows we make use of the Pythia8 parton shower [124] with the A14 tune [266] (`py8tune 21` in the input card). To validate

	`setup-inclusive`	`setup-fiducial`
Z-mass window	$60\,\mathrm{GeV} < m_{Z_1}, m_{Z_2} < 120\,\mathrm{GeV}$	$60\,\mathrm{GeV} < m_{Z_1}, m_{Z_2} < 120\,\mathrm{GeV}$
lepton cuts	$m_{\ell^+\ell^-} > 4\,\mathrm{GeV}$	$p_{\mathrm{T},\ell_1} > 20\,\mathrm{GeV}$, $p_{\mathrm{T},\ell_2} > 10\,\mathrm{GeV}$, $p_{\mathrm{T},\ell_{3,4}} > 5\,\mathrm{GeV}$, $\lvert\eta_\ell\rvert < 2.5$, $m_{\ell^+\ell^-} > 4\,\mathrm{GeV}$

Table 4.8: Inclusive and fiducial cuts used to the define the `setup-inclusive` and `setup-fiducial` phase-space regions [341]. See text for more details.

our calculation and to show where shower effects are crucial, we compare (n)NNLO+PS predictions obtained with MiNNLO$_{\rm PS}$ and (n)NNLO fixed-order predictions obtained with Matrix [81]. Additionally, we consider the inclusion of NLO EW effects. In the Matrix predictions we set $\mu_{\rm R} = \mu_{\rm F} = m_{4\ell}$, and we construct the scale-uncertainty bands with the same canonical seven-point scale variation used for our MiNLO$'$ and MiNNLO$_{\rm PS}$ results.

Moreover, we compare our predictions with the most recent results by the CMS collaboration [341] within the fiducial volume defined in Table 4.8, denoted as `setup-fiducial`. The reconstructed Z-bosons Z_1 and Z_2 are identified as follows: Z_1 corresponds to the opposite-sign same-flavour (OSSF) lepton pair with an invariant mass being the closest to the Z-boson mass, while Z_2 is the remaining OSSF lepton pair. Since here we only consider the different-flavour channel ($e^+e^-\mu^+\mu^-$), the two Z bosons are unambiguously reconstructed and this procedure only amounts to selecting which lepton pair is called Z_1 and which Z_2. Note that in the different-flavour channel the additional $m_{\ell^+\ell^-} > 4\,\text{GeV}$ cut in Table 4.8 has no effect, and just plays a role when considering a same-flavour channel for the process (for instance $e^+e^-e^+e^-$ or $\mu^+\mu^-\mu^+\mu^-$). Besides the fiducial setup, we also consider an inclusive setup (dubbed `setup-inclusive`), where we only require a Z-mass window between 60 GeV and 120 GeV for the two resonances.

In order to provide the most realistic comparison to experimental data, our final predictions include effects from hadronization and multi-particle interactions (MPI). We also include QED showering effects as provided by Pythia8. In order to prevent charged resonances to radiate photons and photons to branch into lepton- or quark-pairs, we set the two flags `TimeShower:QEDshowerByOther` and `TimeShower:QEDshowerByGamma` to `off`.

Finally, we define dressed leptons by adding to the four-momentum of a lepton the four-momenta of all photons within a distance $\Delta R_{\gamma\ell} = \sqrt{\Delta\phi^2_{\gamma\ell} + \Delta\eta^2_{\gamma\ell}} < 0.1$.

Integrated cross section

We start the discussion of our results by first considering integrated cross sections. In Table 4.9 we report predictions both in the inclusive and in the fiducial setup introduced above for various perturbative calculations. Specifically, we consider MiNLO$'$ predictions, and a number of predictions including NNLO corrections, both at fixed order and matched to parton showers through MiNNLO$_{\rm PS}$: besides the complete NNLO predictions (that include the LO loop-induced gg contribution), we provide the NNLO corrections to the $q\bar{q}$ channel (dubbed NNLO$_{q\bar{q}}$) and nNNLO cross sections (as defined before). For completeness, we also quote nNNLO predictions combined with NLO EW corrections, computed at fixed order with Matrix, either using an additive or multiplicative scheme (see discussion at the end of this section). In the latter predictions we also take into account the photon-induced contribution at LO and beyond [21]. In order to compare our predictions to fixed-order results, all MiNNLO$_{\rm PS}$ (and MiNLO$'$) results of Table 4.9

[21] We used the `NNPDF31_nnlo_as_0118_luxqed` [358–360] PDF set to compute fixed-order predictions which include EW corrections, and verified that the (n)NNLO prediction is modified at the few permille level with respect to the prediction obtained with `NNPDF31_nnlo_as_0118`.

are obtained at parton level, without including hadronization, MPI or photon radiation effects. We have checked explicitly that those effects have a negligible impact on the integrated cross sections.

The MiNNLO$_{\rm PS}$ prediction and the NNLO result are in excellent agreement with each other both in the inclusive and in the fiducial setup. The perturbative uncertainty at (n)NNLO(+PS) is at the 2-3% level. In particular, despite the fact that the loop-induced gg process at LO (NLO) contributes only $\sim$ 6-8% ($\sim$ 10-15%) to the NNLO (nNNLO) cross section, the uncertainties of the (n)NNLO results are dominated by the gluon-initiated contribution. The NLO correction for the loop-induced gg channel is particularly sizable, almost doubling the LO contribution entering at α_s^2, as discussed in Section 4.4.2. Accordingly, the nNNLO central prediction is not included in the NNLO uncertainty band.

The MiNLO$'$ result is 8-10% smaller than the MiNNLO$_{\rm PS}$ result. Its uncertainty band, which is considerably larger than the MiNNLO$_{\rm PS}$ one, does not contain the central (n)NNLO+PS prediction, because scale variations cannot account for the additional loop-induced gg process entering at NNLO. We also note that the MiNLO$'$ uncertainty band is larger than the NLO one, and it includes the NLO result. On the contrary, the NLO uncertainty band is very small and neither MiNLO$'$ nor the NNLO central results lie inside it.

Notwithstanding the excellent agreement between the nNNLO(+PS) result and the fiducial cross section measured by CMS, the theoretical predictions should be supplemented with EW corrections. Even though the main focus of this work has been to improve the QCD accuracy of our predictions in such a way to reach a precision competitive with the one of experimental measurements, when NNLO accuracy in QCD is achieved, the size of NLO EW corrections can become comparably important. Therefore, here and in the last part of the following section, we also perform a few studies on the impact of EW effects on some of the observables that we consider.

As described in Ref. [361], NNLO QCD and NLO EW corrections can be combined in an additive or multiplicative scheme [22]. If we neglect the loop-induced channel for a moment, we can enclose the higher-order perturbative corrections to the differential LO cross section $d\sigma_{\rm LO}$ in the factors $\delta_{\rm QCD}$ and $\delta_{\rm EW}$ as follows:

$$d\sigma_{\rm NNLO_{QCD}} = d\sigma_{\rm LO}(1+\delta_{\rm QCD}) \quad d\sigma_{\rm NLO_{EW}} = d\sigma_{\rm LO}(1+\delta_{\rm EW})\,. \tag{4.46}$$

Then, the two combination schemes can be written as:

$$\begin{cases} d\sigma_{\rm NNLO_{QCD}+NLO_{EW}} = d\sigma_{\rm LO}(1+\delta_{\rm QCD}+\delta_{\rm EW}) & \text{additive scheme,} \\ d\sigma_{\rm NNLO_{QCD}\times NLO_{EW}} = d\sigma_{\rm LO}(1+\delta_{\rm QCD})(1+\delta_{\rm EW}) & \text{multiplicative scheme.} \end{cases} \tag{4.47}$$

[22] As discussed in Ref. [361], if we denote by α the EW coupling, the multiplicative scheme presented in Eq. (4.47) introduces extra terms of order $\mathcal{O}(\alpha_s\alpha)$ and $\mathcal{O}(\alpha_s^2\alpha)$ compared to the additive one. These terms provide an estimate of mixed QCD-EW corrections, which can become sizeable when both QCD and EW effects are relevant. Such an estimate is particularly reasonable when QCD and EW corrections factorize, which is generally the case when both of them can be regarded as soft corrections to the same hard underlying process. Anyway, one should bare in mind that factorization can be spoilt in peculiar phase-space regions, where for instance giant QCD K-factors arise.

$\sigma(pp \to e^+e^-\mu^+\mu^-)$ [fb]	`setup-inclusive`	`setup-fiducial`
NLO (MATRIX)	$32.50(1)^{+1.9\%}_{-1.6\%}$	$16.49(1)^{+1.9\%}_{-1.6\%}$
MiNLO′	$31.42(3)^{+6.3\%}_{-5.0\%}$	$16.38(2)^{+6.0\%}_{-5.0\%}$
$\text{NNLO}_{q\bar{q}}$ (MATRIX)	$34.42(4)^{+1.0\%}_{-1.0\%}$	$17.45(3)^{+1.0\%}_{-1.0\%}$
NNLO (MATRIX)	$36.57(4)^{+2.4\%}_{-2.1\%}$	$18.84(3)^{+2.5\%}_{-2.1\%}$
nNNLO (MATRIX)	$38.31(4)^{+2.2\%}_{-2.0\%}$	$19.96(3)^{+2.6\%}_{-2.3\%}$
$\text{nNNLO+NLO}_{\text{EW}}$ (MATRIX)	$36.43(7)^{+2.6\%}_{-2.4\%}$	$19.00(4)^{+2.7\%}_{-2.4\%}$
$\text{nNNLO}\times\text{NLO}_{\text{EW}}$ (MATRIX)	$35.63(7)^{+2.5\%}_{-2.3\%}$	$18.58(4)^{+2.6\%}_{-2.3\%}$
$\text{NNLO}_{q\bar{q}}$+PS ($\text{MiNNLO}_{\text{PS}}$)	$34.36(3)^{+0.8\%}_{-1.0\%}$	$17.45(3)^{+0.9\%}_{-1.0\%}$
NNLO+PS ($\text{MiNNLO}_{\text{PS}}$)	$36.50(3)^{+1.9\%}_{-2.0\%}$	$18.90(3)^{+2.5\%}_{-2.0\%}$
nNNLO+PS ($\text{MiNNLO}_{\text{PS}}$)	$38.35(3)^{+2.1\%}_{-2.0\%}$	$20.04(3)^{+2.5\%}_{-2.0\%}$
Extracted from CMS 13 TeV [341]	$39.4 \pm 0.7_{(\text{stat})}$ $\pm 1.1_{(\text{syst})} \pm 0.9_{(\text{theo})} \pm 0.7_{(\text{lumi})}$	$20.3 \pm 0.4_{(\text{stat})}$ $\pm 0.6_{(\text{syst})} \pm 0.4_{(\text{lumi})}$

Table 4.9: Integrated cross sections at various perturbative orders in both the `setup-inclusive` and `setup-fiducial` region. In brackets we report the statistical uncertainties, while scale uncertainties are reported in percentages. We also report the inclusive and fiducial cross sections measured by the CMS experiment in Ref. [341]. Since the measured inclusive cross section corresponds to on-shell $pp \to ZZ$ production, we have multiplied the measured cross section by a branching fraction of $\text{BR}(Z \to \ell^+\ell^-) = 0.03366$, as quoted in Ref. [341], for each Z boson and by a factor of two to compare with our predictions for $pp \to e^+e^-\mu^+\mu^-$ production. For the measured fiducial cross section the CMS analysis includes both different-flavour ($e^+e^-\mu^+\mu^-$) and same-flavour ($e^+e^-e^+e^-$, $\mu^+\mu^-\mu^+\mu^-$) decay channels of the two Z bosons. We have therefore divided the measured fiducial cross section by a factor of two to compare with our $pp \to e^+e^-\mu^+\mu^-$ predictions.

We notice that the inclusion of NLO EW corrections in the two different schemes has a non-negligible impact on the nNNLO result and reduces the cross section by about 4-6% in the fiducial region, slightly deteriorating the agreement with the experimental measurement. We note that EW effects include photon-initiated processes. These have a negligible impact in the fiducial setup, where the leading lepton has a transverse momentum larger than 20 GeV, and all leptons have a transverse momentum larger than 5 GeV. On the contrary, in the inclusive setup, without a minimal transverse momentum, the photon-initiated contribution features a collinear divergence. To avoid this divergence, the CMS analysis [341] imposed a transverse-momentum cut of 5 GeV on the leptons in the evaluation of the photon-induced component. With this cut, they showed that the photon-induced contribution is less than 1% of the total cross section. For this reason, we set the photon-induced component to zero for the nNNLO+NLO$_{\rm EW}$ and nNNLO×NLO$_{\rm EW}$ results in the inclusive case.

4.4.4 Differential distributions and comparison with CMS data

In this section we present our results for differential distributions. We start by comparing the nNNLO+PS predictions obtained with MiNNLO$_{\rm PS}$ against MiNLO′ and fixed-order nNNLO predictions in the `setup-inclusive`. Then, in the last part of this section we move on to consider the `setup-fiducial` and we compare our MiNNLO$_{\rm PS}$ predictions at nNNLO+PS with the data collected and analyzed by the CMS experiment [341].

Comparison against theoretical predictions

In Figure 4.32 we compare nNNLO+PS predictions for MiNNLO$_{\rm PS}$ with MiNLO′ and nNNLO predictions at fixed order for four different distributions which are non-zero at LO. In particular, we consider the invariant mass of the e^+e^- pair ($m_{e^+e^-}$), the invariant mass ($m_{4\ell}$) and the rapidity ($y_{4\ell}$) of the diboson system, and the rapidity (y_{Z_1}) of the Z boson whose invariant mass is closer to m_Z. As we have already clarified, due to the nature of the comparison, both the MiNNLO$_{\rm PS}$ and the MiNLO′ predictions are obtained at parton level, without including hadronization, MPI or photon radiation effects. We observe a very good agreement between the nNNLO+PS and the nNNLO predictions, both for the central values and for the scale-variation bands. The latter are at the few-percent level across the whole range shown in the plots, becoming larger (about ±5%) at high $m_{4\ell}$. Minor differences are visible in the tails of the distributions, in particular at large $m_{4\ell}$, where the nNNLO-accurate MiNNLO$_{\rm PS}$ and fixed-order predictions however still overlap. Indeed, in the large invariant-mass region scale choices and terms beyond accuracy become increasingly important, as we pointed out for W^+W^- production in Section 4.3.4 and as it was also extensively discussed for $t\bar{t}$ production in Ref. [362, 363]. The MiNLO′ result is in all cases about 15-20% smaller than the nNNLO results, which provide mostly flat corrections to the distributions under consideration, increasing slightly only at large $m_{4\ell}$. We stress that the relatively flat QCD corrections are a feature of the chosen distributions (in the inclusive setup) that does not apply in general, as we shall see below. Although the MiNLO′ uncertainty is a factor of 3 larger

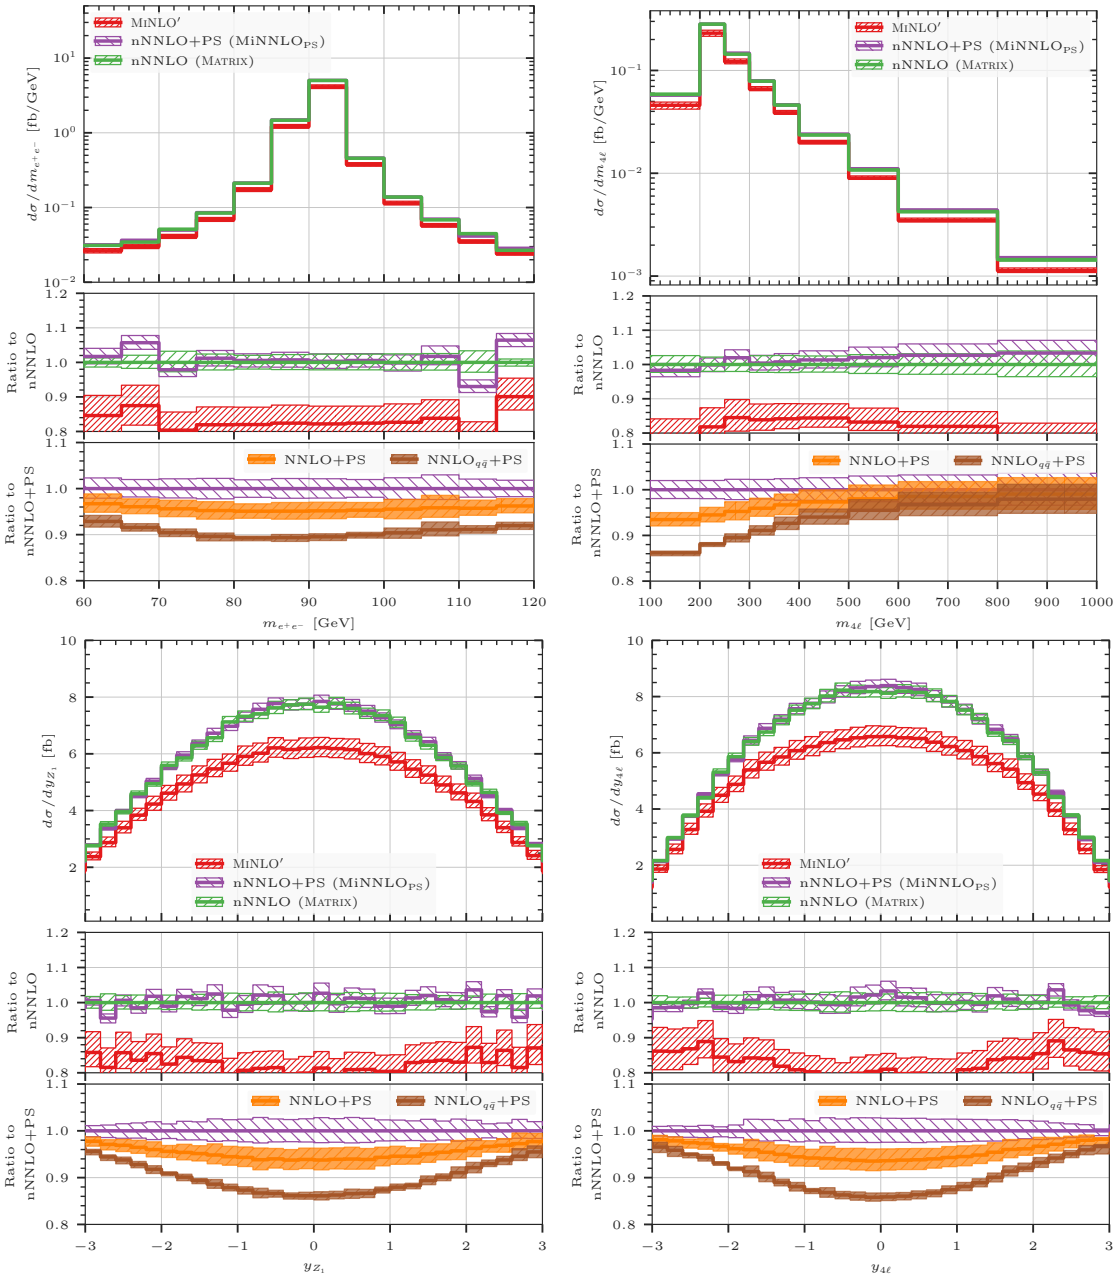

Figure 4.32: Comparison between selected distributions computed with Matrix, MiNNLOPS and MiNLO′. Upper panel: invariant mass of the e^+e^- pair (left) and of the ZZ pair (right); lower panel: rapidity of Z_1 (left) and of the ZZ pair (right).

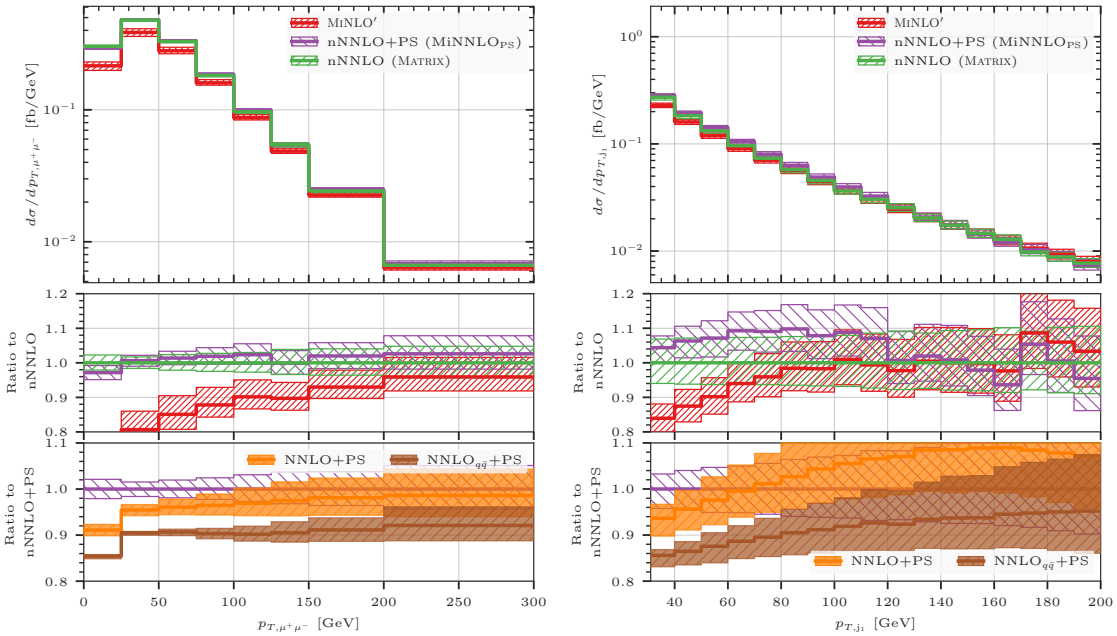

Figure 4.33: Same as Figure 4.32, for the transverse momentum of the $\mu^+\mu^-$ pair (left) and of the leading jet (right).

than the MiNNLO$_{\text{PS}}$ and nNNLO ones, the MiNLO$'$ predictions do not overlap with the nNNLO-accurate results. This is not unexpected since a large part of the difference is caused by the loop-induced gg contribution. Since the latter is missing in the MiNLO$'$ predictions, the MiNLO$'$ scale variation can not account for this new production process, which instead enters the nNNLO results. From the second ratio panel we can appreciate the effect of the loop-induced gg contribution both at LO (comparing NNLO+PS to NNLO$_{q\bar{q}}$+PS) and at NLO (comparing nNNLO+PS to NNLO$_{q\bar{q}}$+PS). It is clear from the plots that due to the gluon flux the impact of the loop-induced gg process is more prominent in certain phase-space regions. The LO (NLO) corrections, which inclusively amount to $\sim$ 6-8% ($\sim$ 10-15%) as pointed out before, contribute more significantly in the bulk regions of the distributions, i.e. at the Z resonance in $m_{e^+e^-}$ as well as for small $m_{4\ell}$ and central rapidities.

In Figure 4.33 we show the same comparison for the transverse momentum of the $\mu^+\mu^-$ pair ($p_{\text{T},\mu^+\mu^-}$) and the transverse momentum of the leading jet (p_{T,j_1}) above 30 GeV. The latter was constructed using the anti-k_{T} algorithm [268] with a jet radius of $R = 0.4$ as implemented in FastJet [269]. While $p_{\text{T},\mu^+\mu^-}$ is already defined at LO, p_{T,j_1} receives its first contribution only at NLO and its accuracy is thus effectively reduced by one perturbative order. The MiNNLO$_{\text{PS}}$ and the nNNLO results for $p_{T,\mu^+\mu^-}$ are in good agreement with each other in the whole range shown here. The MiNLO$'$ result is more than 20% smaller at low values of the transverse momentum, while it agrees with the other two predictions at large values of $p_{\text{T},\mu^+\mu^-}$. Hence, this distribution shows that in general QCD corrections are not uniformly distributed in phase space. By and large, the three

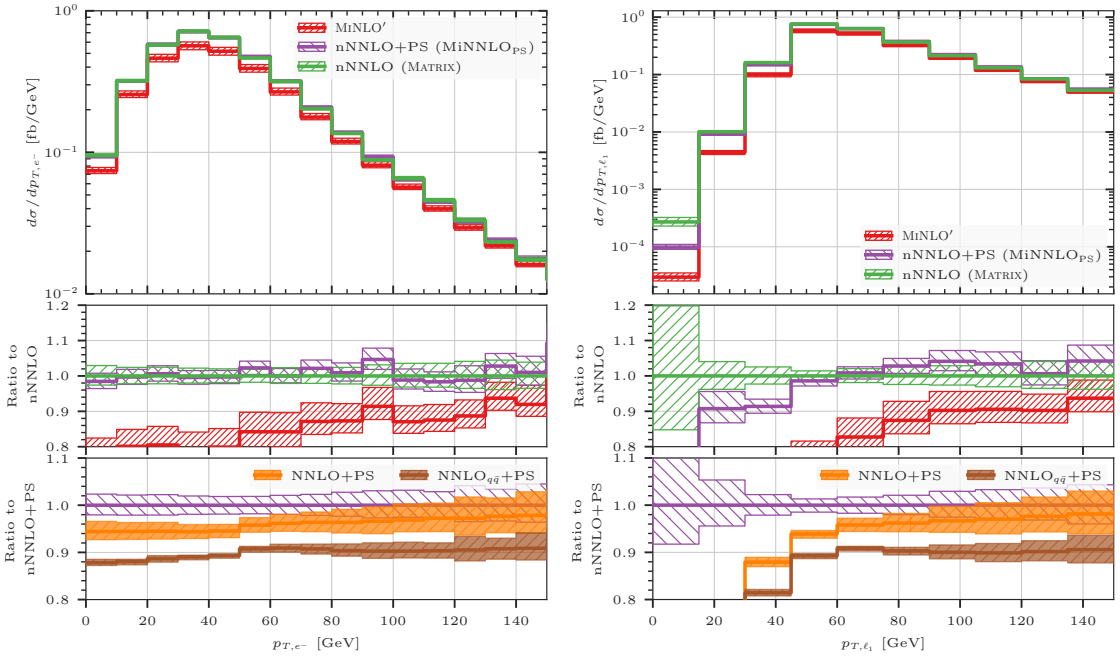

Figure 4.34: Same as Figure 4.32, for the transverse momentum of the electron (left) and of the leading lepton (right).

predictions for the transverse momentum of the leading jet display a good agreement, especially in the tail of the distribution. The level of agreement between nNNLO and MiNNLO$_{\text{PS}}$ is expected as both predictions are effectively nNLO accurate at large p_{T,j_1}. The residual scale uncertainties are at the 5-10% level and they are larger than those in the other distributions, which is a direct consequence of the lower accuracy of the predictions for this observable. Looking at the effect of loop-induced gg contribution in the second ratio panel, we observe a rather peculiar behaviour with the nNNLO+PS corrections being negative with respect to NNLO+PS for $p_{\text{T},j_1} \gtrsim 80\,\text{GeV}$. However, this is completely in line with the results presented in Figure 4.31 and it is a consequence of the fact that the NNLO+PS predictions include only a LO+PS calculation for the loop-induced gg process, which is not expected to accurately describe the high p_{T,j_1} range since it is filled entirely by the parton shower, which has no accuracy in this region. This further underlines the need for including NLO corrections to the loop-induced gg process. Indeed, after including the NLO corrections, the loop-induced gg contribution reduces to 5% (and less) at high p_{T,j_1} (comparing nNNLO to NNLO$_{q\bar{q}}$+PS), which is more reasonable.

In figure 4.34 we show an analogous comparison for the transverse-momentum spectrum of the electron (p_{T,e^-}) and of the leading lepton (p_{T,ℓ_1}). For the p_{T,e^-} distribution we observe excellent agreement over the whole range between the MiNNLO$_{\text{PS}}$ and the nNNLO results, which is fully expected since this distribution should be affected very mildly by resummation/shower effects. We have explicitly checked that a similar level of agreement is obtained when considering the same comparison at NNLO$_{q\bar{q}}$ accuracy,

as opposed to a calculation presented in Ref. [203] in the GENEVA framework (that we have already recalled on different occasions in Chapter 3), where differences between the GENEVA and fixed-order results are observed for $p_{\mathrm{T},e^-} > 40\,\mathrm{GeV}$. When comparing the MiNNLO$_{\mathrm{PS}}$ and the MiNLO$'$ predictions for the p_{T,e^-} spectrum we observe that the effect of both the NNLO$_{q\bar{q}}$ corrections and the loop-induced gg contribution is particularly pronounced in the bulk region of the distribution, where the MiNLO$'$ result is more than 20% smaller than the nNNLO result. On the other hand, the transverse momentum of the leading lepton is subject to shower effects, especially at low p_{T,ℓ_1}, and indeed we observe a difference between the MATRIX results and the MiNNLO$_{\mathrm{PS}}$ predictions below 40 GeV, which becomes increasingly larger the more steeply the MiNNLO$_{\mathrm{PS}}$ distribution falls when p_{T,ℓ_1} approaches zero. Above this value, the shower effects are less pronounced and the two predictions are in good agreement. By comparing the nNNLO+PS predictions to the NNLO+PS and NNLO$_{q\bar{q}}$+PS results we can see that the impact of the loop-induced gg contribution is particularly relevant below 40 GeV, and it is also predominantly responsible for the relatively large shower effects that we observe. In fact, we have checked that for the NNLO$_{q\bar{q}}$+PS result the relative impact of the shower is smaller than for the NLO+PS result in the gg channel, which is expected considering the higher perturbative accuracy (and thereby logarithmic terms) already included at fixed order in the $q\bar{q}$ channel.

Finally, in Figure 4.35 we show predictions for the transverse momentum of the diboson pair ($p_{T,4\ell}$). In this case, we also show the NNLO+N^{3}LL result obtained with MATRIX+RADISH [272], which interfaces MATRIX [81] to the RADISH resummation formalism [105, 274], using $\mu_{\mathrm{R}} = \mu_{\mathrm{F}} = m_{4\ell}$ and $Q_{\mathrm{res}} = m_{4\ell}/2$ for the resummation scale. Since MATRIX+RADISH does not include the contribution stemming from the loop-induced gg channel, we perform this comparison by considering only the $q\bar{q}$-initiated process, i.e. at the NNLO$_{q\bar{q}}$(+PS) level. At small values of the ZZ transverse momentum we observe an excellent agreement between the NNLO+N^{3}LL and the MiNNLO$_{\mathrm{PS}}$ result, especially considering the lower accuracy of the parton shower in that region; MiNNLO$_{\mathrm{PS}}$ is between 5% and 8% larger than the NNLO+N^{3}LL prediction below 10 GeV and has a larger uncertainty band reflecting its lower accuracy. On the other hand, the MiNLO$'$ result is $\mathcal{O}(10\%)$ smaller than the NNLO+N^{3}LL and the MiNNLO$_{\mathrm{PS}}$ predictions and its uncertainty band does not overlap with either of the more accurate results below 40 GeV. Fixed-order calculations actually lead to unphysical results in the small-$p_{\mathrm{T},4\ell}$ region due to large logarithmic corrections, which need to be resummed to all orders. Indeed, the NNLO result diverges at low transverse momentum, and its prediction differs significantly from the ones including resummation effects. At larger values of $p_{\mathrm{T},4\ell}$ the NNLO result is instead in agreement with the NNLO+N^{3}LL, MiNLO$'$ and MiNNLO$_{\mathrm{PS}}$ predictions, as one may expect since all of them have the same formal accuracy in the tail of the distribution.

In conclusion, we observe overall a very good agreement between MiNNLO$_{\mathrm{PS}}$, fixed-order, and analytically resummed results across a variety of distributions, which provides a robust validation of our calculation. The MiNLO$'$ result, despite the considerably larger uncertainty bands, rarely overlaps with the (n)NNLO(+PS) predictions, thus highlighting the importance of higher-order corrections to this process. Moreover, certain

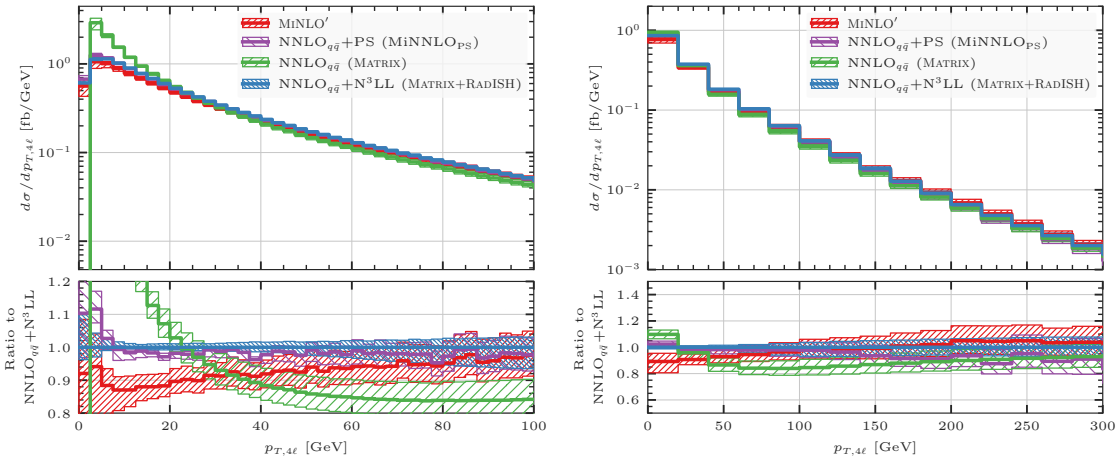

Figure 4.35: Same as Figure 4.32, for the transverse momentum of the ZZ pair for two different ranges of $p_{T,4\ell}$. In both plots, we also show the NNLO+N^{3}LL result computed with Matrix+RadISH [272].

observables require the resummation of large logarithmic contributions, which renders the matching to the parton shower mandatory.

Comparison against data

In this section we compare our MiNNLO$_{\rm PS}$ predictions at nNNLO+PS to the CMS measurement presented in Ref. [341] in the `setup-fiducial` defined in Table 4.8. We have generated the events and estimated the theoretical uncertainties as described in Section 4.4.3. We note that in order to compare against data our predictions include MPI and hadronization effects, as well as QED corrections in the shower approximation.

The comparison between MiNNLO$_{\rm PS}$ predictions and experimental data is presented in Figure 4.36. Altogether, we show predictions for six observables: the invariant mass and the transverse momentum of the diboson pair ($m_{4\ell}$ and $p_{\rm T,4\ell}$), the sum of the four individual transverse-momentum distributions of each final-state lepton (which corresponds to the average of the lepton transverse-momentum distributions), the sum of the two distributions of the transverse momentum of the reconstructed Z bosons (which analogously corresponds to the average of the Z transverse-momentum distributions), and the separation between the two Z bosons in the azimuthal angle ($\Delta\phi_{Z_1,Z_2}$) and in the η–ϕ plane ($\Delta R_{Z_1,Z_2}$). In all cases, except for $\Delta\phi_{Z_1,Z_2}$ that has a kinematical endpoint at $\Delta\phi = \pi$, the last bin shown in the figures also includes the contribution of the overflow.

By and large, we observe a quite remarkable agreement between our predictions and the experimental data. The invariant mass is well described at low $m_{4\ell}$, but there is a tendency of the data to undershoot the prediction at large $m_{4\ell}$, with the last bin being almost two standard deviations away. In this region EW corrections are known to be important and they are only partly included here through the QED shower. Below, we discuss how the inclusion of the NLO EW corrections at fixed order improves the agreement with data in this region. The transverse-momentum distribution of the ZZ

pair is also well described, except for a two-sigma deviation in the last bin, with a remarkable agreement for $p_{\mathrm{T},4\ell}$ values below $\sim 100\,\mathrm{GeV}$, where the all-order corrections provided by the shower are particularly important. The two averaged distributions of p_{T,ℓ_i} and p_{T,Z_i} also compare very well to MiNNLO$_{\mathrm{PS}}$, with only deviations in the tail of the distributions. In the last bins the experimental data are about two standard deviations away from the theoretical predictions, which can again be related to the missing EW corrections, as discussed below. The $\Delta\phi_{Z_1,Z_2}$ and the $\Delta R_{Z_1,Z_2}$ distributions are also very well described by MiNNLO$_{\mathrm{PS}}$, with the data fluctuating (within one sigma, except for one bin with a two-sigma deviation) around the central theoretical prediction across the whole plotted range.

The comparison at the level of integrated cross section in Section 4.4.3 showed that the inclusion of NLO EW effects has a small, but non-negligible impact in the fiducial setup. Since in our comparison with data we include QED effects via parton-shower matching, one may wonder whether the proper inclusion of NLO EW effects in a Monte Carlo context, see e.g. Refs. [364, 365], would further improve the agreement with the data, especially in the tails of distributions where EW logarithms are important. A possible way to assess the impact of the EW corrections beyond the parton shower approximation is to apply to the MiNNLO$_{\mathrm{PS}}$ predictions a differential K-factor accounting for the NLO EW corrections, that are computed at fixed-order accuracy.

We have done this exercise turning off the QED shower in the MiNNLO$_{\mathrm{PS}}$ predictions to avoid double counting. The central rescaled prediction is shown in the lower ratio panels in Figure 4.36. We adopt as our default a factor $K^{(\times)}_{\mathrm{NLO\ EW}}$, defined using the multiplicative scheme nNNLO$\times$NLO$_{\mathrm{EW}}$ [361], which includes an estimate of mixed higher-order corrections, divided by the nNNLO result. Note that for distributions starting at NLO QCD we do not perform this additional comparison, since one would need to compute the EW corrections to the ZZ+1-jet process. We find that the inclusion of NLO EW corrections within this approximation improves the agreement with the experimental data for the tails of the $m_{4\ell}$ and the averaged p_{T,Z_i} distributions, where the effects of EW Sudakov logarithms [23] are expected to be visible. For the averaged p_{T,ℓ_i} distribution their impact is a bit milder, also because the distribution extends to lower values, and there is no significant improvement compared to data. Furthermore, we would like to add some comments on the effects of the QED shower below the $m_{4\ell} \sim 2\,m_Z$ threshold (where QED effects are expected to be sizable) in light of its limited accuracy. First of all, we verified that its effect is below 10% in the first bin of the $m_{4\ell}$ distribution in Figure 4.36,

[23] As also described in a pretty recent review [366], the physical origin of large EW Sudakov corrections as a peculiar EW effect can be easily understood as follows. Large logarithmic terms like $(\alpha/s_W^2)\log^2(s/m_W^2)$ and $(\alpha/s_W^2)\log(s/m_W^2)$ (with α the EW coupling and $s_W = \sin\theta_W$) arise from the exchange of soft/collinear virtual massive gauge bosons at very high energy s. In QED and QCD these terms cancel once the real-emission corrections are integrated out, with at most single logarithmic terms from initial state radiation to be absorbed into bare PDFs (see our discussion in Section 2.1.2). In the context of EW calculations, since the finite gauge boson masses put a physical cutoff on the real integration and since EW gauge bosons are not protected by confinement, they can be experimentally measured (or, to be more precise, reconstructed from their decay products). Therefore, only a small fraction of EW bosons is unresolved and can be integrated out to compensate for the large logarithms arising from virtual corrections.

but it can become instead as large as 35% when considering the `setup-inclusive` for the same bin. Moreover, we notice from Figure 4.36 that the approximation of the QED shower in that region is in good agreement with the result including NLO EW effects. This is not unexpected, because logarithmic contributions due to EW Sudakov effects are small in that region and the first photon emission is the most relevant one for dressed leptons.

The approximate inclusion of NLO EW corrections via a differential K-factor is already enough to show the relevance of moving beyond QCD corrections for an accurate description of LHC observables. Therefore, this simple exercise paves the way for a consistent inclusion of NLO EW effects in our MiNNLO$_{\text{PS}}$ predictions with a complete and consistent matching to QCD and QED showers.

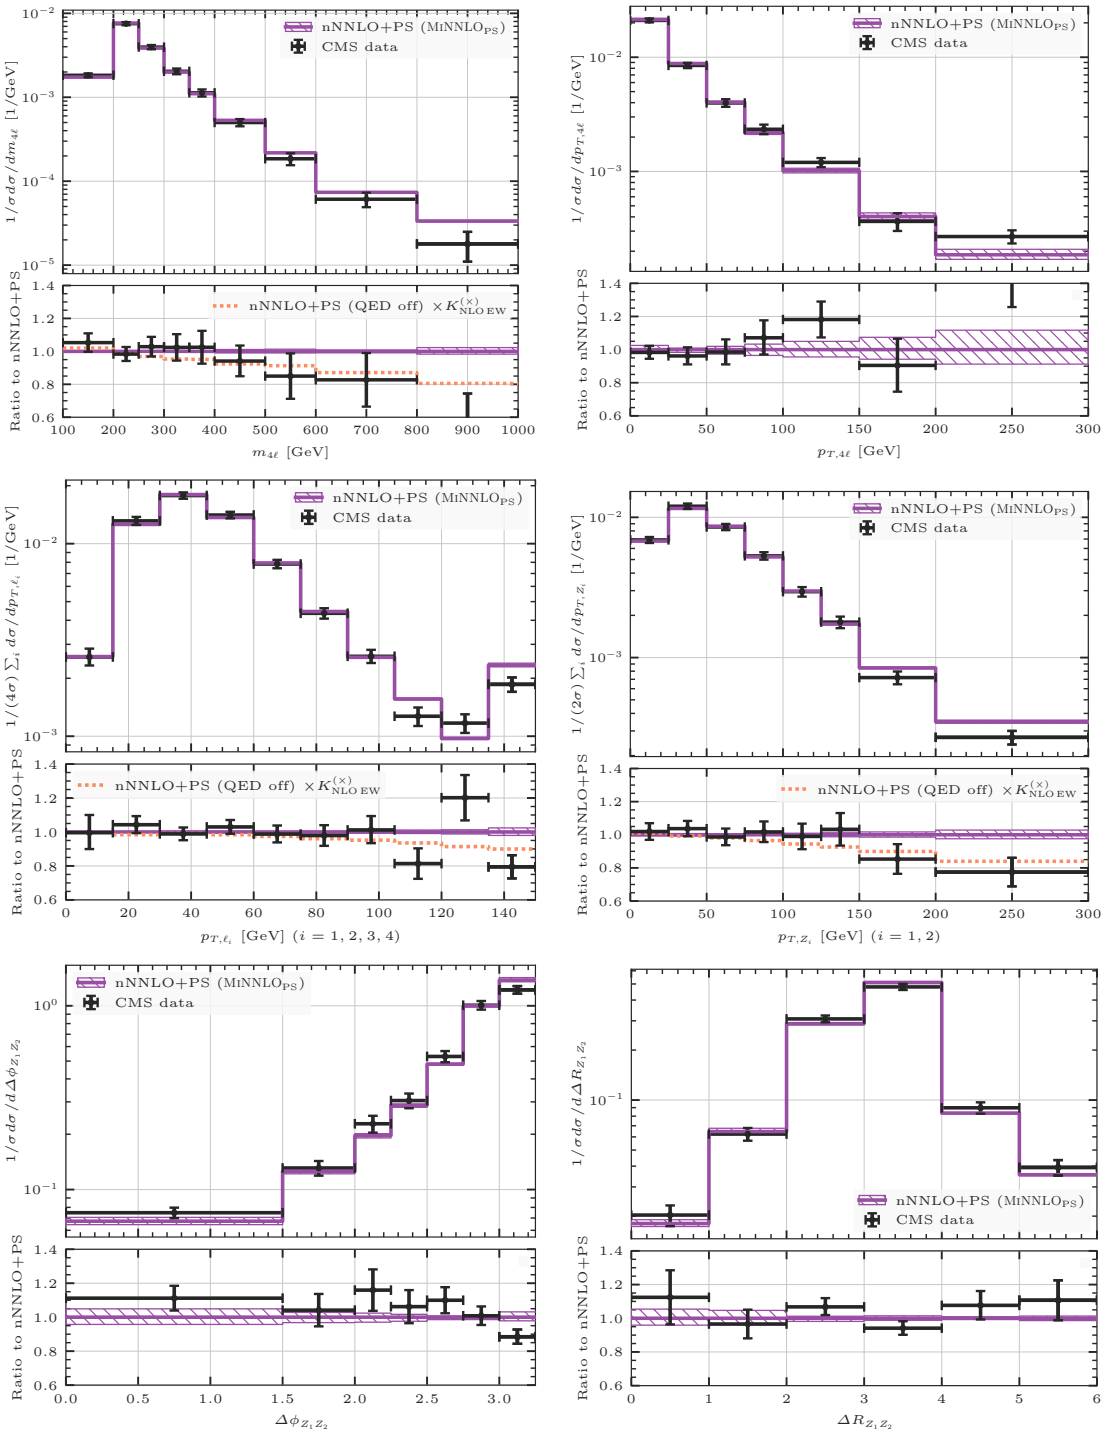

Figure 4.36: Comparison between the MiNNLO$_{\mathrm{PS}}$ predictions and the CMS data of Ref. [341] based on a 137 fb^{-1} 13 TeV analysis for various observables. The MiNNLO$_{\mathrm{PS}}$ predictions include hadronization and MPI effects, as well as QED effects as provided by the Pythia8 parton shower. See text for more details.

5 Conclusions

The leitmovif of this Thesis has been the high-level accuracy of theoretical predictions as an indispensable ingredient for any potential progress in particle physics. In order to benefit as much as possible from the LHC data and to catch up with the precision of the experimental measurements, a realistic and accurate description of particle collisions is required to the theory community. Due to the complexity of a collision event, which is spread over different orders of magnitude in energy, many building blocks enter its simulation: improving the theoretical description of each of them is an extremely important task.

In this work we focused on the consistent combination of two of these constituent blocks: fixed-order calculations and parton showers. In Chapter 2 we have recalled the underlying ideas of the two approaches and we have presented different techniques to simultaneously get the best out of them, namely merging and matching. Despite the vast amount of available methods, event generators capable of reaching NNLO+PS accuracy are by no means standard nowadays. Such a level of accuracy is obtained by matching a NNLO QCD calculation with parton shower, in such a way that NNLO QCD accuracy can be claimed for observables inclusive over QCD radiation, while preserving the logarithmic precision of the parton shower itself.

The results presented in this manuscript paved the way toward NNLO+PS generators for any colour-singlet production process. That has been achieved using the MiNNLO$_{\rm PS}$ method, that we presented in Chapter 3. The MiNNLO$_{\rm PS}$ approach has various positive features. First, NNLO corrections are calculated on-the-fly during the generation of the events, with no need for further re-processing or reweighting of the events. Second, no unphysical merging scale needs to be introduced to separate event samples of different jet multiplicities. Moreover, when combined with transverse-momentum ordered parton showers, the Powheg matching on which MiNNLO$_{\rm PS}$ relies guarantees that the (leading) logarithmic accuracy of the parton shower simulation is preserved.

As a validation of our MiNNLO$_{\rm PS}$ method applied to non-trivial colour-singlet final states, we reported results for diboson production processes in Chapter 4. In Section 4.1, we have presented a novel calculation of NNLO+PS accurate predictions for $Z\gamma$ production at the LHC. That had been the first calculation of a genuine $2 \to 2$ process at this accuracy that does not require an a-posteriori multi-differential reweighting. As a starting point we have implemented NLO+PS generators for both $Z\gamma$ and $Z\gamma$+jet production within the Powheg-Box-Res framework. Then, the $Z\gamma$+jet generator has been extended to include NNLO corrections to $Z\gamma$ production by means of the MiNNLO$_{\rm PS}$ method. The two-loop virtual corrections were obtained from their implementation within the Matrix framework, which has been treated as an external library properly linked to our code. A substantial amount of work has been devoted to render the resulting generator numerically

efficient, due to the complex singular structure of the phase space for processes involving identified final state photons. The results obtained for $Z\gamma$ effectively prepared the stage for all NNLO+PS accurate diboson simulations.

In Section 4.2, we have described an extension of the $Z\gamma$ generator including the effects of anomalous triple gauge couplings at NNLO level. Although our implementation applies also to the $\ell^+\ell^-\gamma$ process, we focused on the $\nu\bar{\nu}\gamma$ final state, due to its particular relevance for BSM searches. We followed the vertex-function approach for a consistent inclusion of the anomalous couplings, in order to directly compare against experimental results, but one should bear in mind that a direct translation from the anomalous coupling to the EFT framework is possible. Reaching NNLO+PS accuracy in this context can be highly relevant to improve the experimental bounds on anomalous triple gauge couplings. Moreover, we showed that increasing the accuracy of $\nu\bar{\nu}\gamma$ simulations is also crucial to reduce the uncertainty on the dominant background in dark-matter searches in the photon plus missing energy channel.

In the remaining part of this work, we moved on to considering applications of the MiNNLO$_{\text{PS}}$ method to production processes for a pair of massive gauge bosons. Results for W^+W^- production have been presented in Section 4.3. We have performed the calculation consistently in the four-flavour scheme with massive bottom quarks. By dropping contributions with final-state bottom quarks, which are regulated by the finite bottom mass, we have generated top-free W^+W^- events, as required to compare against experiments, which apply jet-veto cuts to suppress the top-quark background. Our new MiNNLO$_{\text{PS}}$-based W^+W^- generator superseeds the NNLOPS results of Ref. [211], where NNLO accuracy was achieved by a computationally expensive event reweighting: MiNNLO$_{\text{PS}}$ includes NNLO corrections directly at event generation level, making the simulation way more efficient. We found that the major bottleneck in our computation was the evaluation of the two-loop amplitude. To improve the speed, we have constructed four-dimensional grids of the coefficients that encode all the information required to reproduce the full two-loop contribution. We have then used those grids to obtain the coefficients at any given phase-space configuration through a four-dimensional cubic spline interpolation, so as to reconstruct the two-loop amplitude on-the-fly. As a result, the evaluation time of the two-loop contribution has been reduced by a factor of forty, rendering it subleading with respect to the other parts of the calculation.

Since our W^+W^- generator does not include the luminosity-enhanced contribution coming from the loop-induced gluon-gluon channel, a direct comparison against experiments has not been possible. We accounted for this contribution in the results of Section 4.4, where we presented a Monte Carlo event generator for Z-boson pair production at the LHC. For the $q\bar{q}$-initiated process we have matched NNLO QCD predictions to parton showers using the MiNNLO$_{\text{PS}}$ method. Then, we have included the loop-induced gg-initiated process, which contributes starting from $\mathcal{O}(\alpha_s^2)$, in the Powheg-Box-Res framework at NLO QCD accuracy matched to parton showers. When combined, the ensuing nNNLO+PS results constitute the most accurate theoretical predictions for this process to date. Just as with W^+W^-, the evaluation of the two-loop contributions turned out to be numerically highly demanding. In this case we made full use of the reweighting facility of Powheg and introduced the possibility to evaluate the two-loop contributions

only at the very end of the event generation, considerably speeding up the calculation.

Our MiNNLO$_{\rm PS}$ predictions have been extensively validated in many respects. For observables inclusive over QCD radiation, MiNNLO$_{\rm PS}$ results compare extremely well against fixed-order NNLO calculations, despite the different treatment of terms beyond the accuracy. But the matching to the shower is crucial for observables sensitive to soft-gluon effects, where NNLO results fail to provide a suitable description. For these observables, we also compared MiNNLO$_{\rm PS}$ simulations with high-accurate analytically resummed results for the transverse momentum of the colour-singlet system: despite the lower logarithmic accuracy of the shower, MiNNLO$_{\rm PS}$ results still provide a very reasonable description. Moreover, they can automatically describe in a physically sound way even those observables where multiple classes of enhanced logarithms simultaneously appear, and which are hard to tackle analytically. Finally, we performed various comparisons of our results with experimental data. For $Z\gamma$ production in the $\ell^+\ell^-\gamma$ channel, we found an excellent agreement between our MiNNLO$_{\rm PS}$ predictions and the latest ATLAS 13 TeV data of Ref. [238]. We remind that the few discrepancies emerging for some observables can be easily washed out by including QED effects, and especially photon radiation from the final state leptons. For our nNNLO+PS predictions for ZZ production, we compared against 13 TeV CMS data of Ref. [341]. In this case we also observed a remarkable agreement both at the level of production rates and shapes of kinematical distributions, with nNNLO+PS predictions and CMS data agreeing on almost all bins within one sigma. In the few bins where the differences are at the two-sigma level we have shown that the inclusion of NLO EW corrections removes those differences in most instances. That suggests how consistently accounting for QED and NLO EW effects within MiNNLO$_{\rm PS}$ is a required step to further improve the description of LHC observables in some relevant kinematic regions.

Therefore, from the results of this work it emerges quite clearly that our extension of the MiNNLO$_{\rm PS}$ method to any colour-singlet process is an essential reach on the way to improve current Monte Carlo event simulations. From a theoretical perspective, the achievements presented in this Thesis will represent a starting point for further progresses in this research field. Moreover, we believe that our new event generators, which are already available on http://powhegbox.mib.infn.it within the Powheg-Box-Res framework (see also Appendix B for technical details on the download and usage of the codes), will definitely be valuable tools for upcoming experimental analyses involving gauge-boson pair production processes.

Appendix A

Technical details of the $Z\gamma$ generator

In this Appendix we collect some non-standard settings that we have implemented in order to build our MiNNLO$_{\rm PS}$ $Z\gamma$ generator. In Section A.1 we explicitly present how we handled the non-trivial singular structure of the $Z\gamma$ phase space. In Section A.2 we report some subtleties involved in the projection from the $Z\gamma$+jet to the $Z\gamma$ phase space. Finally, in Section A.3 we comment on the impact of the inclusion of heavy-quark loops. Even though these considerations were originally done for the $\ell^+\ell^-\gamma$ channel, presented in Section 4.1, they are also valid when running our generator in the $\nu\bar{\nu}\gamma$ channel of Section 4.2. Clearly, in the latter case the discussion presented in Section A.1 simplifies, since no generation cuts or related suppression factors are required to avoid the photon-pole divergence or the final-state collinear singularities of the photon with the Z-boson decay products.

A.1 Improving the phase space sampling

Since the Born process for $Z\gamma$ production, especially in the $\ell^+\ell^-\gamma$ channel, involves a number of QED and QCD singularities, we make use of Born and remnant suppression factors to sample the phase space. Additionally, we introduce a number of small technical cuts in the phase-space generation. In this Appendix we give all details about the generation cuts and suppression factors that we have used to obtain the results presented in Section 4.1 and partially also in Section 4.2.

We start by outlining the generation cuts that we employ. First, we introduce a lower cut $p_{\rm T,\gamma}^{\rm cut} = 5\,\text{GeV}$ on the photon transverse momentum, which is required to avoid QED singularities related to collinear photon emissions from the initial states. We also impose a similar cut of $p_{{\rm T},j}^{\rm cut} = 1\,\text{GeV}$ on the transverse momentum of the outgoing QCD partons. A lower cut $m_{\ell\ell}^{\rm cut} = 40\,\text{GeV}$ on the invariant mass of the lepton pair is imposed to avoid singular configurations in $\gamma^* \to \ell^+\ell^-$ splittings. Note that, since the invariant mass of the resonances are preserved when radiation is generated (both within Powheg and the shower), any value $m_{\ell\ell}^{\rm cut}$ equal or below the cut used in the analysis is allowed. Furthermore, we require that the photon is isolated from leptons and QCD partons in the final state. For this purpose, we introduce a cut $m_{\ell\gamma}^2 = 0.1\ \text{GeV}^2$, and we define a smooth isolation as in Eq. (4.4) using $E_{\rm T}^{\rm ref} = \epsilon_\gamma p_{{\rm T},\gamma}$ with $\delta_0 = 0.05$, $\epsilon_\gamma = 0.5$ and $n = 1$. All these generation cuts can be modified via the input card, but for consistency reasons their values should be much smaller than the values used in the fiducial phase-space definition at analysis level. Note that we have explicitly checked that removing the

aforementioned generation cuts (once proper suppression factors are included) does not have any impact beyond the numerical uncertainties for the $Z\gamma$ observables considered in Section 4.1 and 4.2. Nevertheless, we have kept them to avoid potential instabilities in the generation of events [1].

The Born suppression factor that we adopt is constructed in factorized form

$$B_{\text{supp}} = F_{\text{supp}}(p_{\text{T},\gamma}) \cdot G_{\text{supp}}(\Delta R_{\gamma,\,\ell^+}) \cdot G_{\text{supp}}(\Delta R_{\gamma,\,\ell^-}) \cdot H_{\text{supp}}(\Delta R_{\gamma,\,j})\,, \tag{A.1}$$

with

$$F_{\text{supp}}(p_{\text{T},\gamma}) = \frac{(p_{\text{T},\gamma})^2}{(p_{\text{T},\gamma})^2 + (p^0_{\text{T},\gamma})^2}\,, \quad \text{with} \quad p^0_{\text{T},\gamma} = 10\,\text{GeV}\,, \tag{A.2}$$

$$G_{\text{supp}}(\Delta R) = \frac{(\Delta R)^2}{(\Delta R)^2 + (\Delta R_0)^2}\,, \quad \text{with} \quad \Delta R_0 = 0.5\,, \tag{A.3}$$

and

$$H_{\text{supp}}(\Delta R) = \frac{(\Delta R)^2}{(\Delta R)^2 + (\Delta R_0)^2}\,, \quad \text{with} \quad \Delta R_0 = 0.2\,. \tag{A.4}$$

Since we apply an overall Sudakov form factor through MiNLO′/MiNNLO$_{\text{PS}}$, we do not need any suppression related to the outgoing parton for $Z\gamma$+jet production. It is clear that, whenever a singularity is approached, the Born suppression factor in Eq. (A.1) vanishes in such a way that the cross section times Born suppression factor itself remains finite.

As discussed in Section 4.1.1, we have to deal with a remnant contribution which is QCD regular, but QED singular. Accordingly, we introduce a remnant suppression factor of the form

$$\begin{aligned} R_{\text{supp}} = {} & F_{\text{supp}}(p_{\text{T},\gamma}) \cdot G_{\text{supp}}(\Delta R_{\gamma,\ell^+}) \cdot G_{\text{supp}}(\Delta R_{\gamma,\ell^-}) \cdot H_{\text{supp}}(\Delta R_{\gamma,j_1}) \\ & \cdot H_{\text{supp}}(\Delta R_{\gamma,j_2}) \cdot H_{\text{supp}}(\Delta R_{j_1,j_2}) \cdot L_{\text{supp}}(p_{\text{T},j_2})\,, \end{aligned} \tag{A.5}$$

with

$$L_{\text{supp}}(p_{\text{T},j_2}) = \frac{(p_{\text{T},j_2})^2}{(p_{\text{T},j_2})^2 + (p^0_{\text{T},j_2})^2}\,, \quad \text{with} \quad p^0_{\text{T},j_2} = 20\,\text{GeV}\,. \tag{A.6}$$

As usual in Powheg, the Born suppression is evaluated using the Born kinematics, while the remnant suppression is evaluated using the kinematics of the real phase space. We note that in our case it is not necessary to introduce the additional suppression factor $L_{\text{supp}}(p_{\text{T},j_2})$, but we found that results converged more quickly if this additional factor is also included.

[1] Note that in the release of the code of the $Z\gamma$ event generator these generation cuts are absent by default, since they might induce effects in other fiducial setups and their values should be chosen accordingly.

A.2 Projection from the $Z\gamma$+jet to the $Z\gamma$ phase space

The evaluation of the last term in Eq. (3.80) requires a projection from the $Z\gamma$+jet to the $Z\gamma$ kinematics. In this Appendix, we give details about this projection and comment on configurations which after projection have $p_{\mathrm{T},\gamma}$ close to zero.

We denote by p_1 and p_2 the two incoming momenta, and by p_γ, p_Z and p_j the momenta of the photon, the Z boson and the jet in the final state. We define $p_{\mathrm{tot}} = p_1 + p_2 - p_j = p_\gamma + p_Z$. Our projection to the respective $Z\gamma$ configuration is the one adopted for initial-state radiation in POWHEG, which is described in Section 5.1.1 of Ref. [86] (cf. also Appendix A of Ref. [213]). It consists of a longitudinal boost (by β_{L}), such that, after boosting, p_{tot} has no z component (no component along the direction of the collision axis). Then, a second boost (by $\vec{\beta}_{\mathrm{T}}$) in the transverse plane, such that p_{tot} has no transverse component, is applied, followed by a final boost back in the longitudinal direction (by $-\beta_{\mathrm{L}}$). We add a prime to all quantities after the first longitudinal boost and a double prime to those after the second one. The boost vector of the transverse boost is then given by

$$\vec{\beta}_{\mathrm{T}} = \frac{\vec{p}_{\mathrm{T},j}}{E'_{\mathrm{tot}}}\,. \tag{A.7}$$

After the second boost, the transverse momentum of the photon becomes

$$\vec{p}''_{\mathrm{T},\gamma} = \gamma_{\mathrm{T}} \left(\vec{p}_{\mathrm{T},\gamma} + \vec{\beta}_{\mathrm{T}} E'_\gamma\right)\,, \quad \text{with} \quad \gamma_{\mathrm{T}} = \frac{1}{\sqrt{1-\beta_{\mathrm{T}}^2}}\,. \tag{A.8}$$

Therefore, after this boost, the condition $\vec{p}''_{\mathrm{T},\gamma} = 0$ is met if

$$\vec{\beta}_{\mathrm{T}} = -\frac{\vec{p}_{\mathrm{T},\gamma}}{E'_\gamma}\,. \tag{A.9}$$

By comparing Eq. (A.7) and (A.9) we see that $\vec{p}_{\mathrm{T},\gamma}$ and $\vec{p}_{\mathrm{T},j}$ must be anti-aligned. Furthermore, since $E'_{\mathrm{tot}} = E'_\gamma + E'_Z > E'_\gamma$ it follows that $p_{\mathrm{T},j} > p_{\mathrm{T},\gamma}$. Accordingly, any boost that leads to a vanishing transverse momentum of the photon in the $Z\gamma$ configuration has a jet that is harder than the photon in the $Z\gamma$+jet configuration. Since for the photon we impose a transverse-momentum cut, these configurations are free of any large logarithms, and the corrections from $\mathcal{D}_{\ell_{\mathrm{F}}}(p_{\mathrm{T}})$ of Eq. (3.84) entering Eq. (3.79) in this region of phase space are strictly beyond our accuracy and can be safely dropped.

A.3 OpenLoops vs. MCFM implementation: heavy-quark mass effects

In this Appendix, we provide additional information regarding the limitation of the MiNNLO$_{\mathrm{PS}}$ generator when using MCFM amplitudes by comparing the results presented in Sections 4.1.4, 4.1.5, 4.1.6 and 4.1.7 against our alternative implementation that uses OpenLoops amplitudes. As described in detail in Section 4.1, OpenLoops includes

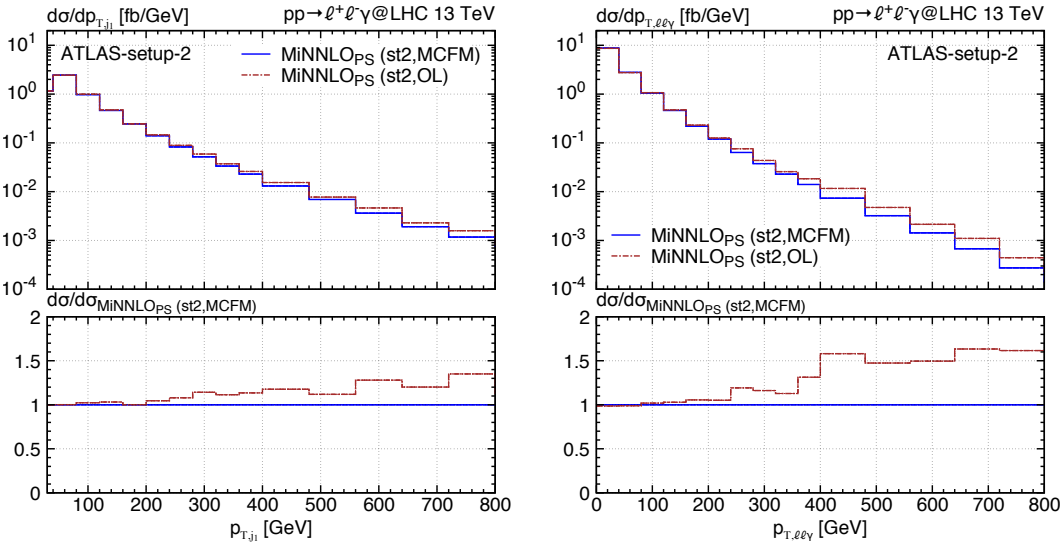

Figure A.1: Comparison of MiNNLO$_{\text{PS}}$ predictions using MCFM (blue, solid) and using OpenLoops (brown, dash-dotted) for the distribution in the transverse momentum of the leading jet (left plot) and of the $Z\gamma$ system (right plot) at Powheg stage 2 in `ATLAS-setup-2` (see Section 4.1.3).

the full dependence on heavy-quark masses in the one-loop amplitudes, while in the MCFM ones bottom-quark loops are included only partially and top-quark loops are omitted entirely. To assess the validity of the approximations done within the MCFM implementation with respect to including the full heavy-quark mass dependence, we consider MiNNLO$_{\text{PS}}$ predictions at Powheg stage 2 using MCFM and OpenLoops. We have compared all distributions that we produced in `ATLAS-setup-1` and `ATLAS-setup-2` (defined in Section 4.1.3) and found that for observables inclusive over QCD radiation heavy-quark mass effects are at the level of the numerical uncertainties and can be safely neglected. At variance, as far as observables that require hard/boosted jets are concerned, heavy-quark mass effects become important and the MCFM and OpenLoops results can differ substantially. For illustration and to provide relevant examples, Figure A.1 compares our MCFM (blue, solid) and OpenLoops (brown, dash-dotted) implementation of MiNNLO$_{\text{PS}}$ for the transverse-momentum spectrum of the leading jet and of the $Z\gamma$ system in `ATLAS-setup-2` at Powheg stage 2. Indeed, in the bulk of the cross section at low transverse momentum heavy-quark mass effects are negligible. For $p_T \gtrsim 150\,\text{GeV}$, however, we can appreciate the increase of the OpenLoops result with respect to MCFM, indicating the importance of heavy-quark mass effects and the limitation of the MCFM implementation.

Appendix B

Usage of the MiNNLO$_{\rm PS}$ code for diboson processes

In this appendix we present some practical details that are required to correctly access to the diboson codes that have been used for the results described in Chapter 4. All processes have been implemented within the POWHEG-BOX-RES framework and are available as sub-folders inside the main `POWHEG-BOX-RES` directory. To keep the discussion as general as possible, we refer to a general colour-singlet final state as F, where F can be $Z\gamma$, W^+W^- or ZZ; accordingly, the corresponding process folder is denoted by `FJ`. All technical details reported in this appendix can be found in the corresponding folders of the processes, at the path `POWHEG-BOX-RES/FJ/Docs/manual_F-MiNNLO.pdf`. We start in Section B.1 with the description of the instructions of how to download and compile the process code correctly. Then, in Section B.2, we briefly recall how to run a process within the POWHEG-BOX framework. Section B.3 includes the description of some process-specific parameters that can be adjusted from the input card of the process. Finally, some general features on how to run the MiNNLO$_{\rm PS}$ machinery are reported in Sections B.4 and B.5.

B.1 Compiling the process

To obtain the desired code do

```
$ svn checkout --username anonymous --password anonymous
svn://powhegbox.mib.infn.it/trunk/User-Processes-RES/FJ
```

inside the `POWHEG-BOX-RES` directory.

To use MiNNLO$_{\rm PS}$, MATRIX [81] is required since it is used by MiNNLO$_{\rm PS}$ as a library. Therefore, one first needs to initialize MATRIX (and compile OPENLOOPS, GINAC, and CLN, or use a compiled version of them [1]) in order for the code to be able to compute the two-loop contributions needed to reach NNLO accuracy for inclusive colour-singlet production. We recall that for W^+W^- and ZZ production the two-loop coefficients are obtained from the public code `VVAMP` [316]. To initialize MATRIX, you have to execute from the `POWHEG-BOX-RES` folder [2]:

[1] The MATRIX code makes use of OPENLOOPS for the generation of tree-level and one-loop amplitudes, of GINAC, a `C++` library for symbolic algebraic manipulation, and of CLN, a `C++` library for arbitrary precision arithmetic.

[2] Note that you need to use `Python` version 2.

```
$ cd MATRIXStuff
$ ./matrix --minnlo_interface
```

which installs OpenLoops 2 [246], Ginac [367] and CLN [368] packages. If you would like to use existing local installations of any of those programs, you need to set the correct paths both in the file `config/MATRIX_configuration` inside the `MATRIXStuff` folder, and later also in the `Makefile` inside the `FJ` folder. If there the automatic installation (through the matrix script) of Ginac or CLN fails you may check inside the `external` folder for `ginac-install` and `cln-install`, respectively, and try to complete the compilation inside those folders manually. Note that for $Z\gamma$ production the installation of Ginac and CLN is actually not strictly necessary, but it is done by default also for that process for completeness, since needed by other diboson processes. If the user really wants to avoid the installation of the two packages, one can also initialize Matrix as follows:

```
$ ./matrix --minnlo_interface --no_cln --no_ginac
```

To compile the specific FJ code, one can now move to the FJ directory and run

```
$ make -j
```

For the code to correctly run, LHAPDF [224] and FastJet [269] need to be installed and both `lhapdf-config` and `fastjet-config` should be recognized by the terminal, or the corresponding variables inside the `Makefile` have to be set appropriately. Note that the first compilation might take several hours to days (especially when running only on a single core) due to the many files compiled for the two-loop amplitudes within `VVamp`. Thus compiling on several cores is strongly recommended.

As described in Section 4.3, for the WWJ process the additional possibility to evaluate the two-loop contribution with an interpolation strategy is available. In order to use the interpolator the user shall download the interpolation grids and unpack them. Depending on the interpolator running mode specified in the `Makefile`, a different set of interpolation grids has to be downloaded. If the variable `INTERPOLATOR` in the `Makefile` is set to `V0`, the interpolator will upload into memory the full set of grids for the two-loop coefficient functions. In this case, the grids should be downloaded and unpacked using:

```
$ wget https://wwwth.mpp.mpg.de/members/wieseman/download/codes/
   WW_MiNNLO/VVamp_interpolation_grids/
   WW_MiNNLO_2loop_grids_full.tar.gz
$ tar xfz WW_MiNNLO_2loop_grids_full.tar.gz
```

(in case of protocol issues, you can try to download the grids from the same link using `curl -LO`). If `INTERPOLATOR=V1` (this is the default setting), the interpolator makes partial use of a set of crossing relations among the two-loop coefficient functions to reduce the number of grids to be loaded into memory. When compiling the code in this mode, a smaller set of grids can be obtained doing

```
$ wget https://wwwth.mpp.mpg.de/members/wieseman/download/codes/
   WW_MiNNLO/VVamp_interpolation_grids/
   WW_MiNNLO_2loop_grids_reduced1.tar.gz
$ tar xfz WW_MiNNLO_2loop_grids_reduced1.tar.gz
```

Finally, with `INTERPOLATOR=V2` the interpolator takes full advantage of the crossing relations mentioned above to reduce the needed grids even further. For this specific case, a dedicated set of grids can be downloaded:

```
$ wget https://wwwth.mpp.mpg.de/members/wieseman/download/codes/
   WW_MiNNLO/VVamp_interpolation_grids/
   WW_MiNNLO_2loop_grids_reduced2.tar.gz
$ tar xfz WW_MiNNLO_2loop_grids_reduced2.tar.gz
```

The option `INTERPOLATOR=V1` is recommended to get the best compromise between memory usage and accuracy of the interpolation procedure. Once the grids are available, in order to be used they have to be unpacked either into the `WWJ` directory or into the current run directory, otherwise the program will abort with a proper error message.

B.2 Generation of events

The running of the code and the generation of diboson events is standard (as for a usual POWHEG process) and we summarize here the main points. First, to perform a full run one can for example do

```
$ cd testrun-lhc
$ ../pwhg_main
```

where `testrun-lhc` can be any directory containing a proper `powheg.input` card. We provide a `powheg.input-save` card inside the `suggested-minnlo-run` folder for a MiNNLO$_{\rm PS}$ run; for $Z\gamma$ production, specific input cards to run MiNNLO$_{\rm PS}$ for $\ell^+\ell^-\gamma$ and $\nu_\ell\bar{\nu}_\ell\gamma$ final states can be found in the dedicated folders `suggested-minnlo-run_llgam` and `suggested-minnlo-run_nunugam`, respectively. At the end of the run, the file `pwgevents.lhe` will contain events for F production in association with one jet in the LHE format. In order to shower them with PYTHIA8 [124] you need to have a separate PYTHIA8 installation, and then simply execute:

```
$ make main-PYTHIA8-lhef
```

from the FJ directory. Then, similarly to the generation of events, move to the directory containing the input card and the generated event file to launch the shower program:

```
$ cd testrun-lhc
$ ../main-PYTHIA8-lhef
```

Note that without parallelization any diboson code will run for several hours/days and still not get very accurate results, since these programs are numerically intensive. Therefore, we do not recommend to run them without the POWHEG parallel feature (described in detail in `POWHEG-BOX-V2/Docs/V2-paper.pdf`). It allows to parallelize different steps of the calculation and to recombine intermediate integration grids and results. Indeed, it is turned on by default in the template input cards inside the `suggested-minnlo-run` directories.

B.3 Process-specific input parameters

We now comment on some important process-specific input settings in the `powheg.input` file. We start from the description of the most relevant $Z\gamma$ inputs in the subsection B.3.1, together with the presentation of specific inputs to include the contribution of aTGCs in subsection B.3.2. Then, we continue with inputs for W^+W^- and ZZ, respectively in subsections B.3.3 and B.3.4.

B.3.1 Zγ–MiNNLO$_{\rm PS}$ input parameters

The $Z\gamma$ code can be run for a Z boson decaying either into charged leptons ($\ell^+\ell^-\gamma$) or neutrinos ($\nu_\ell\bar{\nu}_\ell\gamma$). The user can switch between the two decay modes by choosing the proper value for the input parameter `vdecaymode`:

```
vdecaymode          1     ! Z decay products (default 1): 0 for neutrino decay,
#                         !  1 for charged lepton decay
```

For `vdecaymode 1` different decay modes for the Z boson can be selected:

```
e+e-                1     ! (default 1) if 1, Z decays into electrons (only if
#                         !  vdecaymode 1)
mu+mu-              0     ! (default 0) if 1, Z decays into muons (only if
#                         !  vdecaymode 1)
tau+tau-            0     ! (default 0) if 1, Z decays into taus (only if
#                         !  vdecaymode 1)
leptonic_notau      0     ! (default 0) if 1, Z decays into electrons and muons
#                         !  (only if vdecaymode 1)
leptonic            0     ! (default 0) if 1, Z decays into electrons, muons and
#                         !  taus (only if vdecaymode 1)
```

Moreover, if `massive_leptons 1`, the charged lepton momenta are reshuffled according to their physical masses at event generation level (while they are customarily set to zero in Powheg stage 1 and 2):

```
massive_leptons     0     ! (default 0) if 1, lepton momenta are reshuffled at
#                         !  event generation level to their mass values
e_mass       0.511d-3     ! (default 0.511d-3) electron mass for reshuffling
#                         !  momentum at event generation level (used if
#                         !  massive_leptons=1)
mu_mass      0.1057d0     ! (default 0.1057d0) muon mass for reshuffling momentum
#                         !  at event generation level (used if massive_leptons=1)
tau_mass      1.777d0     ! (default 1.777d0) tau mass for reshuffling momentum
#                         !  at event generation level (used if massive_leptons=1)
```

For `vdecaymode 0` all neutrino flavours are included at event generation level. If one wants to account for all three neutrino families already from Powheg stage 1, you can simply set `sum_over_families 1`:

```
sum_over_families   0     ! (default 0) if 0, a factor of 3 accounting for all
#                         !  neutrino families is included only in the
#                         !  rad_branching  variable at event-generation
#                         !  level (otherwise it is included throughout the
#                         !  calculation)
```

When doing a $Z\gamma$+jet NLO+PS calculation (turning off MiNNLO$_{\rm PS}$ and MiNLO$'$) one may set the renormalization and factorization scales by changing

```
fixedscale          0    ! (default 1) if 0, use dynamical scale below (set
#                        !  by whichscale); if 1, scale is fixed to Z-boson mass
#                        !  (leave this to 0 when MiN(N)LO is used)
whichscale          1    ! (default 0) only used if fixedscale=0:
#                        !  0 = sqrt(M_Z^2+pt_gamma^2), 1 = M_llgamma
#                        !  (leave this to 1 when MiN(N)LO is used)
```

By switching on or off the `fixedscale` input parameter, you can choose between a fixed or dynamical scale, respectively. In the latter case, a specific dynamical scale can be selected by changing the value of `whichscale`. When MiNLO′ or MiNNLO$_{\rm PS}$ is turned on, this scale setting is not used and instead the scales are set as described in Section 3.2 and 3.3.

Since already at LO the $Z\gamma$+1 jet phase space contains both QCD and QED singularities, redirecting the numerical sampling of events into certain regions in phase space through proper suppression factors can significantly improve the convergence of the numerical integration (namely Powheg stage 1 and 2). Moreover, the event generation can be made more efficient by splitting the squared real matrix element into a singular and remnant contribution. These features can be controlled by the following set of input parameters:

```
withdamp            1    ! (default 1) if 1, split real amplitude into a singular
#                        !  and remnant contribution
suppmodel           1    ! (default 1) if 1, the QED divergences are moved into
#                        !  the remnant
bornsuppfact        1    ! (default 1) if 1, the Born suppression factor is
#                        !  included. Weighted events are generated. If 0 no
#                        !  suppression factor is included, and events are
#                        !  unweighted. A generation cut pt_j1_cut>0 must be
#                        !  supplied in the latter case
remnsuppfact        1    ! (default 1) if 1, the remnant suppression factor is
#                        !  included
```

As standard in the `POWHEG-BOX-RES` code, by setting `withdamp 1`, the real amplitude is split into a QCD singular and remnant contribution. If `suppmodel` is set to `1`, the QED divergences of the real amplitude are also included as part of the remnant definition (see discussion in Section 4.1.1). Then, a born and a remnant suppression factor multiplying the cross section during integration and a posteriori divided out again can be used to suppress the event sampling in the QCD and QED singular phase-space regions. These two factors are separately controlled by `bornsuppfact` and `remnsuppfact`, respectively. The exact form of the suppression factors can also be tuned from the input by modifying the parameters (as a reference, see also formulae reported in Appendix A.1):

```
ptj_suppfact     20d0    ! (default 20d0) suppression of regions with jet(s) pt
#                        !  lower than ptj_suppfact -> if anommode is on, the
#                        !  default value is 50d0
powj_suppfact       2    ! (default 2) power controlling the ptj_suppfact term
pta_suppfact     10d0    ! (default 10d0/80d0 for charged lepton/neutrino decay
#                        !  mode) suppression of regions with photon pt lower
#                        !  than pta_suppfact -> if anommode is on, the default
#                        !  value is 600d0
ptnunu_suppfact 150d0    ! (default 150d0) suppression of regions with pt of
#                        !  neutrino couple lower than ptnunu_suppfact -> only
#                        !  used if anommode is on
```

```
powa_suppfact       2    ! (default 2) power controlling the pta_suppfact term
DRal_suppfact   0.5d0    ! (default 0.5d0) suppression of regions with charged
#                        !  lepton-photon DR distance lower than DRal_suppfact
#                        !  (not used for vdecaymode 0)
DRj_suppfact    0.2d0    ! (default 0.2d0) suppression of regions with jet-photon
#                        !  (and jet-jet for remnant) DR distance lower than
#                        !  DRj_suppfact
powdr_suppfact      2    ! (default 2) power controlling the DRal_suppfact and
#                        !  DRj_suppfact terms
```

Each parameter can be used to suppress a specific singular region both in the born and the remnant contribution. Notice that some default values can vary depending on the selected Z decay channel via `vdecaymode` or on the value of the input parameter `anommode` (see Section B.3.2).

To remove the phase-space regions associated with a QCD or a QED singularity of the matrix element, a set of generation cuts can be applied and controlled from the input (again, see discussion in Appendix A.1):

```
pt_j1_cut         1d0     ! (default=1d0) cut on the transverse momentum of the
#                         !  first jet (equivalent to bornktmin)
pt_a_cut          5d0     ! (default=5d0) cut on the transverse momentum of the
#                         !  photon
m_lepg_cut      0.1d0     ! (default=0.1d0/0d0 for charged lepton/neutrino decay
#                         !  mode) cut on lepton-photon invariant mass
invmass_min      40d0     ! (default=40d0/1d-3 for charged lepton/neutrino decay
#                         !  mode) cut on the charged lepton invariant mass to
#                         !  avoid photon pole (not needed in neutrino decay
#                         !  mode, unless for numerical stability)
# Frixione isolation parameters applied at generation level
smooth_dyn          1     ! (default=1) 1: use Frixione isolation as in ATLAS
#                         !              -> E^{ref}_T=smooth_eps*pT_gamma
#                         !             2: use Frixione isolation as in CMS
#                         !              -> E^{ref}_T=smooth_pt0, uncomment
#                         !                 smooth_pt0 below
smooth_eps      0.5d0     ! (default=0.5d0) photon momentum fraction (used when
#                         !  smooth_dyn=1)
#smooth_pt0        1d0     ! (default=1d0) fixed maximal pT inside the photon
#                          !  cone (used when smooth_dyn=2)
smooth_R       0.05d0     ! (default=0.05d0) maximal photon cone size
smooth_n          1d0     ! (default=1d0) exponent of the delta term in Frixione
#                         !  isolation formula
```

Note that the generation cuts are in principle not needed when the suppression factors are used, but they allow for a faster convergence by avoiding potential numerical instabilities. Anyway, it is worth reminding that, with the exception of the invariant-mass cut on the dilepton system `invmass_min` (which is preserved by the born-to-real phase-space remapping and exactly corresponds to an experimental analysis cut), their values need to be much smaller than the cuts applied at analysis level.

The code can be run by computing all relevant matrix elements (with the exception of the two-loop contribution, only available through the Matrix interface) either using MCFM or OpenLoops 2. The relevant lines of the `powheg.input` file are

```
useOL               0     ! if 1, OpenLoops is used for computing the matrix
#                         !  elements, otherwise MCFM is selected
OL_CMS              1     ! (default 1) if 1, use complex-mass scheme in OpenLoops
#                         !  (to compute EW couplings and internal masses)
```

```
OL_nf               6     ! (default 6) number of flavours to use in virtual
#                         !  amplitude computation (available values 5 and 6)
OL_onshellphoton    0     ! (default 1) if 1, alpha(0)=1/137.036 is used for the
#                         !  coupling of on-shell photons (default of OpenLoops 2)
```

By setting `useOL 1`, OpenLoops is used throughout the code. That allows the user to choose the complex-mass scheme (if `OL_CMS 1`), to include exactly the effects of the top mass in the one-loop contributions (if `OL_nf 6`), and to set the electroweak coupling for the on-shell photon to its low energy limit $\alpha(0) \sim 1/137.036$ (if `OL_onshellphoton 1`). Note that, despite the higher flexibility, using OpenLoops can significantly reduce the performance of the code (we noticed that the evaluation of the virtual contribution is about a factor of 10 faster when using MCFM). Therefore, OpenLoops should be used instead of MCFM only if any of the extra features (such as top-quark mass effects) offered by OpenLoops are relevant.

It is possible to obtain directly events with multiple weights associated to scale variation without the need to reweight a posteriori. For instance, if the `powheg.input` file contains the following lines:

```
rwl_file '-'
<initrwgt>
<weightgroup name='First-Weights'>
<weight id='11'> renscfact=1.0 facscfact=1.0 </weight>
<weight id='12'> renscfact=1.0 facscfact=2.0 </weight>
<weight id='21'> renscfact=2.0 facscfact=1.0 </weight>
<weight id='22'> renscfact=2.0 facscfact=2.0 </weight>
<weight id='1H'> renscfact=1.0 facscfact=0.5 </weight>
<weight id='H1'> renscfact=0.5 facscfact=1.0 </weight>
<weight id='HH'> renscfact=0.5 facscfact=0.5 </weight>
</initrwgt>
```

then the usual 7-point scale variation weights are produced. This feature is standard in `POWHEG-BOX-RES` and is also available for the code of all other diboson processes presented in this work.

B.3.2 Anomalous coupling specific parameters for $Z\gamma$–MiNNLO$_{\mathbf{PS}}$

The effects of anomalous triple gauge couplings (aTGCs) have been included for $Z\gamma$ production. Their impact can be studied using the following set of input parameters (see Section 4.2 for further details):

```
anomcoup            1    ! (default 0) if 1, compute aTGC contributions
hZ1                 0    ! (default 0) effective ZgZ CP-odd anomalous coupling
hZ2                 0    ! (default 0) effective ZgZ CP-odd anomalous coupling
hZ3                 0    ! (default 0) effective ZgZ CP-even anomalous coupling
hZ4                 0    ! (default 0) effective ZgZ CP-even anomalous coupling
hg1                 0    ! (default 0) effective Zgg CP-odd anomalous coupling
hg2                 0    ! (default 0) effective Zgg CP-odd anomalous coupling
hg3                 0    ! (default 0) effective Zgg CP-even anomalous coupling
hg4                 0    ! (default 0) effective Zgg CP-even anomalous coupling
anommode            0    ! (default 0) activate resonance aware sampling
#                        !  accounting for extra resonance structure allowed
#                        !  by aTGCs
```

If `anomcoup 1`, the contributions of the aTGCs are turned on in the computation of the amplitudes up to NNLO QCD. The specific values of the aTGCs can be directly set from the input card, as well. Moreover, the dedicated flag `anommode` for aTGC studies can be used to improve the integration convergence in regions of phase space where aTGCs are important, i.e typically in the tails of distributions. By setting `anommode 1`, the s-channel resonance histories associated with the aTGCs are included and accounted for in the resonance-aware sampling. Moreover, some default choices for the generation cuts or the suppression factors are used (see Section B.3.1), unless a different choice is specified from the input. We do not recommend to run without `anommode 1` when including aTGCs.

The possibility to reweight events for different combinations of aTGCs is also allowed. That can be obtained by adding to the standard reweighting list the keyword `anomcoup=1`, followed by the values of the anomalous couplings for which the amplitudes should be recomputed. In the following example the events are reweighted for different CP even $Z^*Z\gamma$ aTGCs for the central scale value (the first line simply computes the SM value by setting `anomcoup=0`):

```
rwl_file '-'
<initrwgt>
<weightgroup name='First-Weights'>
<weight id='aTGChZ30hZ40'> anomcoup=0 hZ3=0.0   hZ4=0.0   </weight>
<weight id='aTGChZ3mhZ40'> anomcoup=1 hZ3=-5e-4 hZ4=0.0   </weight>
<weight id='aTGChZ3phZ40'> anomcoup=1 hZ3=+5e-4 hZ4=0.0   </weight>
<weight id='aTGChZ30hZ4m'> anomcoup=1 hZ3=0.0   hZ4=-5e-7 </weight>
<weight id='aTGChZ30hZ4p'> anomcoup=1 hZ3=0.0   hZ4=+5e-7 </weight>
<weight id='aTGChZ3mhZ4m'> anomcoup=1 hZ3=-5e-4 hZ4=-5e-7 </weight>
<weight id='aTGChZ3phZ4p'> anomcoup=1 hZ3=+5e-4 hZ4=+5e-7 </weight>
</initrwgt>
```

In order to achieve a proper convergence of the code in the phase-space regions where the effect of aTGCs is expected to be more relevant (i.e in the tails of transverse-momentum distributions), the reweighting by different aTGCs at Powheg stage 4 should be performed only if the code has been run with `anommode 1` in all stages.

B.3.3 W^+W^-–MiNNLO$_{\rm PS}$ input parameters

When doing a WW+jet NLO+PS calculation (turning off MiNNLO$_{\rm PS}$ and MiNLO$'$) one may set the renormalization and factorization scales by changing

```
runningscales       0    ! 0 = fixed scale 2*m_W (leave this to 0 when MiN(N)LO
#                        !     is used. All scale settings is taken care of
#                        !     internally)
#                        ! 1 = M_WW, 2 = M_{T,W+}+M_{T,W-}
```

When MiNLO$'$ or MiNNLO$_{\rm PS}$ is turned on, the scale is not used and set as described in Section 3.2 and 3.3.

Several decay modes can be selected by an appropriate flag in the `powheg.input` file:

```
e+e-                0     ! (default 0) if 1, both Ws decay into electrons
mu+mu-              0     ! (default 0) if 1, both Ws decay into muons
tau+tau-            0     ! (default 0) if 1, both Ws decay into taus
e+mu-               1     ! (default 1) if 1, W+ decays into a positron and W-
#                         !  into a muon
```

```
mu+e-                0     ! (default 0) if 1, W+ decays into an anti-muon and W-
#                          !  into an electron
leptonic             0     ! (default 0) if 1, both Ws will decay into all
#                          !  combinations of leptons (e+e-,mu+mu-,e+mu-,mu+e-,
#                          !  tau+tau-,tau+e-,e+tau-,tau+mu-,mu+tau-)
leptonic_notau       0     ! (default 0) if 1, both Ws will decay into all
#                          !  combinations of muons and electrons (e+e-,mu+mu-,
#                          !  e+mu-,mu+e-)
hadronic             0     ! (default 0) if 1, both Ws will decay into all
#                          !  combinations of light quarks (u,d,s,c)
semileptonic         0     ! (default 0) if 1, one W decays into leptons and the
#                          !  other W into light quarks
semileptonic_notau   0     ! (default 0) if 1, one W decays into muons and
#                          !  electrons, and the other W into light quarks
```

If no condition is specified in the input card, the default decay channel is assumed, namely `e+mu-`.

We note that in ref. [195] it was found that closed fermion loops slow down the calculation considerably, yet provide no sizable effect in any considered distribution (within the numerical accuracy considered). Hence, as a default (`GOSAMDIR` set to `GoSamlib_nofboxes` in the `Makefile`) the code runs without including closed fermion loops, but we also provide the possibility to run with fermion loops by setting the variable `GOSAMDIR` to `GoSamlib` in the `Makefile`, and recompiling the code from scratch.

B.3.4 ZZ–MiNNLO$_{\rm PS}$ input parameters

The ZZ-MiNNLO$_{\rm PS}$ code can be run for different leptonic and hadronic final states ($\ell^+\ell^-\ell'^+\ell'^-$, $\ell^+\ell^-\ell^+\ell^-$, $\ell^+\ell^-\nu_{\ell'}\nu_{\ell'}$, $\ell^+\ell^-\nu_\ell\nu_\ell$, $q\bar{q}q'\bar{q}'$, $q\bar{q}q\bar{q}$, $\ell^+\ell^-q\bar{q}$). You can switch between the different Z-boson decays choosing among the following settings:

```
###
# Decay channels of the two Z bosons (turn only one of the following to 1)
#
# 4l-DF-ZZ - each Z boson decays to a pair of leptons of different flavour (DF)
e+e-mu+mu-           1     ! (default 0) if 1, one Z to electrons, one Z to muons
e+e-tau+tau-         0     ! (default 0) if 1, one Z to electrons, one Z to taus
mu+mu-tau+tau-       0     ! (default 0) if 1, one Z to muons, one Z to taus
4l_DF                0     ! (default 0) if 1, ZZ decay to all combinations of
#                          !  leptons of DF (e-e+mu-mu+, e-e+tau-tau+,
#                          !  mu-mu+tau-tau+; e-e+mu-mu+ up to stage 3)

# 4l-SF-ZZ - both Z boson decay to a pair of leptons of same flavour (SF)
e+e-e+e-             0     ! (default 0) if 1, both Zs decay to electrons
mu+mu-mu+mu-         0     ! (default 0) if 1, both Zs decay to muons
tau+tau-tau+tau-     0     ! (default 0) if 1, both Zs decay to taus
4l_SF                0     ! (default 0) if 1, ZZ decay to four electrons, four
#                          !  muons or four taus (e-e+e-e+, mu-mu+mu-mu+,
#                          !  tau-tau+tau-tau+; uses e-e+e-e+ up to stage 3)
4l_notau_SF          0     ! (default 0) if 1, ZZ decay to four electrons or four
#                          !  muons (e-e+e-e+, mu-mu+mu-mu+; note: uses
#                          !  e-e+e-e+ up to stage 3)

# 2l2nu-DF-ZZ - one Z boson decays to leptons the other to neutrinos (only ZZ
# contributions)
e+e-nunu_DF          0     ! (default 0) if 1, one Z to electrons, one Z to
#                          !  neutrinos
mu+mu-nunu_DF        0     ! (default 0) if 1, one Z to muons, one Z to neutrinos
```

```
tau+tau-nunu_DF     0     ! (default 0) if 1, one Z to taus, one Z to neutrinos
2l2nu_DF            0     ! (default 0) if 1, one Z to all leptons, one Z to
#                         !  neutrinos (uses e-e+nunu_DF up to stage 3)
2l2nu_notau_DF      0     ! (default 0) if 1, one Z to electrons or muons, one Z
#                         !  to neutrinos (uses e-e+nunu_DF up to stage 3)
n_neutrinos_DF      3     ! (default 3) sets the number of considered neutrino
#                         !  flavours: if 2 only neutrinos with flavour different
#                         !  of the leptons; if 3 all three neutrinos included,
#                         !  but the neutrino with SF as the lepton treated as
#                         !  DF, i.e. ZZ only

# 2l2nu-SF-ZZWW - one Z boson decays to leptons the other to neutrinos with same
# flavour (ZZ & WW)
e+e-nunu_SF         0     ! (default 0) if 1, one Z to electrons, one Z to
#                         !  electron neutrinos
mu+mu-nunu_SF       0     ! (default 0) if 1, one Z to muons, one Z to muon
#                         !  neutrinos
tau+tau-nunu_SF     0     ! (default 0) if 1, one Z to taus, one Z to tau
#                         !  neutrinos
2l2nu_SF            0     ! (default 0) if 1, one Z to all leptons, one Z to
#                         !  corresponding neutrinos (it uses e-e+nunu_SF
#                         !  up to stage 3)
2l2nu_notau_SF      0     ! (default 0) if 1, one Z to electrons or muons, one Z
#                         !  to corresponding neutrinos (it uses e-e+nunu_SF
#                         !  up to stage 3)

# 4q-DF-ZZ, 4q-SF-ZZ, 2l2q-ZZ and 2q2nu-ZZ - final states including hadronic Z
# boson decays
4q_DF               0     ! (default 0) if 1, ZZ decay to all combinations of
#                         !  light quarks of DF (uses e-e+mu-mu+ up to stage 3)
4q_SF               0     ! (default 0) if 1, ZZ decay to all combinations of
#                         !  light quarks of SF (uses e-e+e-e+ up to stage 3)
2l2q                0     ! (default 0) if 1, one Z to all charged leptons, one Z
#                         !  to all light quarks (uses e-e+mu-mu+ up to stage 3)
2q2nu               0     ! (default 0) if 1, one Z to all neutrinos, one Z to
#                         !  all light quarks (uses e-e+nunu_DF up to stage 3)
```

Moreover, if `massive_leptons 1`, the charged lepton momenta are reshuffled according to their physical masses at event generation level, as already presented for $Z\gamma$ in the subsection B.3.1. The same holds for the charm and bottom quarks, whose momenta are reshuffled at event generation level according to the mass values set from the input:

```
c_mass          1.40d0     ! (default 1.40d0) charm mass for reshuffling momentum
#                          !  at event generation level
b_mass          4.92d0     ! (default 4.92d0) bottom mass for reshuffling momentum
#                          !  at event generation level
```

As for the other diboson codes, when doing a ZZ+jet NLO+PS calculation (turning off MiNNLO$_{\rm PS}$ and MiNLO$'$) you may set the renormalization and factorization scales by changing

```
fixedscale          0    ! (default 1), if 0, use dynamical scale below (set by
#                        !  whichscale); if 1, scale is fixed to Z-boson mass
#                        !  (leave this to 0 when MiN(N)LO is used)
whichscale          1    ! (default 0), only used if fixedscale=0
#                        !  0 = sqrt(M_ZZ^2+pt_ZZ^2), 1 = M_ZZ (leave this to 1
#                        !  when MiN(N)LO is used)
```

The code computes all relevant matrix elements using OpenLoops 2 (with the exception of the two-loop contribution, which is obtained through the Matrix interface, using

`VVamp`). The settings are hard-coded to the default values. In particular, the complex-mass scheme is used (`use_cms 1` in the OpenLoops 2 interface).

As far as the electroweak (EW) and other input parameters of the calculation are concerned, the user can set them directly from the input file. In particular, the scheme for evaluating the EW parameters can be chosen, where the G_μ scheme (`ewscheme 1`) is the preferred choice.

```
ewscheme            1     ! (default 1) if 1, (complex) Gmu scheme used, if 2
#                         !  alpha scheme used
gfermi     1.16639d-5     ! Fermi constant used as input for Gmu scheme
#alpha 0.00756246890d0     ! alpha used as input for alpha scheme (used only
#                          !  when ewscheme=2)
zmass       91.1876d0     ! Z-boson mass (also input for Gmu scheme)
wmass        80.385d0     ! W-boson mass (also input for Gmu scheme)
hmass           125d0     ! Higgs mass
tmass         173.2d0     ! top-quark mass
bmass             0d0     ! bottom-quark mass (should not be changed, since
#                         !  5-flavour computation)
zwidth       2.4952d0     ! Z-boson width
wwidth       2.0854d0     ! W-boson width
twidth      1.44262d0     ! Top-quark width
hwidth      0.00407d0     ! Higgs width
```

The G_μ scheme (`ewscheme 1`) takes as inputs the complex masses m_Z (computed from `zmass` and `zwidth`) and m_W (computed from `wmass` and `wwidth`) as well as the constant G_μ (`gfermi`) to compute the EW coupling and the mixing angle, while in the $\alpha(0)$ scheme the value of the input keyword `alpha` is used as input instead of G_μ.

B.4 MiNNLO$_{\mathrm{PS}}$-specific inputs

In this section we present the parameters in the `powheg.input` specifically related to the MiNNLO$_{\mathrm{PS}}$ implementation and so valid for all processes presented in this work (and even the other MiNNLO$_{\mathrm{PS}}$ processes implemented so far). One can activate independently MiNLO$'$ or MiNNLO$_{\mathrm{PS}}$ by setting the proper parameters in the input card:

```
minlo               1    ! (default 1) if 1, activate MiNLO'
minnlo              1    ! (default 0) if 1, activate MiNNLO_PS
```

Note that, if `minnlo` is switched on, `minlo` will be turned on by default as well.

The MiNNLO$_{\mathrm{PS}}$ method achieves NNLO accuracy in the regions inclusive over QCD radiation, while it is only NLO accurate in the regions of high transverse momentum p_{T} of the colour-singlet F. As already explained in Section 3.3.2, to avoid spurious contributions one would like to switch off the genuine NNLO terms in regions where p_{T} is higher than a scale $Q \sim M$, where M is the invariant mass of F. This results in switching off the logarithms of the ratio of the two scales p_{T} and Q and the jacobian factor multiplying the NNLO corrections in the MiNNLO$_{\mathrm{PS}}$ formula. To do that, one is free to choose different prescriptions for modified logarithms. Those can be selected by setting `modlog_p` in the input card:

```
modlog_p         -1d0    ! (default -1d0; only works when minlo and/or minnlo is
```

```
#                          !  set to 1)
#                          ! if >0d0, activate modified logs and set the exponent
#                          ! if -1d0, activate piecewise modified log version 1
#                          ! if -2d0, activate piecewise modified log version 2
#                          !  piecewise modified log: log(Q/pT) for pT<Q/2;
#                          !  zero for pT>Q; smoothly interpolated in-between
#                          ! if 0d0, -log(Q/pT) * theta(Q-pT)  (not recommended)
```

By choosing a positive value for `modlog_p`, one selects the power of the ratio of $p_{\rm T}$ and Q entering the modified logarithms and the jacobian factor: the higher the power, the faster the spurious NNLO contributions will go to zero at high $p_{\rm T}$. The values `-1d0` and `-2d0` trigger the usage of a piecewise-defined modified logarithm, which is the standard logarithm below $Q/2$ and exactly zero above Q. In-between those values, two different functional forms may be chosen to smoothly interpolate between $Q/2$ and Q, keeping the modified logarithms and its derivative continuous, by setting `modlog_p -1d0` or `modlog_p -2d0` (see code and Ref. [223] for further details). By setting `modlog_p 0` (not recommended), the standard logarithm with a theta function is used, that sets both the logarithms and the jacobian factor to zero for $p_{\rm T} > Q$. The setting `modlog_p -1d0` is the currently recommended choice.

To further suppress these spurious NNLO contributions in regions of high transverse momentum $p_{\rm T}$, the hard scale Q entering logarithms in the MiNNLO$_{\rm PS}$ procedure can be chosen different from the invariant mass M of the colour-singlet final state by multiplying it by an integer factor whose value can be fixed by changing `kappaQ` in the input card:

```
kappaQ          1d0      ! (default 1d0) integer factor multiplying the hard
#                        !               scale Q entering the minnlo procedure
```

Notice that the default value is `kappaQ 1d0`. While setting `kappaQ` to values less than one reduces the effect of NNLO contributions at high $p_{\rm T}$ of the colour-singlet system, it is not advisable to use any values above one.

In order for the renormalization and factorization scales to smoothly approach the non-perturbative region at low $p_{\rm T}$, as described in Section 3.3.2 (see Eq. (3.85)), one may set `Q0` different from zero:

```
Q0                0d0    ! (default 0.) cutoff of profiled scales
```

The higher `Q0` is set the higher is the damping effect. Note that scales above about 2 GeV should not be used, and our current recommendation is to keep the damping off, unless there are instabilities related to the Sudakov, which have not been observed for any $q\bar{q}$ initiated process so far.

Finally, at high $p_{\rm T}$, where predictions are only NLO accurate, one can decide to use factorization and renormalization scales equal to the invariant mass of the colour singlet or to the $p_{\rm T}$ of the colour singlet. The latter would be more appropriate for energetic F production. One can choose between the two options by setting `largeptscales`:

```
largeptscales       1     ! (default 1) if 0, at large pt, use muR=muF=~Q in
#                         !  fixed-order part; if 1, at large pt, use muR=muF=pt
#                         !  in the fixed-order part
```

B.5 Running modes for MiNNLO$_{\rm PS}$

We provide different ways of running the MiNNLO$_{\rm PS}$ code for a diboson process. This code optimization has been first developed for ZZ production in Ref. [220], and then made available to all processes. The user can decide whether the two-loop amplitudes, needed to achieve NNLO accurate predictions for inclusive F observables, should be included at different Powheg stages or just via event reweighting, but also whether to compute MiNLO$'$ via event reweighting. These optimization options are crucial for the computation of the two-loop contribution for massive diboson processes. On the other hand, since the two-loop amplitudes for $Z\gamma$ production are not particularly time-consuming, such optimization for the two-loop amplitudes are not really needed in that case, but it might still be useful to obtain MiNLO$'$ via event reweighting, instead of doing a separate run. These options can be selected via `run_mode` in the `powheg.input`:

```
run_mode            1    ! (default 1) Select running mode for treatment of
#                        !  2-loop and D-terms.  This provides the possibility to
#                        !  speed up the code by either including 2-loop through
#                        !  reweighting at event generation level, or evaluating
#                        !  2-loop *only* from stage 2 on, or computing MiNLO
#                        !  through reweighting!
#                        !  1 - standard running mode (2-loop fully included)
#                        !  2 - compute MiNLO by turning off D-terms
#                        !  3 - skip 2-loop computation at stage 1 only
#                        !  4 - switch off 2-loop computation entirely (and
#                        !      include it via reweighting, see below)
```

By setting `run_mode 1`, the code will run normally and the two-loop contribution will be included in all stages of the calculation. With option `2`, the NNLO corrections of the MiNNLO$_{\rm PS}$ formula are set to zero, so that the code effectively generates MiNLO$'$ results, i.e. it runs exactly as if the `minlo` parameter but not the `minnlo` parameter was on. In run mode `3`, the Powheg integration grids are generated without the evaluation of the 2-loop amplitudes, which are included only starting from Powheg stage 2. Finally, run mode `4` turns off the 2-loop contributions entering the NNLO terms of MiNNLO$_{\rm PS}$, which allows the user to include them only at the very last Powheg stage through event reweighting. In particular, besides setting the input flag `run_mode 4`, this is obtained by resetting the keyword `run_mode` to `1` in the reweighting list, for instance:

```
<initrwgt>
<weightgroup name='First-Weights'>
<weight id='11'> renscfact=1.0 facscfact=1.0 run_mode=1 </weight>
<weight id='12'> renscfact=1.0 facscfact=2.0 run_mode=1 </weight>
<weight id='21'> renscfact=2.0 facscfact=1.0 run_mode=1 </weight>
<weight id='22'> renscfact=2.0 facscfact=2.0 run_mode=1 </weight>
<weight id='1H'> renscfact=1.0 facscfact=0.5 run_mode=1 </weight>
<weight id='H1'> renscfact=0.5 facscfact=1.0 run_mode=1 </weight>
<weight id='HH'> renscfact=0.5 facscfact=0.5 run_mode=1 </weight>
</initrwgt>
```

In addition (or alternatively) one can obtain MiNLO$'$ via reweighting by resetting `run_mode` to `2` in the reweighting block; for instance add:

```
<weight id='11-MiNLO'> renscfact=1.0 facscfact=1.0 run_mode=2 </weight>
<weight id='12-MiNLO'> renscfact=1.0 facscfact=2.0 run_mode=2 </weight>
```

```
<weight id='21-MiNLO'> renscfact=2.0 facscfact=1.0 run_mode=2 </weight>
<weight id='22-MiNLO'> renscfact=2.0 facscfact=2.0 run_mode=2 </weight>
<weight id='1H-MiNLO'> renscfact=1.0 facscfact=0.5 run_mode=2 </weight>
<weight id='H1-MiNLO'> renscfact=0.5 facscfact=1.0 run_mode=2 </weight>
<weight id='HH-MiNLO'> renscfact=0.5 facscfact=0.5 run_mode=2 </weight>
```

As described in Section 4.3.1, for the W^+W^- code a further optimization option is available to the user, that can decide whether the two-loop amplitudes should be computed using their time-consuming, but exact expressions, which rely on the public code `VVAMP`, or making use of the interpolation approach, which is recommend to substantially speed up the runtime. The two possibilities are controlled by the input parameter `use_interpolator`:

```
use_interpolator    1    ! (default 0) if 1, activates interpolator instead
#                        !  of MATRIX+VVamp
```

In case the user wants to make use of the interpolator, the set of precomputed grids required for this approach need to be separately downloaded and unpacked, as described at the end of Section B.1.

Bibliography

[1] R. K. Ellis, W. J. Stirling and B. R. Webber, QCD and collider physics, vol. 8. Cambridge University Press, 2, 2011, 10.1017/CBO9780511628788.

[2] Particle Data Group collaboration, P. Zyla et al., Review of Particle Physics, PTEP **2020** (2020) 083C01.

[3] C. A. Baker et al., An Improved experimental limit on the electric dipole moment of the neutron, Phys. Rev. Lett. **97** (2006) 131801, [hep-ex/0602020].

[4] T. D. Cohen, Does chiral perturbation theory rule out QCD-based solutions to the strong CP problem?, Phys. Rev. D **99** (2019) 094007, [1811.04833].

[5] P. W. Higgs, Broken Symmetries and the Masses of Gauge Bosons, Phys. Rev. Lett. **13** (1964) 508–509.

[6] J. Goldstone, Field Theories with Superconductor Solutions, Nuovo Cim. **19** (1961) 154–164.

[7] CMS collaboration, S. Chatrchyan et al., Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC, Phys. Lett. B **716** (2012) 30–61, [1207.7235].

[8] ATLAS collaboration, G. Aad et al., Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, Phys. Lett. B **716** (2012) 1–29, [1207.7214].

[9] M. Kobayashi and T. Maskawa, CP Violation in the Renormalizable Theory of Weak Interaction, Prog. Theor. Phys. **49** (1973) 652–657.

[10] E. Fermi, Tentativo di una teoria dell'emissione dei raggi beta, Ric. Sci. **4** (1933) 491–495.

[11] C. D. Anderson and H. L. Anderson, Unraveling the Particle Content of Cosmic Rays, in 1st International Symposium on the History of Particle Physics, 1980.

[12] L. M. Brown, The idea of the neutrino, Phys. Today **31N9** (1978) 23–28.

[13] G. Zweig, An SU(3) model for strong interaction symmetry and its breaking. Version 1, .

[14] C. G. Callan, Jr. and D. J. Gross, High-energy electroproduction and the constitution of the electric current, Phys. Rev. Lett. **22** (1969) 156–159.

[15] D. J. Gross and F. Wilczek, Ultraviolet Behavior of Nonabelian Gauge Theories, Phys. Rev. Lett. **30** (1973) 1343–1346.

[16] H. D. Politzer, Reliable Perturbative Results for Strong Interactions?, Phys. Rev. Lett. **30** (1973) 1346–1349.

[17] D. d'Enterria et al., The strong coupling constant: State of the art and the decade ahead, 2203.08271.

[18] Z.-G. Zhao, R-values from low-energy e+ e- annihilation, Int. J. Mod. Phys. A **15** (2000) 3739–3769.

[19] JADE collaboration, W. Bartel et al., Observation of Planar Three Jet Events in e+ e- Annihilation and Evidence for Gluon Bremsstrahlung, Phys. Lett. B **91** (1980) 142–147.

[20] L3 collaboration, B. Adeva et al., A Test of QCD based on four jet events from Z0 decays, Phys. Lett. B **248** (1990) 227–234.

[21] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, Experimental Test of Parity Conservation in β Decay, Phys. Rev. **105** (1957) 1413–1414.

[22] S. L. Glashow, The renormalizability of vector meson interactions, Nucl. Phys. **10** (1959) 107–117.

[23] A. Salam and J. C. Ward, Weak and electromagnetic interactions, Nuovo Cim. **11** (1959) 568–577.

[24] S. Weinberg, A Model of Leptons, Phys. Rev. Lett. **19** (1967) 1264–1266.

[25] G. 't Hooft, Renormalizable Lagrangians for Massive Yang-Mills Fields, Nucl. Phys. B **35** (1971) 167–188.

[26] G. 't Hooft and M. J. G. Veltman, Regularization and Renormalization of Gauge Fields, Nucl. Phys. B **44** (1972) 189–213.

[27] Gargamelle Neutrino collaboration, F. J. Hasert et al., Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment, Phys. Lett. B **46** (1973) 138–140.

[28] UA1 collaboration, G. Arnison et al., Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV, Phys. Lett. B **122** (1983) 103–116.

[29] UA2 collaboration, M. Banner et al., Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider, Phys. Lett. B **122** (1983) 476–485.

[30] UA1 collaboration, G. Arnison et al., Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c**2 at the CERN SPS Collider, Phys. Lett. B **126** (1983) 398–410.

[31] E598 collaboration, J. J. Aubert et al., Experimental Observation of a Heavy Particle J, Phys. Rev. Lett. **33** (1974) 1404–1406.

[32] SLAC-SP-017 collaboration, J. E. Augustin et al., Discovery of a Narrow Resonance in e^+e^- Annihilation, Phys. Rev. Lett. **33** (1974) 1406–1408.

[33] S. L. Glashow, J. Iliopoulos and L. Maiani, Weak Interactions with Lepton-Hadron Symmetry, Phys. Rev. D **2** (1970) 1285–1292.

[34] KTeV collaboration, A. Alavi-Harati et al., Observation of direct CP violation in $K_{S,L} \to \pi\pi$ decays, Phys. Rev. Lett. **83** (1999) 22–27, [hep-ex/9905060].

[35] NA48 collaboration, V. Fanti et al., A New measurement of direct CP violation in two pion decays of the neutral kaon, Phys. Lett. B **465** (1999) 335–348, [hep-ex/9909022].

[36] S. W. Herb et al., Observation of a Dimuon Resonance at 9.5-GeV in 400-GeV Proton-Nucleus Collisions, Phys. Rev. Lett. **39** (1977) 252–255.

[37] M. L. Perl et al., Evidence for Anomalous Lepton Production in e+ - e- Annihilation, Phys. Rev. Lett. **35** (1975) 1489–1492.

[38] Belle collaboration, K. Abe et al., Observation of large CP violation in the neutral B meson system, Phys. Rev. Lett. **87** (2001) 091802, [hep-ex/0107061].

[39] BaBar collaboration, B. Aubert et al., Observation of CP violation in the B^0 meson system, Phys. Rev. Lett. **87** (2001) 091801, [hep-ex/0107013].

[40] CDF collaboration, F. Abe et al., Evidence for top quark production in $\bar{p}p$ collisions at $\sqrt{s} = 1.8$ TeV, Phys. Rev. Lett. **73** (1994) 225–231, [hep-ex/9405005].

[41] DONUT collaboration, K. Kodama et al., Observation of tau neutrino interactions, Phys. Lett. B **504** (2001) 218–224, [hep-ex/0012035].

[42] B. W. Lee, C. Quigg and H. B. Thacker, Weak Interactions at Very High-Energies: The Role of the Higgs Boson Mass, Phys. Rev. D **16** (1977) 1519.

[43] F. Zwicky, Die Rotverschiebung von extragalaktischen Nebeln, Helv. Phys. Acta **6** (1933) 110–127.

[44] A. H. Guth, The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems, Phys. Rev. D **23** (1981) 347–356.

[45] A. D. Sakharov, Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe, Pisma Zh. Eksp. Teor. Fiz. **5** (1967) 32–35.

[46] E. Akhmedov, Quantum mechanics aspects and subtleties of neutrino oscillations, in International Conference on History of the Neutrino: 1930-2018, 1, 2019. 1901.05232.

[47] BaBar collaboration, J. P. Lees et al., Evidence for an excess of $\bar{B} \to D^{(*)}\tau^-\bar{\nu}_\tau$ decays, Phys. Rev. Lett. **109** (2012) 101802, [1205.5442].

[48] LHCb collaboration, R. Aaij et al., Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \to D^{*+}\tau^-\bar{\nu}_\tau)/\mathcal{B}(\bar{B}^0 \to D^{*+}\mu^-\bar{\nu}_\mu)$, Phys. Rev. Lett. **115** (2015) 111803, [1506.08614].

[49] Belle collaboration, M. Huschle et al., Measurement of the branching ratio of $\bar{B} \to D^{(*)}\tau^-\bar{\nu}_\tau$ relative to $\bar{B} \to D^{(*)}\ell^-\bar{\nu}_\ell$ decays with hadronic tagging at Belle, Phys. Rev. D **92** (2015) 072014, [1507.03233].

[50] Muon g-2 collaboration, B. Abi et al., Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm, Phys. Rev. Lett. **126** (2021) 141801, [2104.03281].

[51] CDF collaboration, T. Aaltonen et al., High-precision measurement of the W boson mass with the CDF II detector, Science **376** (2022) 170–176.

[52] J. C. Collins, D. E. Soper and G. F. Sterman, Factorization for Short Distance Hadron - Hadron Scattering, Nucl. Phys. B **261** (1985) 104–142.

[53] J. C. Collins, D. E. Soper and G. F. Sterman, Soft Gluons and Factorization, Nucl. Phys. B **308** (1988) 833–856.

[54] S. Ferrario Ravasio, G. Limatola and P. Nason, Infrared renormalons in kinematic distributions for hadron collider processes, JHEP **06** (2021) 018, [2011.14114].

[55] F. Caola, S. F. Ravasio, G. Limatola, K. Melnikov and P. Nason, On linear power corrections in certain collider observables, JHEP **01** (2022) 093, [2108.08897].

[56] G. C. Wick, The evaluation of the collision matrix, Phys. Rev. **80** (Oct, 1950) 268–272.

[57] H. Lehmann, K. Symanzik and W. Zimmermann, Zur formulierung quantisierter feldtheorien, Il Nuovo Cimento (1955-1965) **I** (1955) 205–225.

[58] G. 't Hooft and M. Veltman, Regularization and renormalization of gauge fields, Nuclear Physics B **44** (1972) 189–213.

[59] H. D. Politzer, Reliable perturbative results for strong interactions?, Phys. Rev. Lett. **30** (Jun, 1973) 1346–1349.

[60] D. J. Gross and F. Wilczek, Asymptotically free gauge theories. i, Phys. Rev. D **8** (Nov, 1973) 3633–3652.

[61] S. Catani and M. H. Seymour, A General algorithm for calculating jet cross-sections in NLO QCD, Nucl. Phys. B **485** (1997) 291–419, [hep-ph/9605323].

[62] T. Kinoshita, Mass singularities of Feynman amplitudes, J. Math. Phys. **3** (1962) 650–677.

[63] T. D. Lee and M. Nauenberg, Degenerate systems and mass singularities, Phys. Rev. **133** (Mar, 1964) B1549–B1562.

[64] J. C. Collins, D. E. Soper and G. F. Sterman, Factorization of Hard Processes in QCD, Adv. Ser. Direct. High Energy Phys. **5** (1989) 1–91, [hep-ph/0409313].

[65] J. C. Collins, D. E. Soper and G. F. Sterman, Transverse Momentum Distribution in Drell-Yan Pair and W and Z Boson Production, Nucl. Phys. B **250** (1985) 199–224.

[66] S. Catani, M. Ciafaloni and F. Hautmann, GLUON CONTRIBUTIONS TO SMALL x HEAVY FLAVOR PRODUCTION, Phys. Lett. B **242** (1990) 97–102.

[67] S. Catani, M. Ciafaloni and F. Hautmann, High-energy factorization and small x heavy flavor production, Nucl. Phys. B **366** (1991) 135–188.

[68] M. Constantinou et al., Parton distributions and lattice-QCD calculations: Toward 3D structure, Prog. Part. Nucl. Phys. **121** (2021) 103908, [2006.08636].

[69] L. Del Debbio, T. Giani and C. J. Monahan, Notes on lattice observables for parton distributions: nongauge theories, JHEP **09** (2020) 021, [2007.02131].

[70] M. Constantinou, The x-dependence of hadronic parton distributions: A review on the progress of lattice QCD, Eur. Phys. J. A **57** (2021) 77, [2010.02445].

[71] Y. L. Dokshitzer, Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics., Sov. Phys. JETP **46** (1977) 641–653.

[72] V. N. Gribov and L. N. Lipatov, Deep inelastic e p scattering in perturbation theory, Sov. J. Nucl. Phys. **15** (1972) 438–450.

[73] G. Altarelli and G. Parisi, Asymptotic Freedom in Parton Language, Nucl. Phys. B **126** (1977) 298–318.

[74] W. Furmanski and R. Petronzio, Lepton - Hadron Processes Beyond Leading Order in Quantum Chromodynamics, Z. Phys. C **11** (1982) 293.

[75] S. Catani and M. Grazzini, An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC, Phys. Rev. Lett. **98** (2007) 222002, [hep-ph/0703012].

[76] R. Bonciani, S. Catani, M. Grazzini, H. Sargsyan and A. Torre, The q_T subtraction method for top quark production at hadron colliders, Eur. Phys. J. C **75** (2015) 581, [1508.03585].

[77] R. Boughezal, C. Focke, X. Liu and F. Petriello, W-boson production in association with a

jet at next-to-next-to-leading order in perturbative QCD, Phys. Rev. Lett. **115** (2015) 062002, [1504.02131].

[78] R. Boughezal, X. Liu and F. Petriello, N-jettiness soft function at next-to-next-to-leading order, Phys. Rev. D **91** (2015) 094035, [1504.02540].

[79] J. Gaunt, M. Stahlhofen, F. J. Tackmann and J. R. Walsh, N-jettiness Subtractions for NNLO QCD Calculations, JHEP **09** (2015) 058, [1505.04794].

[80] I. W. Stewart, F. J. Tackmann and W. J. Waalewijn, N-Jettiness: An Inclusive Event Shape to Veto Jets, Phys. Rev. Lett. **105** (2010) 092002, [1004.2489].

[81] M. Grazzini, S. Kallweit and M. Wiesemann, Fully differential NNLO computations with MATRIX, Eur. Phys. J. **C78** (2018) 537, [1711.06631].

[82] MATRIX is the abbreviation of Munich Automates qT subtraction and Resummation to Integrate X-sections by M. Grazzini, S. Kallweit, M. Wiesemann, http://matrix.hepforge.org.

[83] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. Giele, X. Liu, F. Petriello and C. Williams, Color singlet production at NNLO in MCFM, Eur. Phys. J. C **77** (2017) 7, [1605.08011].

[84] J. M. Campbell, R. K. Ellis and S. Seth, Non-local slicing approaches for NNLO QCD in MCFM, 2202.07738.

[85] S. Alioli, C. W. Bauer, C. Berggren, F. J. Tackmann, J. R. Walsh and S. Zuberi, Matching Fully Differential NNLO Calculations and Parton Showers, JHEP **06** (2014) 089, [1311.0286].

[86] S. Frixione, P. Nason and C. Oleari, Matching NLO QCD computations with Parton Shower simulations: the POWHEG method, JHEP **11** (2007) 070, [0709.2092].

[87] S. Frixione, Z. Kunszt and A. Signer, Three jet cross-sections to next-to-leading order, Nucl. Phys. B **467** (1996) 399–442, [hep-ph/9512328].

[88] S. Frixione, A General approach to jet cross-sections in QCD, Nucl. Phys. B **507** (1997) 295–314, [hep-ph/9706545].

[89] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and T. A. Morgan, Precise QCD predictions for the production of a Z boson in association with a hadronic jet, Phys. Rev. Lett. **117** (2016) 022001, [1507.02850].

[90] J. Currie, T. Gehrmann and J. Niehues, Precise QCD predictions for the production of dijet final states in deep inelastic scattering, Phys. Rev. Lett. **117** (2016) 042001, [1606.03991].

[91] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and J. Pires, Precise predictions for dijet production at the LHC, Phys. Rev. Lett. **119** (2017) 152001, [1705.10271].

[92] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss and T. A. Morgan, The NNLO QCD corrections to Z boson production at large transverse momentum, JHEP **07** (2016) 133, [1605.04295].

[93] T. Gehrmann et al., Jet cross sections and transverse momentum distributions with NNLOJET, PoS **RADCOR2017** (2018) 074, [1801.06415].

[94] M. Czakon, A novel subtraction scheme for double-real radiation at NNLO, Phys. Lett. B **693** (2010) 259–268, [1005.0274].

[95] M. Cacciari, F. A. Dreyer, A. Karlberg, G. P. Salam and G. Zanderighi, Fully Differential Vector-Boson-Fusion Higgs Production at Next-to-Next-to-Leading Order, Phys. Rev. Lett. **115** (2015) 082002, [1506.02660].

[96] F. Caola, K. Melnikov and R. Röntsch, Nested soft-collinear subtractions in NNLO QCD computations, Eur. Phys. J. C **77** (2017) 248, [1702.01352].

[97] F. Caola, G. Luisoni, K. Melnikov and R. Röntsch, NNLO QCD corrections to associated WH production and $H \to b\bar{b}$ decay, Phys. Rev. D **97** (2018) 074022, [1712.06954].

[98] V. Del Duca, C. Duhr, A. Kardos, G. Somogyi and Z. Trócsányi, Three-Jet Production in Electron-Positron Collisions at Next-to-Next-to-Leading Order Accuracy, Phys. Rev. Lett. **117** (2016) 152004, [1603.08927].

[99] V. Del Duca, C. Duhr, A. Kardos, G. Somogyi, Z. Szőr, Z. Trócsányi and Z. Tulipánt, Jet production in the CoLoRFulNNLO method: event shapes in electron-positron collisions, Phys. Rev. D **94** (2016) 074019, [1606.03453].

[100] Z. Tulipánt, A. Kardos and G. Somogyi, Energy–energy correlation in electron–positron annihilation at NNLL + NNLO accuracy, Eur. Phys. J. C **77** (2017) 749, [1708.04093].

[101] F. Herzog, Geometric IR subtraction for final state real radiation, JHEP **08** (2018) 006, [1804.07949].

[102] L. Magnea, E. Maina, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, Local analytic sector subtraction at NNLO, JHEP **12** (2018) 107, [1806.09570].

[103] M. Beneke, Renormalons, Phys. Rept. **317** (1999) 1–142, [hep-ph/9807443].

[104] A. Banfi, G. P. Salam and G. Zanderighi, Principles of general final-state resummation and automated implementation, JHEP **03** (2005) 073, [hep-ph/0407286].

[105] W. Bizon, P. F. Monni, E. Re, L. Rottoli and P. Torrielli, Momentum-space resummation for transverse observables and the Higgs $p_\perp$ at N^3LL+NNLO, JHEP **02** (2018) 108, [1705.09127].

[106] T. Becher, A. Broggio and A. Ferroglia, Introduction to Soft-Collinear Effective Theory, vol. 896. Springer, 2015, 10.1007/978-3-319-14848-9.

[107] A. Banfi, H. McAslan, P. F. Monni and G. Zanderighi, A general method for the resummation of event-shape distributions in e^+e^- annihilation, JHEP **05** (2015) 102, [1412.2126].

[108] M. Dasgupta, F. A. Dreyer, K. Hamilton, P. F. Monni and G. P. Salam, Logarithmic accuracy of parton showers: a fixed-order study, JHEP **09** (2018) 033, [1805.09327].

[109] M. Dasgupta, F. A. Dreyer, K. Hamilton, P. F. Monni, G. P. Salam and G. Soyez, Parton showers beyond leading logarithmic accuracy, Phys. Rev. Lett. **125** (2020) 052002, [2002.11114].

[110] M. van Beekveld, S. Ferrario Ravasio, G. P. Salam, A. Soto-Ontoso, G. Soyez and R. Verheyen, PanScales parton showers for hadron collisions: formulation and fixed-order studies, 2205.02237.

[111] M. van Beekveld, S. F. Ravasio, K. Hamilton, G. P. Salam, A. Soto-Ontoso, G. Soyez and R. Verheyen, PanScales showers for hadron collisions: all-order validation, 2207.09467.

[112] S. Höche, Introduction to parton-shower event generators, in Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders, pp. 235–295, 2015. 1411.4085. DOI.

[113] S. Amoroso et al., Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report, in 11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches, 3, 2020. 2003.01700.

[114] S. Höche and S. Prestel, Triple collinear emissions in parton showers, Phys. Rev. D **96** (2017) 074017, [1705.00742].

[115] J. C. Collins, Spin Correlations in Monte Carlo Event Generators, Nucl. Phys. B **304** (1988) 794–804.

[116] I. G. Knowles, Spin Correlations in Parton - Parton Scattering, Nucl. Phys. B **310** (1988) 571–588.

[117] I. G. Knowles, Angular Correlations in QCD, Nucl. Phys. B **304** (1988) 767–793.

[118] I. G. Knowles, A Linear Algorithm for Calculating Spin Correlations in Hadronic Collisions, Comput. Phys. Commun. **58** (1990) 271–284.

[119] J. Bellm et al., Herwig 7.2 release note, Eur. Phys. J. C **80** (2020) 452, [1912.06509].

[120] G. C. Fox and S. Wolfram, A Model for Parton Showers in QCD, Nucl. Phys. B **168** (1980) 285–295.

[121] T. Sjostrand and P. Z. Skands, Transverse-momentum-ordered showers and interleaved multiple interactions, Eur. Phys. J. C **39** (2005) 129–154, [hep-ph/0408302].

[122] S. Gieseke, P. Stephens and B. Webber, New formalism for QCD parton showers, JHEP **12** (2003) 045, [hep-ph/0310083].

[123] T. Sjostrand, S. Mrenna and P. Z. Skands, PYTHIA 6.4 Physics and Manual, JHEP **05** (2006) 026, [hep-ph/0603175].

[124] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen and P. Z. Skands, An introduction to PYTHIA 8.2, Comput. Phys. Commun. **191** (2015) 159–177, [1410.3012].

[125] G. Gustafson and U. Pettersson, Dipole Formulation of QCD Cascades, Nucl. Phys. B **306** (1988) 746–758.

[126] W. T. Giele, D. A. Kosower and P. Z. Skands, A simple shower and matching algorithm, Phys. Rev. D **78** (2008) 014026, [0707.3652].

[127] L. Lonnblad, ARIADNE version 4: A Program for simulation of QCD cascades implementing the color dipole model, Comput. Phys. Commun. **71** (1992) 15–31.

[128] S. Schumann and F. Krauss, A Parton shower algorithm based on Catani-Seymour dipole factorisation, JHEP **03** (2008) 038, [0709.1027].

[129] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert and J. Winter, Event generation with SHERPA 1.1, JHEP **02** (2009) 007, [0811.4622].

[130] T. Sjostrand, A Model for Initial State Parton Showers, Phys. Lett. B **157** (1985) 321–325.

[131] J. Campbell, J. Huston and F. Krauss, The Black Book of Quantum Chromodynamics: A Primer for the LHC Era. Oxford University Press, 12, 2017.

[132] Z. Nagy and D. E. Soper, Evolution of parton showers and parton distribution functions, Phys. Rev. D **102** (2020) 014025, [2002.04125].

[133] M. Bahr et al., Herwig++ Physics and Manual, Eur. Phys. J. C **58** (2008) 639–707, [0803.0883].

[134] S. Catani, B. R. Webber and G. Marchesini, QCD coherent branching and semiinclusive processes at large x, Nucl. Phys. B **349** (1991) 635–654.

[135] P. Nason, A New method for combining NLO QCD with shower Monte Carlo algorithms, JHEP **11** (2004) 040, [hep-ph/0409146].

[136] P. Nason, Shower Monte Carlo programs, Frascati Phys. Ser. **49** (2009) 34–67.

[137] Z. Nagy and D. E. Soper, A parton shower based on factorization of the quantum density matrix, JHEP **06** (2014) 097, [1401.6364].

[138] N. Fischer, S. Prestel, M. Ritzmann and P. Skands, Vincia for Hadron Colliders, Eur. Phys. J. C **76** (2016) 589, [1605.06142].

[139] S. Höche and S. Prestel, The midpoint between dipole and parton showers, Eur. Phys. J. C **75** (2015) 461, [1506.05057].

[140] P. Richardson and S. Webster, Spin Correlations in Parton Shower Simulations, Eur. Phys. J. C **80** (2020) 83, [1807.01955].

[141] A. Karlberg, G. P. Salam, L. Scyboz and R. Verheyen, Spin correlations in final-state parton showers and jet observables, Eur. Phys. J. C **81** (2021) 681, [2103.16526].

[142] K. Hamilton, A. Karlberg, G. P. Salam, L. Scyboz and R. Verheyen, Soft spin correlations in final-state parton showers, JHEP **03** (2022) 193, [2111.01161].

[143] S. Plätzer, Summing Large-N Towers in Colour Flow Evolution, Eur. Phys. J. C **74** (2014) 2907, [1312.2448].

[144] K. Hamilton, R. Medves, G. P. Salam, L. Scyboz and G. Soyez, Colour and logarithmic accuracy in final-state parton showers, JHEP **03** (2021) 041, [2011.10054].

[145] M. De Angelis, J. R. Forshaw and S. Plätzer, Resummation and Simulation of Soft Gluon Effects beyond Leading Color, Phys. Rev. Lett. **126** (2021) 112001, [2007.09648].

[146] S. Höche and D. Reichelt, Numerical resummation at subleading color in the strongly ordered soft gluon limit, Phys. Rev. D **104** (2021) 034006, [2001.11492].

[147] J. R. Forshaw, J. Holguin and S. Plätzer, Rings and strings: a basis for understanding subleading colour and QCD coherence beyond the two-jet limit, JHEP **05** (2022) 190, [2112.13124].

[148] M. Bengtsson and T. Sjostrand, Coherent Parton Showers Versus Matrix Elements: Implications of PETRA - PEP Data, Phys. Lett. B **185** (1987) 435.

[149] P. Nason and G. Ridolfi, A Positive-weight next-to-leading-order Monte Carlo for Z pair hadroproduction, JHEP **08** (2006) 077, [hep-ph/0606275].

[150] W. T. Giele, D. A. Kosower and P. Z. Skands, Higher-Order Corrections to Timelike Jets, Phys. Rev. D **84** (2011) 054003, [1102.2126].

[151] P. Nason and G. P. Salam, Multiplicative-accumulative matching of NLO calculations with parton showers, JHEP **01** (2022) 067, [2111.03553].

[152] S. Frixione and B. R. Webber, Matching NLO QCD computations and parton shower simulations, JHEP **06** (2002) 029, [hep-ph/0204244].

[153] K. Danziger, S. Höche and F. Siegert, Reducing negative weights in Monte Carlo event generation with Sherpa, 2110.15211.

[154] R. Frederix, S. Frixione, S. Prestel and P. Torrielli, On the reduction of negative weights in MC@NLO-type matching procedures, JHEP **07** (2020) 238, [2002.12716].

[155] T. Ježo and P. Nason, On the Treatment of Resonances in Next-to-Leading Order Calculations Matched to a Parton Shower, JHEP **12** (2015) 065, [1509.09071].

[156] S. Jadach, W. Płaczek, S. Sapeta, A. Siódmok and M. Skrzypek, Matching NLO QCD with parton shower in Monte Carlo scheme — the KrkNLO method, JHEP **10** (2015) 052, [1503.06849].

[157] S. Jadach, A. Kusina, W. Placzek, M. Skrzypek and M. Slawinska, Inclusion of the QCD next-to-leading order corrections in the quark-gluon Monte Carlo shower, Phys. Rev. D **87** (2013) 034029, [1103.5015].

[158] M. Backes, A. Butter, T. Plehn and R. Winterhalder, How to GAN Event Unweighting, SciPost Phys. **10** (2021) 089, [2012.07873].

[159] N. Lavesson and L. Lonnblad, Merging parton showers and matrix elements: Back to basics, JHEP **04** (2008) 085, [0712.2966].

[160] M. L. Mangano, M. Moretti and R. Pittau, Multijet matrix elements and shower evolution in hadronic collisions: $Wb\bar{b}$ + n jets as a case study, Nucl. Phys. B **632** (2002) 343–362, [hep-ph/0108069].

[161] S. Catani, F. Krauss, R. Kuhn and B. R. Webber, QCD matrix elements + parton showers, JHEP **11** (2001) 063, [hep-ph/0109231].

[162] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock and B. R. Webber, New clustering algorithm for multi - jet cross-sections in e+ e- annihilation, Phys. Lett. B **269** (1991) 432–438.

[163] A. Schalicke and F. Krauss, Implementing the ME+PS merging algorithm, JHEP **07** (2005) 018, [hep-ph/0503281].

[164] J. Alwall et al., Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions, Eur. Phys. J. C **53** (2008) 473–500, [0706.2569].

[165] K. Hamilton, P. Richardson and J. Tully, A Modified CKKW matrix element merging approach to angular-ordered parton showers, JHEP **11** (2009) 038, [0905.3072].

[166] L. Lonnblad, Correcting the color dipole cascade model with fixed order matrix elements, JHEP **05** (2002) 046, [hep-ph/0112284].

[167] L. Lonnblad and S. Prestel, Unitarising Matrix Element + Parton Shower merging, JHEP **02** (2013) 094, [1211.4827].

[168] S. Mrenna and P. Richardson, Matching matrix elements and parton showers with HERWIG and PYTHIA, JHEP **05** (2004) 040, [hep-ph/0312274].

[169] H. Brooks and C. T. Preuss, Efficient multi-jet merging with the Vincia sector shower, Comput. Phys. Commun. **264** (2021) 107985, [2008.09468].

[170] R. Frederix and S. Frixione, Merging meets matching in MC@NLO, JHEP **12** (2012) 061, [1209.6215].

[171] S. Hoeche, F. Krauss, M. Schonherr and F. Siegert, QCD matrix elements + parton showers: The NLO case, JHEP **04** (2013) 027, [1207.5030].

[172] L. Lönnblad and S. Prestel, Merging Multi-leg NLO Matrix Elements with Parton Showers, JHEP **03** (2013) 166, [1211.7278].

[173] S. Alioli, C. W. Bauer, C. J. Berggren, A. Hornig, F. J. Tackmann, C. K. Vermilion, J. R. Walsh and S. Zuberi, Combining Higher-Order Resummation with Multiple NLO Calculations and Parton Showers in GENEVA, JHEP **09** (2013) 120, [1211.7049].

[174] K. Hamilton, P. Nason and G. Zanderighi, MINLO: Multi-Scale Improved NLO, JHEP **10** (2012) 155, [1206.3572].

[175] K. Hamilton, P. Nason, C. Oleari and G. Zanderighi, Merging H/W/Z + 0 and 1 jet at NLO with no merging scale: a path to parton shower + NNLO matching, JHEP **05** (2013) 082, [1212.4504].

[176] R. Frederix and K. Hamilton, Extending the MINLO method, JHEP **05** (2016) 042, [1512.02663].

[177] S. Alioli, P. Nason, C. Oleari and E. Re, A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX, JHEP **06** (2010) 043, [1002.2581].

[178] J. M. Campbell, R. K. Ellis, P. Nason and E. Re, Top-Pair Production and Decay at NLO Matched with Parton Showers, JHEP **04** (2015) 114, [1412.1828].

[179] T. Ježo, J. M. Lindert, P. Nason, C. Oleari and S. Pozzorini, An NLO+PS generator for $t\bar{t}$ and Wt production and decay including non-resonant and interference effects, Eur. Phys. J. C **76** (2016) 691, [1607.04538].

[180] E. Boos et al., Generic User Process Interface for Event Generators, in 2nd Les Houches Workshop on Physics at TeV Colliders, 9, 2001. hep-ph/0109068.

[181] P. Nason, MINT: A Computer program for adaptive Monte Carlo integration and generation of unweighted distributions, 0709.2085.

[182] S. Kawabata, A New version of the multidimensional integration and event generation package BASES/SPRING, Comput. Phys. Commun. **88** (1995) 309–326.

[183] G. P. Lepage, VEGAS: AN ADAPTIVE MULTIDIMENSIONAL INTEGRATION PROGRAM, 3, 1980.

[184] S. Alioli, P. Nason, C. Oleari and E. Re, Vector boson plus one jet production in POWHEG, JHEP **01** (2011) 095, [1009.5594].

[185] T. Melia, P. Nason, R. Rontsch and G. Zanderighi, W^+W^+ plus dijet production in the POWHEGBOX, Eur. Phys. J. C **71** (2011) 1670, [1102.4846].

[186] J. Collins, Foundations of perturbative QCD, vol. 32. Cambridge University Press, 11, 2013.

[187] S. Catani, D. de Florian and M. Grazzini, Universality of non-leading logarithmic contributions in transverse momentum distributions, Nucl. Phys. B **596** (2001) 299–312, [hep-ph/0008184].

[188] J. P. Isaacson, ResBos2: Precision Resummation for the LHC Era. PhD thesis, Michigan State U., 2017. 10.25335/M5DG44.

[189] C. Balazs and C. P. Yuan, Soft gluon effects on lepton pairs at hadron colliders, Phys. Rev. D **56** (1997) 5558–5583, [hep-ph/9704258].

[190] F. Landry, R. Brock, P. M. Nadolsky and C. P. Yuan, Tevatron Run-1 Z boson data and Collins-Soper-Sterman resummation formalism, Phys. Rev. D **67** (2003) 073016, [hep-ph/0212159].

[191] CDF collaboration, T. Aaltonen et al., Limits on Anomalous Trilinear Gauge Couplings in $Z\gamma$ Events from $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. Lett. **107** (2011) 051802, [1103.2990].

[192] J. Isaacson, Y. Fu and C. P. Yuan, ResBos2 and the CDF W Mass Measurement, 2205.02788.

[193] Z.-B. Kang, A. Prokudin, N. Sato and J. Terry, Efficient Fourier Transforms for Transverse Momentum Dependent Distributions, Comput. Phys. Commun. **258** (2021) 107611, [1906.05949].

[194] W. Bizoń, E. Re and G. Zanderighi, NNLOPS description of the $H \to b\bar{b}$ decay with MiNLO, JHEP **06** (2020) 006, [1912.09982].

[195] K. Hamilton, T. Melia, P. F. Monni, E. Re and G. Zanderighi, Merging WW and WW+jet with MINLO, JHEP **09** (2016) 057, [1606.07062].

[196] S. Höche, Y. Li and S. Prestel, Drell-Yan lepton pair production at NNLO QCD with parton showers, Phys. Rev. D **91** (2015) 074015, [1405.3607].

[197] S. Höche, Y. Li and S. Prestel, Higgs-boson production through gluon fusion at NNLO QCD with parton showers, Phys. Rev. D **90** (2014) 054011, [1407.3773].

[198] V. Bertone and S. Prestel, Combining N3LO QCD calculations and parton showers for hadronic collision events, 2202.01082.

[199] S. Alioli, C. W. Bauer, C. Berggren, F. J. Tackmann and J. R. Walsh, Drell-Yan production at NNLL'+NNLO matched to parton showers, Phys. Rev. D **92** (2015) 094020, [1508.01475].

[200] S. Alioli, A. Broggio, S. Kallweit, M. A. Lim and L. Rottoli, Higgsstrahlung at NNLL'+NNLO matched to parton showers in GENEVA, Phys. Rev. D **100** (2019) 096016, [1909.02026].

[201] S. Alioli, A. Broggio, A. Gavardi, S. Kallweit, M. A. Lim, R. Nagar, D. Napoletano and L. Rottoli, Precise predictions for photon pair production matched to parton showers in GENEVA, JHEP **04** (2021) 041, [2010.10498].

[202] S. Alioli, A. Broggio, A. Gavardi, S. Kallweit, M. A. Lim, R. Nagar, D. Napoletano and L. Rottoli, Resummed predictions for hadronic Higgs boson decays, JHEP **04** (2021) 254, [2009.13533].

[203] S. Alioli, A. Broggio, A. Gavardi, S. Kallweit, M. A. Lim, R. Nagar and D. Napoletano, Next-to-next-to-leading order event generation for Z boson pair production matched to parton shower, Phys. Lett. B **818** (2021) 136380, [2103.01214].

[204] T. Cridge, M. A. Lim and R. Nagar, Wγ production at NNLO+PS accuracy in Geneva, Phys. Lett. B **826** (2022) 136918, [2105.13214].

[205] S. Alioli, C. W. Bauer, A. Broggio, A. Gavardi, S. Kallweit, M. A. Lim, R. Nagar, D. Napoletano and L. Rottoli, Matching NNLO predictions to parton showers using N3LL color-singlet transverse momentum resummation in geneva, Phys. Rev. D **104** (2021) 094020, [2102.08390].

[206] J. M. Campbell, S. Höche, H. T. Li, C. T. Preuss and P. Skands, Towards NNLO+PS Matching with Sector Showers, 2108.07133.

[207] K. Hamilton, P. Nason, E. Re and G. Zanderighi, NNLOPS simulation of Higgs boson production, JHEP **10** (2013) 222, [1309.0017].

[208] A. Karlberg, E. Re and G. Zanderighi, NNLOPS accurate Drell-Yan production, JHEP **09** (2014) 134, [1407.2940].

[209] W. Astill, W. Bizon, E. Re and G. Zanderighi, NNLOPS accurate associated HW production, JHEP **06** (2016) 154, [1603.01620].

[210] W. Astill, W. Bizoń, E. Re and G. Zanderighi, NNLOPS accurate associated HZ production with $H \to b\bar{b}$ decay at NLO, JHEP **11** (2018) 157, [1804.08141].

[211] E. Re, M. Wiesemann and G. Zanderighi, NNLOPS accurate predictions for W^+W^- production, JHEP **12** (2018) 121, [1805.09857].

[212] J. C. Collins and D. E. Soper, Angular Distribution of Dileptons in High-Energy Hadron Collisions, Phys. Rev. D **16** (1977) 2219.

[213] P. F. Monni, P. Nason, E. Re, M. Wiesemann and G. Zanderighi, MiNNLO$_{PS}$: a new method to match NNLO QCD to parton showers, JHEP **05** (2020) 143, [1908.06987].

[214] P. F. Monni, E. Re and M. Wiesemann, MiNNLO$_{\rm PS}$: optimizing $2 \to 1$ hadronic processes, Eur. Phys. J. C **80** (2020) 1075, [2006.04133].

[215] S. Zanoli, M. Chiesa, E. Re, M. Wiesemann and G. Zanderighi, Next-to-next-to-leading order event generation for VH production with $H \to b\bar{b}$ decay, JHEP **07** (2022) 008, [2112.04168].

[216] U. Haisch, D. J. Scott, M. Wiesemann, G. Zanderighi and S. Zanoli, NNLO event generation for $pp \to Zh \to \ell^+\ell^- b\bar{b}$ production in the SM effective field theory, 2204.00663.

[217] D. Lombardi, M. Wiesemann and G. Zanderighi, Advancing MiNNLO$_{\rm PS}$ diboson processes: Zγ production at NNLO+PS, JHEP **06** (2021) 095, [2010.10478].

[218] D. Lombardi, M. Wiesemann and G. Zanderighi, Anomalous couplings in Zγ events at NNLO+PS and improving $\nu\bar{\nu}\gamma$ backgrounds in dark-matter searches, Phys. Lett. B **824** (2022) 136846, [2108.11315].

[219] D. Lombardi, M. Wiesemann and G. Zanderighi, W^+W^- production at NNLO+PS with MINNLO$_{PS}$, JHEP **11** (2021) 230, [2103.12077].

[220] L. Buonocore, G. Koole, D. Lombardi, L. Rottoli, M. Wiesemann and G. Zanderighi, ZZ production at nNNLO+PS with MiNNLO$_{PS}$, JHEP **01** (2022) 072, [2108.05337].

[221] A. Gavardi, C. Oleari and E. Re, NNLO+PS Monte Carlo simulation of photon pair production with MiNNLOPS, 2204.12602.

[222] J. Mazzitelli, P. F. Monni, P. Nason, E. Re, M. Wiesemann and G. Zanderighi, Next-to-Next-to-Leading Order Event Generation for Top-Quark Pair Production, Phys. Rev. Lett. **127** (2021) 062001, [2012.14267].

[223] J. Mazzitelli, P. F. Monni, P. Nason, E. Re, M. Wiesemann and G. Zanderighi, Top-pair production at the LHC with MINNLO$_{PS}$, JHEP **04** (2022) 079, [2112.12135].

[224] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht, M. Schönherr and G. Watt, LHAPDF6: parton density access in the LHC precision era, Eur. Phys. J. C **75** (2015) 132, [1412.7420].

[225] G. P. Salam and J. Rojo, A Higher Order Perturbative Parton Evolution Toolkit (HOPPET), Comput. Phys. Commun. **180** (2009) 120–156, [0804.3755].

[226] B. Cabouat and T. Sjöstrand, Some Dipole Shower Studies, Eur. Phys. J. C **78** (2018) 226, [1710.00391].

[227] CMS collaboration, S. Chatrchyan et al., Measurement of $W\gamma$ and $Z\gamma$ production in pp collisions at $\sqrt{s} = 7$ TeV, Phys. Lett. B **701** (2011) 535–555, [1105.2758].

[228] ATLAS collaboration, G. Aad et al., Measurement of Wgamma and Zgamma production in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS Detector, JHEP **09** (2011) 072, [1106.1592].

[229] ATLAS collaboration, G. Aad et al., Measurement of $W\gamma$ and $Z\gamma$ production cross sections in pp collisions at $\sqrt{s} = 7$ TeV and limits on anomalous triple gauge couplings with the ATLAS detector, Phys. Lett. B **717** (2012) 49–69, [1205.2531].

[230] CMS collaboration, S. Chatrchyan et al., Measurement of the $W\gamma$ and $Z\gamma$ Inclusive Cross Sections in pp Collisions at $\sqrt{s} = 7$ TeV and Limits on Anomalous Triple Gauge Boson Couplings, Phys. Rev. D **89** (2014) 092005, [1308.6832].

[231] CMS collaboration, S. Chatrchyan et al., Measurement of the Production Cross Section for $Z\gamma \to \nu\bar{\nu}\gamma$ in pp Collisions at $\sqrt{s} = 7$ TeV and Limits on $ZZ\gamma$ and $Z\gamma\gamma$ Triple Gauge Boson Couplings, JHEP **10** (2013) 164, [1309.1117].

[232] ATLAS collaboration, G. Aad et al., Measurements of $W\gamma$ and $Z\gamma$ production in pp collisions at $\sqrt{s}$=7 TeV with the ATLAS detector at the LHC, Phys. Rev. D **87** (2013) 112003, [1302.1283].

[233] ATLAS collaboration, G. Aad et al., Search for new resonances in $W\gamma$ and $Z\gamma$ final states in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, Phys. Lett. B **738** (2014) 428–447, [1407.8150].

[234] CMS collaboration, V. Khachatryan et al., Measurement of the Zγ Production Cross Section in pp Collisions at 8 TeV and Search for Anomalous Triple Gauge Boson Couplings, JHEP **04** (2015) 164, [1502.05664].

[235] CMS collaboration, V. Khachatryan et al., Measurement of the Z$\gamma \to \nu\bar{\nu}\gamma$ production cross section in pp collisions at $\sqrt{s} = 8$ TeV and limits on anomalous ZZγ and Z$\gamma\gamma$ trilinear gauge boson couplings, Phys. Lett. B **760** (2016) 448–468, [1602.07152].

[236] ATLAS collaboration, G. Aad et al., Measurements of $Z\gamma$ and $Z\gamma\gamma$ production in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, Phys. Rev. D **93** (2016) 112002, [1604.05232].

[237] ATLAS collaboration, M. Aaboud et al., Measurement of the $Z\gamma \to \nu\bar{\nu}\gamma$ production cross section in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector and limits on anomalous triple gauge-boson couplings, JHEP **12** (2018) 010, [1810.04995].

[238] ATLAS collaboration, G. Aad et al., Measurement of the $Z(\to \ell^+\ell^-)\gamma$ production

cross-section in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, JHEP **03** (2020) 054, [1911.04813].

[239] CMS collaboration, A. M. Sirunyan et al., Search for the decay of a Higgs boson in the $\ell\ell\gamma$ channel in proton-proton collisions at $\sqrt{s} = 13$ TeV, JHEP **11** (2018) 152, [1806.05996].

[240] ATLAS collaboration, G. Aad et al., A search for the $Z\gamma$ decay mode of the Higgs boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Lett. B **809** (2020) 135754, [2005.05382].

[241] T. Becher and T. Neumann, Fiducial q_T resummation of color-singlet processes at $N^3LL+NNLO$, JHEP **03** (2021) 199, [2009.11437].

[242] J. Campbell and T. Neumann, Precision Phenomenology with MCFM, JHEP **12** (2019) 034, [1909.09117].

[243] S. Pozzorini, N. Schär and M. F. Zoller, Two-loop tensor integral coefficients in OpenLoops, JHEP **05** (2022) 161, [2201.11615].

[244] S. Pozzorini, N. Schär and M. F. Zoller, Towards two-loop automation in OpenLoops, 7, 2022. 2207.07468.

[245] F. Cascioli, P. Maierhöfer and S. Pozzorini, Scattering Amplitudes with Open Loops, Phys. Rev. Lett. **108** (2012) 111601, [1111.5206].

[246] F. Buccioni, S. Pozzorini and M. Zoller, On-the-fly reduction of open loops, Eur. Phys. J. **C78** (2018) 70, [1710.11452].

[247] F. Buccioni, J.-N. Lang, J. M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang and M. F. Zoller, OpenLoops 2, Eur. Phys. J. C **79** (2019) 866, [1907.13071].

[248] L. J. Dixon, Z. Kunszt and A. Signer, Helicity amplitudes for $O(\alpha_s)$ production of W^+W^-, $W^\pm Z$, ZZ, $W^\pm\gamma$, or $Z\gamma$ pairs at hadron colliders, Nucl. Phys. **B531** (1998) 3–23, [hep-ph/9803250].

[249] D. De Florian and A. Signer, W gamma and Z gamma production at hadron colliders, Eur. Phys. J. C **16** (2000) 105–114, [hep-ph/0002138].

[250] J. M. Campbell, H. B. Hartanto and C. Williams, Next-to-leading order predictions for $Z\gamma$+jet and Z $\gamma\gamma$ final states at the LHC, JHEP **11** (2012) 162, [1208.0566].

[251] J. M. Campbell, T. Neumann and C. Williams, $Z\gamma$ Production at NNLO Including Anomalous Couplings, JHEP **11** (2017) 150, [1708.02925].

[252] J. M. Campbell and R. Ellis, Top-quark loop corrections in Z+jet and Z + 2 jet production, JHEP **01** (2017) 020, [1610.02189].

[253] W. Beenakker, G. J. van Oldenborgh, A. Denner, S. Dittmaier, J. Hoogland, R. Kleiss, C. G. Papadopoulos and G. Passarino, The Fermion loop scheme for finite width effects in e^+e^- annihilation into four fermions, Nucl. Phys. B **500** (1997) 255–298, [hep-ph/9612260].

[254] A. Denner, S. Dittmaier, M. Roth and D. Wackeroth, Predictions for all processes e+ e- —> 4 fermions + gamma, Nucl. Phys. B **560** (1999) 33–65, [hep-ph/9904472].

[255] S. Frixione, Isolated photons in perturbative QCD, Phys. Lett. B **429** (1998) 369–374, [hep-ph/9801442].

[256] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, Diphoton production at the LHC: a QCD study up to NNLO, JHEP **04** (2018) 142, [1802.02095].

[257] L. Barze, M. Chiesa, G. Montagna, P. Nason, O. Nicrosini, F. Piccinini and V. Prosperi, Wγ production in hadronic collisions using the POWHEG+MiNLO method, JHEP **12** (2014) 039, [1408.5766].

[258] T. Jezo, M. Klasen and F. König, Prompt photon production and photon-hadron jet correlations with POWHEG, JHEP **11** (2016) 033, [1610.02275].

[259] M. Klasen, C. Klein-Bösing and H. Poppenborg, Prompt photon production and photon-jet correlations at the LHC, JHEP **03** (2018) 081, [1709.04154].

[260] S. Alioli, P. Nason, C. Oleari and E. Re, NLO vector-boson production matched with shower in POWHEG, JHEP **07** (2008) 060, [0805.4802].

[261] T. Gehrmann and L. Tancredi, Two-loop QCD helicity amplitudes for $q\bar{q} \to W^{\pm}\gamma$ and $q\bar{q} \to Z^0\gamma$, JHEP **02** (2012) 004, [1112.1531].

[262] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, Universality of transverse-momentum resummation and hard factors at the NNLO, Nucl. Phys. **B881** (2014) 414–443, [1311.1654].

[263] P. Nason and C. Oleari, Generation cuts and Born suppression in POWHEG, 1303.3922.

[264] NNPDF collaboration, R. D. Ball et al., Parton distributions for the LHC Run II, JHEP **04** (2015) 040, [1410.8849].

[265] T. Gehrmann and E. Remiddi, Numerical evaluation of harmonic polylogarithms, Comput. Phys. Commun. **141** (2001) 296–312, [hep-ph/0107173].

[266] ATLAS Pythia 8 tunes to 7 TeV datas, ATL-PHYS-PUB-2014-021, 11, 2014.

[267] M. Grazzini, S. Kallweit and D. Rathlev, $W\gamma$ and $Z\gamma$ production at the LHC in NNLO QCD, JHEP **07** (2015) 085, [1504.01330].

[268] M. Cacciari, G. P. Salam and G. Soyez, The anti-k_t jet clustering algorithm, JHEP **04** (2008) 063, [0802.1189].

[269] M. Cacciari, G. P. Salam and G. Soyez, FastJet User Manual, Eur. Phys. J. C **72** (2012) 1896, [1111.6097].

[270] A. Denner, S. Dittmaier, M. Hecht and C. Pasold, NLO QCD and electroweak corrections to $Z + \gamma$ production with leptonic Z-boson decays, JHEP **02** (2016) 057, [1510.08742].

[271] S. Catani and B. R. Webber, Infrared safe but infinite: Soft gluon divergences inside the physical region, JHEP **10** (1997) 005, [hep-ph/9710333].

[272] S. Kallweit, E. Re, L. Rottoli and M. Wiesemann, Accurate single- and double-differential resummation of colour-singlet processes with MATRIX+RADISH: W^+W^- production at the LHC, JHEP **12** (2020) 147, [2004.07720].

[273] M. Wiesemann, L. Rottoli and P. Torrielli, The Zγ transverse-momentum spectrum at NNLO+N^{3}LL, Phys. Lett. B **809** (2020) 135718, [2006.09338].

[274] P. F. Monni, E. Re and P. Torrielli, Higgs transverse-momentum resummation in direct space, Phys. Rev. Lett. **116** (2016) 242001, [1604.02191].

[275] P. F. Monni, L. Rottoli and P. Torrielli, Higgs transverse momentum with a jet veto: a double-differential resummation, vol. 124. 2020, 10.1103/PhysRevLett.124.252001.

[276] ATLAS collaboration, G. Aad et al., Search for dark matter in association with an energetic photon in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, JHEP **02** (2021) 226, [2011.05259].

[277] U. Baur and E. L. Berger, Probing the weak boson sector in $Z\gamma$ production at hadron colliders, Phys. Rev. D **47** (1993) 4889–4904.

[278] K. Hagiwara, R. D. Peccei, D. Zeppenfeld and K. Hikasa, Probing the Weak Boson Sector in e+ e- —> W+ W-, Nucl. Phys. B **282** (1987) 253–307.

[279] G. J. Gounaris, J. Layssac and F. M. Renard, Signatures of the anomalous Z_γ and ZZ production at the lepton and hadron colliders, Phys. Rev. D **61** (2000) 073013, [hep-ph/9910395].

[280] C. Degrande, N. Greiner, W. Kilian, O. Mattelaer, H. Mebane, T. Stelzer, S. Willenbrock and C. Zhang, Effective Field Theory: A Modern Approach to Anomalous Couplings, Annals Phys. **335** (2013) 21–32, [1205.4231].

[281] C. Degrande, A basis of dimension-eight operators for anomalous neutral triple gauge boson interactions, JHEP **02** (2014) 101, [1308.6323].

[282] ATLAS collaboration, M. Aaboud et al., Measurement of the $Z\gamma \to \nu\bar{\nu}\gamma$ production cross section in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector and limits on anomalous triple gauge-boson couplings, JHEP **12** (2018) 010, [1810.04995].

[283] J. Ellis, H.-J. He and R.-Q. Xiao, Probing Neutral Triple Gauge Couplings at the LHC and Future Hadron Colliders, 2206.11676.

[284] G. J. Gounaris, J. Layssac and F. M. Renard, Off-shell structure of the anomalous Z and γ selfcouplings, Phys. Rev. D **62** (2000) 073012, [hep-ph/0005269].

[285] G. J. Gounaris, J. Layssac and F. M. Renard, New and standard physics contributions to anomalous Z and gamma selfcouplings, Phys. Rev. D **62** (2000) 073013, [hep-ph/0003143].

[286] T. Gehrmann, E. W. N. Glover, T. Huber, N. Ikizlerli and C. Studerus, Calculation of the quark and gluon form factors to three loops in QCD, JHEP **06** (2010) 094, [1004.3653].

[287] P. A. Baikov, K. G. Chetyrkin, A. V. Smirnov, V. A. Smirnov and M. Steinhauser, Quark and gluon form factors to three loops, Phys. Rev. Lett. **102** (2009) 212002, [0902.3519].

[288] R. N. Lee, A. V. Smirnov and V. A. Smirnov, Analytic Results for Massless Three-Loop Form Factors, JHEP **04** (2010) 020, [1001.2887].

[289] ALEPH, DELPHI, L3, OPAL, LEP Electroweak collaboration, S. Schael et al., Electroweak Measurements in Electron-Positron Collisions at W-Boson-Pair Energies at LEP, Phys. Rept. **532** (2013) 119–244, [1302.3415].

[290] D0 collaboration, V. M. Abazov et al., $Z\gamma$ production and limits on anomalous $ZZ\gamma$ and $Z\gamma\gamma$ couplings in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, Phys. Rev. D **85** (2012) 052001, [1111.3684].

[291] D. R. Green, P. Meade and M.-A. Pleier, Multiboson interactions at the LHC, Rev. Mod. Phys. **89** (2017) 035008, [1610.07572].

[292] ATLAS collaboration, M. Aaboud et al., Search for dark matter at $\sqrt{s} = 13$ TeV in final states containing an energetic photon and large missing transverse momentum with the ATLAS detector, Eur. Phys. J. C **77** (2017) 393, [1704.03848].

[293] CMS collaboration, A. M. Sirunyan et al., Search for new physics in final states with a single photon and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV, JHEP **02** (2019) 074, [1810.00196].

[294] ATLAS collaboration, M. Aaboud et al., Search for dark matter and other new

phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector, JHEP **01** (2018) 126, [1711.03301].

[295] CMS collaboration, A. M. Sirunyan et al., Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV, Phys. Rev. D **97** (2018) 092005, [1712.02345].

[296] ATLAS collaboration, M. Aaboud et al., Search for dark matter produced in association with bottom or top quarks in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, Eur. Phys. J. C **78** (2018) 18, [1710.11412].

[297] CMS collaboration, A. M. Sirunyan et al., Search for dark matter in events with energetic, hadronically decaying top quarks and missing transverse momentum at $\sqrt{s} = 13$ TeV, JHEP **06** (2018) 027, [1801.08427].

[298] CMS collaboration, A. M. Sirunyan et al., Search for dark matter and unparticles in events with a Z boson and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV, JHEP **03** (2017) 061, [1701.02042].

[299] ATLAS collaboration, M. Aaboud et al., Search for dark matter in events with a hadronically decaying vector boson and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, JHEP **10** (2018) 180, [1807.11471].

[300] ATLAS collaboration, M. Aaboud et al., Search for dark matter in association with a Higgs boson decaying to two photons at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Rev. D **96** (2017) 112004, [1706.03948].

[301] CMS collaboration, A. M. Sirunyan et al., Search for associated production of dark matter with a Higgs boson decaying to $b\bar{b}$ or $\gamma\gamma$ at $\sqrt{s} = 13$ TeV, JHEP **10** (2017) 180, [1703.05236].

[302] J. Abdallah et al., Simplified Models for Dark Matter Searches at the LHC, Phys. Dark Univ. **9-10** (2015) 8–23, [1506.03116].

[303] O. Buchmueller, M. J. Dolan, S. A. Malik and C. McCabe, Characterising dark matter searches at colliders and direct detection experiments: Vector mediators, JHEP **01** (2015) 037, [1407.8257].

[304] D. Abercrombie et al., Dark Matter benchmark models for early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum, Phys. Dark Univ. **27** (2020) 100371, [1507.00966].

[305] I. Brivio, M. B. Gavela, L. Merlo, K. Mimasu, J. M. No, R. del Rey and V. Sanz, ALPs Effective Field Theory and Collider Signatures, Eur. Phys. J. C **77** (2017) 572, [1701.05379].

[306] Sherpa collaboration, E. Bothmann et al., Event Generation with Sherpa 2.2, SciPost Phys. **7** (2019) 034, [1905.09127].

[307] T. Melia, P. Nason, R. Röntsch and G. Zanderighi, W+W-, WZ and ZZ production in the POWHEG BOX, JHEP **11** (2011) 078, [1107.5051].

[308] S. Kallweit and M. Wiesemann, ZZ production at the LHC: NNLO predictions for $2\ell 2\nu$ and 4ℓ signatures, Phys. Lett. **B786** (2018) 382–389, [1806.05941].

[309] S. Kallweit, J. Lindert, S. Pozzorini and M. Schönherr, NLO QCD+EW predictions for $2\ell 2\nu$ diboson signatures at the LHC, JHEP **11** (2017) 120, [1705.00598].

[310] T. Gehrmann, M. Grazzini, S. Kallweit, P. Maierhöfer, A. von Manteuffel, S. Pozzorini,

D. Rathlev and L. Tancredi, W^+W^- Production at Hadron Colliders in Next to Next to Leading Order QCD, Phys. Rev. Lett. **113** (2014) 212001, [1408.5243].

[311] M. Grazzini, S. Kallweit, S. Pozzorini, D. Rathlev and M. Wiesemann, W^+W^- production at the LHC: fiducial cross sections and distributions in NNLO QCD, JHEP **08** (2016) 140, [1605.02716].

[312] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn, D. L. Rainwater and T. Stelzer, MadGraph/MadEvent v4: The New Web Generation, JHEP **09** (2007) 028, [0706.2334].

[313] J. M. Campbell, R. Ellis, R. Frederix, P. Nason, C. Oleari and C. Williams, NLO Higgs Boson Production Plus One and Two Jets Using the POWHEG BOX, MadGraph4 and MCFM, JHEP **07** (2012) 092, [1202.5475].

[314] G. Cullen et al., G*OSAM*-2.0: a tool for automated one-loop calculations within the Standard Model and beyond, Eur. Phys. J. C **74** (2014) 3001, [1404.7096].

[315] T. Gehrmann, A. von Manteuffel and L. Tancredi, The two-loop helicity amplitudes for $q\bar{q}' \to V_1V_2 \to 4$ leptons, JHEP **09** (2015) 128, [1503.04812].

[316] The VVamp project, by T. Gehrmann, A. von Manteuffel, and L. Tancredi, is publicly available, http://vvamp.hepforge.org.

[317] L. J. Dixon, A brief introduction to modern amplitude methods, in Theoretical Advanced Study Institute in Elementary Particle Physics: Particle Physics: The Higgs Boson and Beyond, pp. 31–67, 2014. 1310.5353. DOI.

[318] T. Gehrmann, T. Huber and D. Maitre, Two-loop quark and gluon form-factors in dimensional regularisation, Phys. Lett. **B622** (2005) 295–302, [hep-ph/0507061].

[319] R. J. Gonsalves, Dimensionally Regularized Two Loop On-shell Quark Form Factor, Phys. Rev. D **28** (1983) 1542.

[320] W. van Neerven, Dimensional Regularization of Mass and Infrared Singularities in Two Loop On-shell Vertex Functions, Nucl. Phys. B **268** (1986) 453–488.

[321] G. Kramer and B. Lampe, Integrals for Two Loop Calculations in Massless QCD, J. Math. Phys. **28** (1987) 945.

[322] The Btwxt general-purpose, N-dimensional interpolation library, by N. Kruis, T. Scimone and P. Sullivan, https://github.com/bigladder/btwxt.

[323] G. Birkhoff and H. L. Garabedian, Smooth surface interpolation, Journal of Mathematics and Physics **39** (1960) 258–268.

[324] E. Catmull and R. Rom, A class of local interpolating splines, in Computer Aided Geometric Design (R. E. Barnhill and R. F. Riesenfiled, eds.), pp. 317 – 326. Academic Press, 1974. DOI.

[325] Particle Data Group collaboration, C. Patrignani et al., Review of Particle Physics, Chin. Phys. **C40** (2016) 100001.

[326] Matrix+RadISH is an interface to RadISH within Matrix by S. Kallweit, E. Re, L. Rottoli, M. Wiesemann, https://matrix.hepforge.org/matrix+radish.html.

[327] P. Skands, S. Carrazza and J. Rojo, Tuning PYTHIA 8.1: the Monash 2013 Tune, Eur. Phys. J. C **74** (2014) 3024, [1404.5630].

[328] ATLAS collaboration, M. Aaboud et al., Measurement of the W^+W^- production cross

section in pp collisions at a centre-of-mass energy of $\sqrt{s} = 13$ TeV with the ATLAS experiment, Phys. Lett. **B773** (2017) 354–374, [1702.04519].

[329] ATLAS collaboration, M. Aaboud et al., Measurement of fiducial and differential W^+W^- production cross-sections at $\sqrt{s} = 13$ TeV with the ATLAS detector, Eur. Phys. J. **C79** (2019) 884, [1905.04242].

[330] CMS collaboration, Measurement of the WW cross section pp collisions at sqrt(s)=13 TeV, https://cds.cern.ch/record/2160868.

[331] CMS collaboration, A. M. Sirunyan et al., W^+W^- boson pair production in proton-proton collisions at $\sqrt{s} = 13$ TeV, Phys. Rev. D **102** (2020) 092001, [2009.00119].

[332] M. Grazzini, S. Kallweit, M. Wiesemann and J. Y. Yook, W^+W^- production at the LHC: NLO QCD corrections to the loop-induced gluon fusion channel, Phys. Lett. B **804** (2020) 135399, [2002.01877].

[333] I. W. Stewart and F. J. Tackmann, Theory Uncertainties for Higgs and Other Searches Using Jet Bins, Phys. Rev. D **85** (2012) 034011, [1107.2117].

[334] A. Banfi, P. F. Monni, G. P. Salam and G. Zanderighi, Higgs and Z-boson production with a jet veto, Phys. Rev. Lett. **109** (2012) 202001, [1206.4998].

[335] A. J. Barr, Measuring slepton spin at the LHC, JHEP **02** (2006) 042, [hep-ph/0511115].

[336] CMS collaboration, V. Khachatryan et al., Measurement of the ZZ production cross section and $Z \to \ell^+\ell^-\ell'^+\ell'^-$ branching fraction in pp collisions at $\sqrt{s}$=13 TeV, Phys. Lett. B **763** (2016) 280–303, [1607.08834].

[337] ATLAS collaboration, M. Aaboud et al., $ZZ \to \ell^+\ell^-\ell'^+\ell'^-$ cross-section measurements and search for anomalous triple gauge couplings in 13 TeV pp collisions with the ATLAS detector, Phys. Rev. **D97** (2018) 032005, [1709.07703].

[338] CMS collaboration, A. M. Sirunyan et al., Measurements of the pp $\to$ ZZ production cross section and the $Z \to 4\ell$ branching fraction, and constraints on anomalous triple gauge couplings at $\sqrt{s} = 13$ TeV, Eur. Phys. J. C **78** (2018) 165, [1709.08601].

[339] CMS collaboration, A. M. Sirunyan et al., Measurement of differential cross sections for Z boson pair production in association with jets at $\sqrt{s} = 8$ and 13 TeV, Phys. Lett. B **789** (2019) 19–44, [1806.11073].

[340] ATLAS collaboration, M. Aaboud et al., Measurement of ZZ production in the $\ell\ell\nu\nu$ final state with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV, JHEP **10** (2019) 127, [1905.07163].

[341] CMS collaboration, A. M. Sirunyan et al., Measurements of pp $\to$ ZZ production cross sections and constraints on anomalous triple gauge couplings at $\sqrt{s} = 13$ TeV, Eur. Phys. J. C **81** (2021) 200, [2009.01186].

[342] ATLAS collaboration, M. Aaboud et al., Constraints on off-shell Higgs boson production and the Higgs boson total width in $ZZ \to 4\ell$ and $ZZ \to 2\ell 2\nu$ final states with the ATLAS detector, Phys. Lett. **B786** (2018) 223–244, [1808.01191].

[343] CMS collaboration, A. M. Sirunyan et al., Measurements of the Higgs boson width and anomalous HVV couplings from on-shell and off-shell production in the four-lepton final state, Phys. Rev. D **99** (2019) 112003, [1901.00174].

[344] U. Haisch and G. Koole, Probing Higgs portals with matrix-element based kinematic discriminants in $ZZ \to 4\ell$ production, JHEP **04** (2022) 166, [2201.09711].

[345] M. Grazzini, S. Kallweit and D. Rathlev, ZZ production at the LHC: fiducial cross sections and distributions in NNLO QCD, Phys. Lett. **B750** (2015) 407–410, [1507.06257].

[346] F. Caola, K. Melnikov, R. Röntsch and L. Tancredi, QCD corrections to ZZ production in gluon fusion at the LHC, Phys. Rev. **D92** (2015) 094028, [1509.06734].

[347] M. Grazzini, S. Kallweit, M. Wiesemann and J. Y. Yook, ZZ production at the LHC: NLO QCD corrections to the loop-induced gluon fusion channel, JHEP **03** (2019) 070, [1811.09593].

[348] A. Denner, J.-N. Lang and S. Uccirati, NLO electroweak corrections in extended Higgs Sectors with RECOLA2, JHEP **07** (2017) 087, [1705.06053].

[349] A. Denner, J.-N. Lang and S. Uccirati, Recola2: REcursive Computation of One-Loop Amplitudes 2, Comput. Phys. Commun. **224** (2018) 346–361, [1711.07388].

[350] A. von Manteuffel and L. Tancredi, The two-loop helicity amplitudes for $gg \to V_1V_2 \to 4$ leptons, JHEP **1506** (2015) 197, [1503.08835].

[351] B. Agarwal, S. P. Jones and A. von Manteuffel, Two-loop helicity amplitudes for $gg \to ZZ$ with full top-quark mass effects, JHEP **05** (2021) 256, [2011.15113].

[352] C. Brønnum-Hansen and C.-Y. Wang, Top quark contribution to two-loop helicity amplitudes for Z boson pair production in gluon fusion, JHEP **05** (2021) 244, [2101.12095].

[353] S. Alioli, S. Ferrario Ravasio, J. M. Lindert and R. Röntsch, Four-lepton production in gluon fusion at NLO matched to parton showers, Eur. Phys. J. C **81** (2021) 687, [2102.07783].

[354] F. Cascioli, S. Höche, F. Krauss, P. Maierhöfer, S. Pozzorini and F. Siegert, Precise Higgs-background predictions: merging NLO QCD and squared quark-loop corrections to four-lepton + 0,1 jet production, JHEP **01** (2014) 046, [1309.0500].

[355] J. M. Campbell, R. K. Ellis, E. Furlan and R. Röntsch, Interference effects for Higgs boson mediated Z-pair plus jet production, Phys. Rev. D **90** (2014) 093008, [1409.1897].

[356] C. Li, Y. An, C. Charlot, R. Covarelli, Z. Guan and Q. Li, Loop-induced ZZ production at the LHC: An improved description by matrix-element matching, Phys. Rev. D **102** (2020) 116003, [2006.12860].

[357] NNPDF collaboration, R. D. Ball et al., Parton distributions from high-precision collider data, Eur. Phys. J. **C77** (2017) 663, [1706.00428].

[358] A. Manohar, P. Nason, G. P. Salam and G. Zanderighi, How bright is the proton? A precise determination of the photon parton distribution function, Phys. Rev. Lett. **117** (2016) 242002, [1607.04266].

[359] A. V. Manohar, P. Nason, G. P. Salam and G. Zanderighi, The Photon Content of the Proton, JHEP **12** (2017) 046, [1708.01256].

[360] NNPDF collaboration, V. Bertone, S. Carrazza, N. P. Hartland and J. Rojo, Illuminating the photon content of the proton within a global PDF analysis, SciPost Phys. **5** (2018) 008, [1712.07053].

[361] M. Grazzini, S. Kallweit, J. M. Lindert, S. Pozzorini and M. Wiesemann, NNLO QCD + NLO EW with Matrix+OpenLoops: precise predictions for vector-boson pair production, JHEP **02** (2020) 087, [1912.00068].

[362] M. Czakon, D. Heymes and A. Mitov, Dynamical scales for multi-TeV top-pair production at the LHC, JHEP **04** (2017) 071, [1606.03350].

[363] F. Caola, F. A. Dreyer, R. W. McDonald and G. P. Salam, Framing energetic top-quark pair production at the LHC, JHEP **07** (2021) 040, [2101.06068].

[364] S. Bräuer, A. Denner, M. Pellen, M. Schönherr and S. Schumann, Fixed-order and merged parton-shower predictions for WW and WWj production at the LHC including NLO QCD and EW corrections, JHEP **10** (2020) 159, [2005.12128].

[365] M. Chiesa, C. Oleari and E. Re, NLO QCD+NLO EW corrections to diboson production matched to parton shower, Eur. Phys. J. C **80** (2020) 849, [2005.12146].

[366] A. Denner and S. Dittmaier, Electroweak Radiative Corrections for Collider Physics, Phys. Rept. **864** (2020) 1–163, [1912.06823].

[367] C. W. Bauer, A. Frink and R. Kreckel, Introduction to the GiNaC framework for symbolic computation within the C++ programming language, J. Symb. Comput. **33** (2000) 1, [cs/0004015].

[368] B. Haible and R. B. Kreckel, CLN: Class Library for Numbers, https://www.ginac.de/CLN/.

List of Figures

List of Tables

Acknowledgements

This Thesis is the result of a four-year PhD, during which I have learnt a lot about physics, as well as about life and especially about myself. As a result, all people who have supported me through this journey would deserve to be mentioned in this list of acknowledgements ... but as soon as I started to write it down, I quickly realised it was becoming too long. Therefore, I apologise to everyone who finds a place in my heart, but not on this page.

First of all, I am extremely grateful to Giulia Zanderighi and Marius Wiesemann, since, if I have learnt anything about QCD and particle phenomenology, all the credit goes to them. In particular, I want to thank Giulia for having been much more than simply my supervisor in these difficult years which, among others, have been deeply marked by a wordlwide pandemic: I would like her to know I will not forget all her help and support. I also want to show special gratitude to Marius for having been an inspiring model to follow and for having always believed in me, sometimes even more than I did. But I would also like to thank Andreas Weiler and Ulrich Haisch for all important discussions and advices during my doctorate.

It is important to recognize that my entire PhD would have not been the same without the big family of the Max-Planck of Munich. I can not help remembering the nice months with my old office mates, Simone Zoia, Christoph Dlapa and Leila Maestri, but also with my new pheno-friends, Silvia Zanoli and Alessandro Ratti. Moreover, even if we shared the office just for few months, I also want to remember here the great time spent with Ryan Wood, Federico Silvetti and Matteo Tresoldi during my visit at Cern, in Geneva.

Last but not least, I want to thank my family and all of my friends, both the ones still in Italy and the ones found in Munich. In particular, I would like to explicitly mention my parents, Gianni and Rosa, that always act behind the scenes: nothing of what I did would have been possible without your presence. Thanks!